A Note on the Complexity of Real Algebraic Hypersurfaces

Michael Kerber[1], Michael Sagraloff[2]

[1] IST (Institute of Science and Technology) Austria, 3400 Klosterneuburg, Austria
[2] Max-Planck-Institute for Informatics, 66123 Saarbrücken, Germany

Abstract. Given an algebraic hypersurface $\mathcal{O}$ in $\mathbb{R}^d$, how many simplices are necessary for a simplicial complex isotopic to $\mathcal{O}$? We address this problem and the variant where all vertices of the complex must lie on $\mathcal{O}$. We give asymptotically tight worst-case bounds for algebraic plane curves. Our results gradually improve known bounds in higher dimensions; however, the question for tight bounds remains unsolved for $d \geq 3$.

Key words. Algebraic curves, algebraic surfaces, triangulation, isotopy

1. Introduction

A standard technique to process non-linear curves and surfaces in geometric systems is to approximate them in terms of a piecewise linear object (a simplicial complex). A main goal is to preserve the topological properties of the input. Furthermore, geometric properties, such as the position of singular or "extremal" points of the object are often of interest. For algebraic curves and surfaces as inputs, the former problem is usually called *topology computation*, the latter *topological-geometric analysis* of the object.[1]

We consider the following question: *How many simplices are needed to embed a simplicial complex in $\mathbb{R}^d$ that is isotopic to a real algebraic hypersurface in $\mathbb{R}^d$ of degree n?* Our main contribution is to provide sharp bounds for the planar case ($d = 2$): for a topologically correct representation, $\Omega(n^2)$ line segments are needed in the worst case, and we give an algorithm producing $O(n^2)$ line segments for all cases. Although the idea is simple, it seemingly does not appear in the literature yet. For geometric-topological representations, we construct a class of curves such that $\Omega(n^3)$ line segments are necessary. This proves that the *cylindrical algebraic decomposition* [8] ("Find the critical x-coordinates of the curve; compute the fiber at these coordinates and at separating points in between; connect the fiber points by straight-line segments." – compare Fig. 2) is asymptotically optimal. This is surprising because the vertical decomposition strategy seems to introduce much more line segments than actually necessary.

Our results can be partially generalized in higher dimensions. This allows a gradual improvement of lower and upper bounds that can be derived easily from cylindrical algebraic decomposition. Nevertheless, our bounds fail to be tight already for algebraic surfaces: For the topological approximation, we get a lower bound of $\Omega(n^3)$ and an upper bound of $O(n^5)$

[1] all terms will be formally defined in Section 2

triangles. For the geometric-topological approximation, the bounds are $\Omega(n^4)$ and $O(n^7)$, respectively. These gaps increase in higher dimensions because the lower bounds grow single exponentially in the dimension, whereas the upper bounds grow double exponentially.

Related work: Efficient techniques for topology computation of algebraic curves (e.g, see [9, 12], and references therein) and surfaces [4,1] have been presented in case where the defining polynomial f has integer coefficients. For the planar case, the complexity of the problem has been upper bounded by $O(N^{12})$ [11,14], where N is defined as the maximum of the degree of f and the bitsize of its coefficients. However, our question of how many segment/triangles are needed in principal to capture the topology of the object seems to be untreated in this context.

We remark that similar problems have been extensively studied for 2-manifolds. For instance, Nakamoto and Ota [16] show that any closed compact 2-manifold of genus g can be triangulated using $\Theta(g)$ vertices. An often discussed concept in this context is an *irreducible triangulation* of a 2-manifold, that is, a triangulation where no edge can be contracted without changing the topology. It has been shown that only finitely many irreducible triangulations exist [2], and they have been enumerated explicitly for the torus [15]. Although these results aim in a somewhat similar direction, algebraic surfaces are in general not 2-manifolds and need different techniques to be analyzed.

2. Basic notation and definitions

A *homeomorphism* between two sets $X,Y \subset \mathbb{R}^d$ is a bijective, continuous map $h : X \to Y$ whose inverse is continuous as well. X and Y are *isotopic* if they are "connected by homeomorphism", that is, there exists a continuous map $\psi : [0,1] \times X \to \mathbb{R}^d$ such that $\psi(0,\cdot) = \mathrm{id}_x$, $\psi(1,X) = Y$, and $\psi(t_0,x) : X \to \psi(t_0,X)$ is a homeomorphism for any $t_0 \in [0,1]$. ψ is called an *isotopy* between X and Y; see also [7,6] for more details. We assume that the reader is familiar with the definition of a simplicial complex. We only consider d-dimensional complexes that are embedded in $\mathbb{R}^d$ by fixing their vertices, and we identify the complex and the induced point set.

A *(real) algebraic hypersurface* $\mathscr{O}$ in $\mathbb{R}^d$ is the (real) solution set of an equation $f = 0$ with $f \in \mathbb{R}[x_1,\ldots,x_d]$. We also denote the real *vanishing set* of a polynomial f by $V(f) := \{x \in \mathbb{R}^d : f(x) = 0\}$. Hypersurfaces in dimensions 2 and 3 are called *algebraic curves* and *algebraic surfaces*, respectively. The *degree* of $\mathscr{O}$ is defined by the degree of f. An *isolated point* $p \in \mathbb{R}^d$ is a point on $\mathscr{O}$ such that an open neighborhood of p in $\mathbb{R}^d$ does not contain any further point of $\mathscr{O}$.

For a compact hypersurface $\mathscr{O} \subset \mathbb{R}^d$, we call an *isocomplex* of $\mathscr{O}$ to be a simplicial complex $S \subset \mathbb{R}^d$ that is isotopic to $\mathscr{O}$. We call a *stable isocomplex* to be an isocomplex that is stable at vertices, that is, there exists an isotopy ψ between $\mathscr{O}$ and S such that for each vertex v of S, $\psi(t,v) = v$ for any $t \in [0,1]$. Computing the topology of $\mathscr{O}$ means to compute an isocomplex, computing a geometric-topological analysis means to compute a stable isocomplex.

For unbounded hypersurfaces, one can define a (stable) isocomplex with respect to a compact region C to be a complex isotopic to $\mathscr{O} \cap C$. For simplicity, we restrict to the case of compact hypersurfaces in this work; however, the obtained bounds also hold for hypersurfaces restricted to any axis-aligned bounding box in $\mathbb{R}^d$.

3. Bounds for algebraic plane curves

3.1. Stable isocomplexes

Our main idea for deriving lower bounds is to construct algebraic hypersurfaces with many isolated points. We can even fix the location of each isolated point to a ball of arbitrary small radius.

Theorem 1. *For $d,n \in \mathbb{N}$, set $c := \binom{\lfloor n/2 \rfloor + d}{d} - d$. Then, for any $\varepsilon > 0$, and any set of points $p_1,\dots,p_c \in \mathbb{Q}^d$, there exists a hypersurface $C \subset \mathbb{R}^d$ of degree n such that for any p_i, C contains an isolated point $p'_i \in \mathbb{R}^d$ with $\|p_i - p'_i\|_2 < \varepsilon$.*

Proof. The idea is to construct d polynomials $f_1,\dots,f_d$ of degree $\lfloor n/2 \rfloor$ that all interpolate the points $p_1,\dots,p_c$, and to consider the hypersurface defined by $f := f_1^2 + \ldots + f_d^2$. Obviously, $\deg f \le n$, and $V(f) = V(f_1) \cap \ldots \cap V(f_d)$. If $V(f)$ is zero-dimensional (that means, contains only finitely many isolated points), the theorem is proven. However, for certain $p_1,\dots,p_c$, $V(f)$ is not zero-dimensional for any choice of interpolation polynomials $f_1,\dots,f_n$; the remainder of the proof argues that we can always use the described construction after a small perturbation of the initial points.

Firstly, almost all choices of d hypersurfaces $g_1,\dots,g_d$ in $\mathbb{R}^d$ (of degree $\lfloor n/2 \rfloor$) yield a zero-dimensional intersection: consider the coefficients of the polynomials as indeterminates, then the (multivariate) resultant R_g [10] with respect to any variable, say x_1, is a polynomial in x_1 that does not vanish identically (we write R_g because the resultant is parameterized in $g_1,\dots,g_d$).

We next consider c points $q_1,\dots,q_c$ in $\mathbb{C}^d$ with indeterminate coordinates. We force d hypersurfaces, each of degree $\lfloor n/2 \rfloor$, with indeterminate coefficients to pass through them. As a consequence, each coefficient can be re-expressed in dependency of the coordinates of the q_i, plus additional degrees of freedom. The same also holds true for the coefficient of the resultant polynomial R_q of these hyperplanes (we write R_q because the resultant is parametrized in the points $q_1,\dots,q_c$). We will show next that R_q does not vanish identically by showing that it does not vanish for at least one concrete choice of $q_1,\dots,q_c$.

The degree of R_g is $\lfloor n/2 \rfloor^d$. Choose d hypersurfaces $g_1,\dots,g_d$ such that the leading term of R_g does not vanish. Then, there exist (cf. [13] for a refined version of Bézout's Theorem) $\lfloor n/2 \rfloor^d$ intersection points in the projective space $\mathbb{P}(\mathbb{C}^d)$ (counted with multiplicities), and we can w.l.o.g. assume that all these points are distinct and lie in the affine space $\mathbb{C}^d$. It is a simple proof that $\lfloor n/2 \rfloor^d \ge c = \binom{\lfloor n/2 \rfloor + d}{d}$ for all $n,d \in \mathbb{N}$ (by induction on d). So, we can pick c of the common intersection points to take the role of the points $q_1,\dots,q_c$ from above, and set the other degrees of freedom such that we obtain $g_1,\dots,g_d$. With this choice, $R_q = R_g \ne 0$, thus, R_q defines a lower-dimensional variety in $\mathbb{C}^d$. It follows that R_q does not vanish for almost any choice of base points $q_1,\dots,q_c$.

Thus, for given points $p_1,\dots,p_c \in \mathbb{Q}^d$, we find points $p'_1,\dots,p'_c$ in an ε-ball around them such that there are hypersurfaces $f_1,\dots,f_d$ interpolating them and such that the resultant of $f_1,\dots,f_d$ does not vanish identically. It remains to argue that $p'_1,\dots,p'_c$ can be chosen with real coordinates, but this follows immediately, since otherwise, the resultant variety would contain an open ball of $\mathbb{R}^d$, and consequently, it would contain the whole $\mathbb{R}^d$, which is impossible.

For constant d, the theorem says that we can choose $\Theta(n^d)$ arbitrary rational points and construct an algebraic hypersurface of degree n with isolated points close to them.

Theorem 2. *There exists an algebraic curve $\mathscr{C} \subset \mathbb{R}^2$ of degree n such that any stable isocomplex for $\mathscr{C}$ has $\Omega(n^3)$ vertices.*

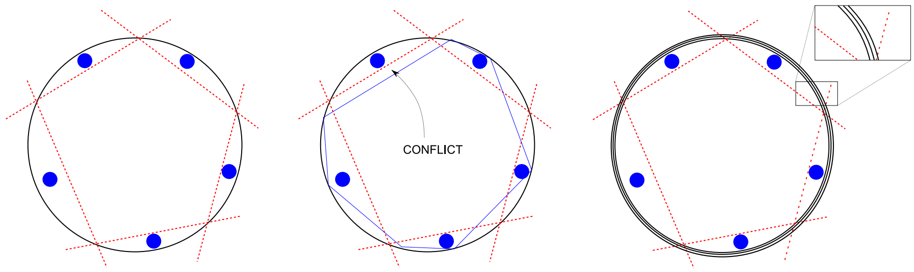

Fig. 1. Illustration of the construction in the proof of Theorem 2

Proof. We prove the claim by constructing a suitable curve $\mathscr{C}$. Assume first that the unit circle is a component of $\mathscr{C}$. Any isocomplex of $\mathscr{C}$ must contain a sequence of points on the unit circle which forms a cycle in the complex. We cut out $c' := \binom{\lfloor n/4 \rfloor + 2}{2} - 2$ disjoint regions of the unit disc by intersecting the disc with c' different lines. We place a disc of size ε in each of the regions and force an isolated point of the curve $\mathscr{C}$ to lie inside each disc (Fig. 1 left). By Theorem 1, this is possible if $\mathscr{C}$ is of degree at least $n/2$.

The isotopic cycle for the unit circle component contains a vertex in each of the regions: If there is no such vertex, the cycle misses the region completely, so the isolated point is outside the cycle, contradicting the properties of a stable isocomplex (Fig. 1 middle). Hence, at least $c' = \Omega(n^2)$ vertices are placed on the unit circle.

Finally, we take a collection of $n/4$ concentric circles to be part of $\mathscr{C}$ (instead of just the unit circle) such that the lines chosen as above still cut out c' disjoint regions for any of the circles (Fig. 1 right). This is clearly possible if all concentric circles have radius close enough to 1. The argument from above now works separately for each of the circles, thus, each one is divided into $\Omega(n^2)$ line segments under the isotopy.

To summarize, the final curve consists of two components: one curve of degree $n/2$ that forces the isolated singularities in the regions and a collection of $n/4$ circles (of total degree $n/2$). The union is of degree n, and any stable isocomplex requires $\Omega(n^2)$ vertices per circle, so $\Omega(n^3)$ vertices are required in total.

The upper bound of $O(n^3)$ vertices follows immediately from standard theory of algebraic curves and cylindrical algebraic decomposition [8].

Lemma 1. *For any algebraic curve $\mathscr{C} \subset \mathbb{R}^2$ of degree n, there exists a stable isocomplex with $O(n^3)$ cells.*

Proof. An algebraic curve $\mathscr{C}$ of degree n has up to $n(n-1)$ x-critical points p, that is, $f(p) = f_y(p) = 0$ by Bezout's Theorem [13]. The projections of these points decompose the x-axes into $O(n^2)$ delineable sets. This means that the fiber above each cell in the decomposition consists of finitely many (at most n) function graphs. Inserting points in between two consecutive projections and lifting each of the points in one dimension leads to a stable isocomplex of $\mathscr{C}$ with $O(n^3)$ points. See also Figure 2.

3.2. General Isocomplexes

We next remove the stability requirement on the isocomplex. Considering an arrangement of n lines in generic position, we observe that each pair intersect in a point. The union of n lines defines an algebraic curve of degree n with $\binom{n}{2}$ singularities. It follows:

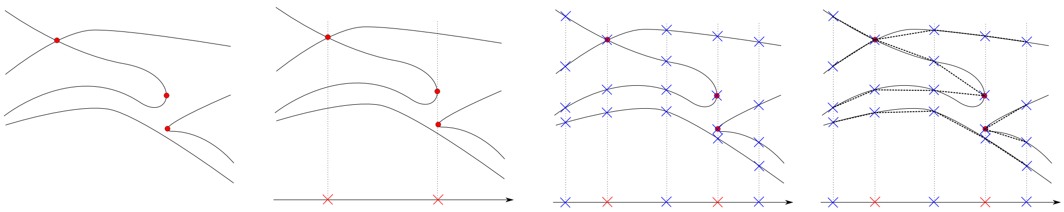

Fig. 2. Illustration of the construction in Lemma 1.

Proposition 1. *For any $n \in \mathbb{N}$, there exists an algebraic curve $\mathcal{O} \subset \mathbb{R}^2$ of degree n such that any isocomplex for $\mathcal{O}$ has $\Omega(n^2)$ vertices.*

In order to establish the upper bound of $O(n^2)$ for isocomplexes of algebraic curves, we show first that an algebraic curve decomposes into up to $O(n^2)$ points and smooth, x-monotone segments.

Definition 1. *Let $\mathcal{O} \subset \mathbb{R}^2$ be an algebraic curve without vertical segments. For a point $p \in \mathbb{R}^2$, the* branch numbers *of p are a pair of integers (ℓ_p, r_p) denoting the number of paths of the curve entering from the left hand side and from the right hand side, respectively. A point is called* event point *if its branch numbers do not equal $(1,1)$.*

Lemma 2. *For an event point p, we set b_p the sum of its branch numbers. Then, the sum of the b_p's for all event points is bounded by $2n(n-1)$.*

Proof. For a point $p = (x_0, y_0)$ on an algebraic plane curve $\mathcal{O} = V(f)$, we consider the Taylor expansion

$$f(x,y) = \sum_{i=0}^{n} (a_{i0}(x-x_0)^i + a_{i1}(x-x_0)^{i-1}(y-y_0) + \ldots + a_{ii}(y-y_0)^i)$$

of f at p. The smallest i such that at least one of the coefficients a_{ij}, $0 \leq j \leq i$, differs from zero is denoted the *multiplicity* $m_{\mathcal{O}}(p)$ of $\mathcal{O}$ at p. From this definition, it follows that $\mathcal{O}' := V(f_y)$ has multiplicity $m_{\mathcal{O}'}(p) \geq m_{\mathcal{O}}(p) - 1$ at p. Furthermore, the *intersection multiplicity* $\mathrm{int}(\mathcal{O}_1, \mathcal{O}_2, p)$ of two algebraic curves $\mathcal{O}_1 = V(f)$ and $\mathcal{O}_2 = V(g)$ at a point $p \in \mathbb{C}^2$ is defined as the dimension of the vector space $\mathbb{C}[x,y]_p/(f,g)$ where $\mathbb{C}[x,y]_p$ is the localization of the polynomial ring $\mathbb{C}[x,y]$ at p [3]. It holds that $m_{\mathcal{O}_1}(p) \cdot m_{\mathcal{O}_2}(p) \leq \mathrm{int}(\mathcal{O}_1, \mathcal{O}_2, p)$ with equality occurring iff f and g have no tangent line in common at p. Furthermore, due to Bézout's Theorem, the sum $\sum_{p \in \mathcal{O}_1 \cap \mathcal{O}_2} \mathrm{int}(\mathcal{O}_1, \mathcal{O}_2, p)$ of all intersection multiplicities is bounded by $\deg(f) \cdot \deg(g)$.

If $p = (x_0, y_0)$ is not an intersection point of $\mathcal{O}$ and $\mathcal{O}' := V(f_y)$, then p is adjacent to exactly two arcs of $\mathcal{O}$ which are orthogonal to the gradient $\nabla f(p) = (f_x(p), f_y(p))$ at p. Thus, the branch numbers for p are $(1,1)$. An event point $p = (x_0, y_0)$ is an intersection point of $\mathcal{O}$ and $\mathcal{O}'$ and, hence, $\mathrm{int}(\mathcal{O}, \mathcal{O}', p) \geq 1$ for each event point. The arithmetic mean $(\ell_p + r_p)/2$ of the two branch numbers ℓ_p and r_p at p constitutes a lower bound on the multiplicity of $\mathcal{O}$ at p; this follows from the fact that, for arbitrary small ε, there exists lines $L_x = V(x - x_0 + \varepsilon_x)$ and $L_y = V(y - y_0 + \varepsilon_y)$, $|\varepsilon_x|, |\varepsilon_y| < \varepsilon$, that both intersect $\mathcal{O}$ in at least $(\ell_p + r_p)/2$ points. This shows that the first $\lceil (\ell_p + r_p)/2 \rceil$ −order terms of the Taylor expansion of f at p vanish and, thus,

$$\sum_{p \text{ event point}} (\ell_p + r_p) \leq 2 \cdot \sum_{p \text{ event point}} m_{\mathcal{O}}(p) \leq 2 \cdot \sum_{p \text{ event point}} m_{\mathcal{O}}(p) \cdot m_{\mathcal{O}'}(p)$$
$$\leq 2 \cdot \sum_{p \in \mathcal{O}} \mathrm{int}(\mathcal{O}, \mathcal{O}', p) \leq 2n(n-1)$$

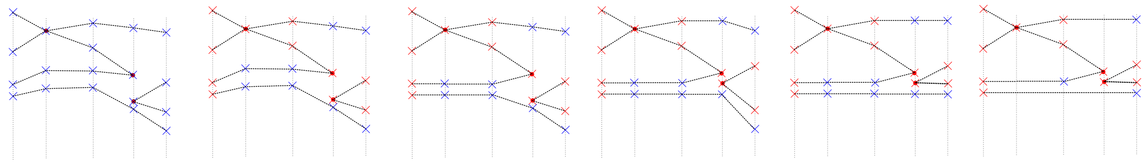

Fig. 3. Starting with the stable isocomplex of size $O(n^3)$ we straighten edges which connect two non critical points. Finally adjacent straight line connections are removed. The size of the so obtained isocomplex reduces to the number of arcs of $\mathscr{O}$ connecting two critical points, that is, $O(n^2)$.

Theorem 3. *For any algebraic curve $\mathscr{O} \subset \mathbb{R}^2$ of degree n, there exists an isocomplex with $O(n^2)$ simplices.*

Proof. We consider the isocomplex returned by a cylindrical algebraic decomposition algorithm. It returns $O(n^2)$ many fibers of the curves (with respect to some projection direction) and connects the fiber points by straight-line segments. Since any fiber has at most n points, the complexity is $O(n^3)$. We can assume that no segment is vertical and consider the complex as a directed graph from left to right, with the fiber points as vertices. In particular, it makes sense to talk about the *in-degree* of a vertex as the number of edges that enter from the left hand side. We re-embed the graph into the plane with the following rules. (1) Each vertex remains at the same x-coordinate, and the vertical ordering of the vertices at the same x-coordinate remains unchanged. (2) Each edge from a vertex of in-degree 1 to another vertex of in-degree 1 must be horizontal.

Properties (1) ensures that the result is isotopic to the original complex. A complex with properties (1) and (2) can be computed by a simple plane sweep algorithm (Fig. 3). Vertices adjacent to exactly two horizontal edges are removed afterwards, and the edges are merged. Let $\mathscr{C}_h$ denote this new complex. By construction, any maximal smooth x-monotone segment of the curve is represented by a polyline in $\mathscr{C}_h$ with two bends, running horizontally between the two bends. The number of edges is thus at most three times the number of segments of the curve that leave a critical point. Their number is be bounded by $O(n^2)$ according to Lemma 2 and, thus, the complexity of $\mathscr{C}_h$ is also $O(n^2)$.

4. Higher dimensions

We show to what extent our results for curves can be generalized to higher dimensions. Throughout this section, we consider $d \geq 2$ to be a fixed constant – this yields bounds of the form $\Omega/O(n^{h(d)})$ for some function h in d. However, one should keep in mind that the constants hidden in the O-notation depend on d. Furthermore, we still assume for simplicity that the considered hypersurface is bounded in each coordinate.

Stable isocomplexes: The construction of Theorem 2 can be immediately transfered to arbitrary dimensions:

Theorem 4. *For any $n \in \mathbb{N}$ and $d \geq 2$, there exists an algebraic hypersurface $\mathscr{O} \subset \mathbb{R}^d$ of degree n such that any stable isocomplex for $\mathscr{O}$ has $\Omega(n^{d+1})$ vertices.*

Proof. We cut out $\Theta(n^d)$ disjoint sections of the unit d-sphere and place an isolated point in each region, using Theorem 1. A vertex on the unit sphere must be placed in each region to ensure an isotopy. The bound follows by repeating the same argument on $n/2$ concentric d-spheres of radius close to 1.

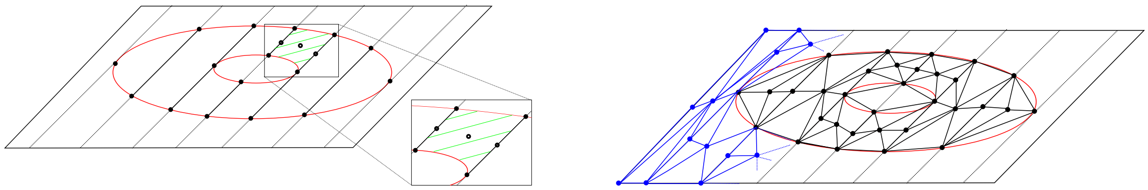

Fig. 4. Illustration of Theorem 6 for a torus: The silhouette $\mathcal{O}_R$ is a plane curve consisting of 2 circles (in red). $\mathcal{O}_R$ and the fibers at critical points and points in between decompose the projected bounding box B_R into quadrilaterals. Inserting points at the boundary and the interior of each of these quadrilaterals leads to a triangulation T_R of B_R. A subset S_R (in black, right figure) of T_R constitutes a stable isocomplex for $\mathcal{O}_R$.

Also the upper bound construction can be generalized; however, the exponent increases exponentially with d. A similar construction idea has also been used in [3, Thm 5.43]. A detailed description for the special case of 3 dimensions can be found in [4].

Theorem 5. *For a hypersurface $\mathcal{O} \subset \mathbb{R}^d$ of degree n, there exists a stable isocomplex with $O(n^{2^d-1})$ simplices.*

We will prove the theorem by proving a stronger statement.

Theorem 6. *Given an algebraic hypersurface $\mathcal{O} \subset \mathbb{R}^d$ of degree n with axis-aligned bounding box B. Then, there exist simplical complexes T and $S \subset T$ such that*

- *S is a stable isocomplex of $\mathcal{O}$*
- *T triangulates B*
- *T has $O(n^{2^d-1})$ simplices (and so has S)*

Proof. We do induction on the dimension d. In every dimension, we will first construct a stable isocomplex S and extend it to a triangulation T of B in a second step without increasing the complexity.

For $d = 2$, we construct a stable isocomplex as explained in Lemma 1. It consists of $O(n^3)$ simplices. Recall that in the construction, we have introduced $O(n^2)$ fibers, one for each critical x-coordinate, and one in between two consecutive coordinates. On each fiber, we add a point between two consecutive points on the fiber, and connect the point with its neighbors by a vertical segment. This decomposes the bounding box into trapezoids which are bounded on top and bottom by exactly one edge, and, on the left and right, by vertical segments (cf. Figure 3 and Figure 4). Summing up the number of vertical segments for all trapezoids in between two consecutive fibers, we have at most $O(n)$ segments since, on each fiber, we introduced $O(n)$ points. We pick a point p in the middle of a face and triangulate the face by connecting p with every vertex on the boundary and adding triangles accordingly. Then, at most $O(n)$ triangles are added for all trapezoids within two consecutive fibers, thus the bound of $O(n^3)$ holds for T. We have skipped the description of how to triangulate the boundary of the bounding box, but this is straight-forward by adding a fiber at the left and right boundary, considering the corners as fiber points. We skip further details.

For arbitrary $d > 2$, we consider the silhouette hypersurface $\mathcal{O}_R := \mathrm{res}_{x_d}(f, \frac{\partial f}{\partial x_d})$, where f is the defining equation for $\mathcal{O}$. Let B_R be the projection of B to the first $d-1$ variables. By the induction hypothesis, there exists a triangulation T_R of B_R containing a stable isocomplex S_R of $\mathcal{O}_R$ as a subset (cf. Figure 4). Moreover, T_R has $O(n^{2^d-2})$ simplices because $\mathcal{O}_R$ is of

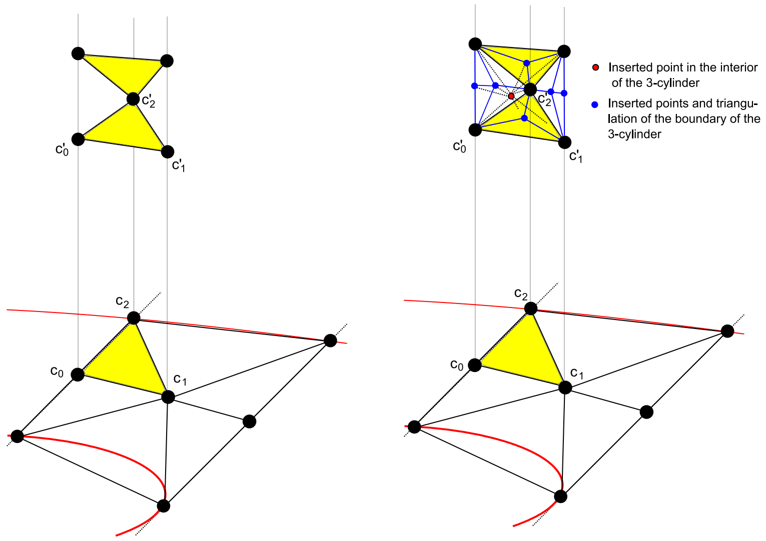

Fig. 5. Illustration of Theorem 6 for a torus ctd.: We consider the ”lift” of one of the triangles in the triangulation S_R. The triangle $\Delta c_0c_1c_2$ lifts to two triangles ℓ_1 and ℓ_2 in 3-space (cf. left figure). We insert points at the boundary and in the interior of the 3-cylinder between ℓ_1 and ℓ_2 (cf. right figure) to obtain a triangulation in 3-space.

degree $O(n^2)$. We first construct an isotopy from $\mathcal{O}$ to a stable isocomplex S. We proceed in two steps. For the first one, let ϕ_R be the isotopy between $\mathcal{O}_R$ and S_R. We can easily extend ϕ_R to an isotopy ϕ'_R from B_R to itself such that the vertices of T_R remain fixed during the transformation (the extension is not unique, but the choice does not matter for the argument). Note that $\mathcal{O}$ is delineable over any cell of $\mathcal{O}_R$, that is, the lift of each cell consists of disjoint function graphs. We can let ϕ'_R act on $\mathcal{O}$ as follows: For a point $p=(p_1,\ldots,p_d)\in\mathcal{O}$, we define $\phi'_R(p,t)=(\phi'_R((p_1,\ldots,p_{d-1}),t),p_d)$, that is, we leave the d-th coordinate fixed. This transforms $\mathcal{O}$ into some $\mathcal{O}'$ (which is not an algebraic set anymore). By construction, $\mathcal{O}'$ is delineable with respect to T_R, that is, the lift of each cell of T_R with respect to $\mathcal{O}'$ consists of (up to n) disjoint function graphs. We consider a k-simplex Δ of T_R with vertices $c_0,\ldots,c_k$, and one of its lifts, called c. c is uniquely defined by its “corners” $c'_0,\ldots,c'_k$ where c'_i is some lift of c_i (cf. Figure 5). In the second phase of the isotopy, we transform c to the simplex defined by $c'_0,\ldots,c'_k$. We can do so simultaneously for every lift without changing the vertical order, and without moving any vertex. We can also rule out the case that two lifts are mapped to the same simplex since by the way we construct $\mathcal{O}_R$, the lifts of any cell of T_R differ in at least one vertex. This implies that the transformation indeed is an isotopy between $\mathcal{O}'$ and S. We let S denote the isocomplex obtained by the described two-step transformation.

Finally, we complete S to a triangulation T of B. For a k-simplex Δ of T_R and two consecutive lifts ℓ_1,ℓ_2, we define the $(k+1)$-cylinder C between ℓ_1 and ℓ_2 to be the $(k+1)$-dimensional area between ℓ_1 and ℓ_2. Notice that the boundary of C might also contain lifts of vertices, edges, etc. By induction, we can assume that, above each of the $(k-1)-$simplices on the boundary of Δ, there exists a triangulation of the corresponding fiber with $O(n)$ many k-simplices (notice that $k\leq d$ can be treated as a constant!). Thus, the boundary of all $(k+1)$-cylinders above Δ admits a triangulation with $O(n)$ many elements.

We now traverse the simplices of T_R in increasing dimension, and triangulate the cylinders above each simplex. Let Δ be a k-simplex as above and denote $C_1,\ldots,C_m$, $m\leq n$, the cylinders above Δ. Then, from the above consideration, there exists a triangulation of the boundary of each C_i and the total number of k-simplices for all C_i is bounded by $O(n)$. We place a point p_i in the interior of each of the cylinders C_i, and construct simplices connecting p_i with all its boundary simplices (in other word, we construct a simplicial complex whose link is the boundary of the cylinder). By handling the boundary of the bounding box B in a similar fashion, this strategy yields a complex T that triangulates B.

Regarding the complexity of T, the number of simplices created above one k-simplex is at most $O(n)$ since each of k-simplices on the boundary of a cylinder yields at most a constant number of $(k+1)-$simplices. Hence, our bound follows from the induction hypothesis that T_R has $O(n^{2^d-2})$ simplices.

General isocomplexes: Again, the simple lower bound from Proposition 1 transfers directly into higher dimensions by considering n hyperplanes in generic position: Each set of d such hyperplanes intersects in a common point. The union of the hyperplanes yields an algebraic hyperplane of degree n with $\binom{n}{d}$ singularities. It follows:

Proposition 2. *For any $n \in \mathbb{N}$ and $d \geq 2$, there exists an algebraic hypersurface $\mathcal{O} \subset \mathbb{R}^d$ of degree n such that any isocomplex for $\mathcal{O}$ has $\Omega(n^d)$ vertices.*

Using Theorem 3 as a base case, we can improve the upper bound from Theorem 5 slightly for general isocomplexes:

Theorem 7. *Given a (compact) hypersurface $\mathcal{O}$ with axis-aligned bounding box B. Then, there exists simplical complexes T, $S \subset T$ such that*

- *S is an isocomplex of $\mathcal{O}$*
- *T triangulates B*
- *T has $O(n^{3/4 \cdot 2^d - 1})$ simplices*

In particular, there exists an isocomplex for $\mathcal{O}$ with $O(n^{3/4 \cdot 2^d - 1})$ cells.

Proof. We prove the claim by induction on d. For $d = 2$, Theorem 3 yields an isocomplex S with $O(n^2)$ simplices. To complete it to a triangulation of B, we first consider a trapezoidal decompostion [5] of B with respect to S, that means, from every vertex, we draw vertical rays upwards and downwards until we intersect another cell of S, or the boundary of B. This introduces 2 vertical segments, and up to 2 new vertices and does not increase the complexity. We barycentrically subdivide the trapezoidal decomposition, that means, we decompose each edge into two sub edges and a point in its interior (we need this for technical reasons in the induction because otherwise, it can happen that two lifts over an edge are transformed into the same 1-simplex in $\mathbb{R}^3$). Finally, we need to triangulate each trapezoid, which works in analogy to Theorem 6. For higher dimensions, we use exactly the same construction as in Theorem 6, and the same proof applies (note that in particular, the triangulation S that we construct is stable in the coordinates $x_1, x_3, x_4, \ldots, x_d$, only the x_2 coordinate changes).

We remark that, in [4], it was shown that there exists a stratification of an algebraic surface $\mathcal{O}$ in $\mathbb{R}^3$ with $O(n^5)$ many simply connected components. More precisely, these components are lifts of the $O(n^4)$ simply connected components of the arrangement induced by the projected silhouette curve $\mathcal{O}' \subset \mathbb{R}^2$. Our ”straightening idea” as presented in Theorem 3 shows that each of these components can be triangulated by a constant number of triangles.

5. Conclusion

Our main contribution is to establish tight bounds for the size of stable or general isocomplexes for algebraic curves of degree n. Our $O(n^2)$-bound for a general isocomplex also improves the complexity bound for an isocomplex in higher dimensions, but due to the projection strategy used in the construction, the bound remains double exponential in d. We believe that this upper bound is not tight – it might be possible to improve it by a triangulation method not based on projection. However, already for algebraic surfaces, it seems difficult to come up with a simplification algorithm which provably reduces the complexity and preserves the topology at the same time.

Another interesting variant of the problem is to further constrain the hypersurface, for instance, considering algebraic curves of degree n with a bounded bitsize, with a bounded number of singularities, or similar. We remark that, although our constructions yield reducible curves, the same bounds can be achieved with little extra effort when restricting to irreducible curves [14].

References

1. L. Alberti, B. Mourrain, J.-P. Técourt: "Isotopic triangulation of a real algebraic surface". *J. Symb. Comp.* **44** (2009) 1291–1310.

2. D. Barnette, A. Edelson: "All 2-manifolds have finitely many minimal triangulations". *Israel J. Math* **67** (1989) 123–128.

3. S. Basu, R. Pollack, M.-F. Roy: *Algorithms in Real Algebraic Geometry, Algorithms and Computation in Mathematics*, vol. 10. Springer, 2nd edn., 2006.

4. E. Berberich, M. Kerber, M. Sagraloff: "An Efficient Algorithm for the Stratification and Triangulation of Algebraic Surfaces". *Computational Geometry: Theory and Applications* **43** (2010) 257–278. Special issue on SoCG'08.

5. M. de Berg, M. van Kreveld, M. Overmars, O. Schwarzkopf: *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, Germany, 2nd edn., 2000.

6. J.-D. Boissonnat, M. T. (Editors): *Effective Computational Geometry for Curves and Surfaces*. Springer, 2006.

7. G. Bredon: *Topology and Geometry*. Springer, New-York, 1997.

8. B. F. Caviness, J. R. Johnson (eds.): *Quantifier Elimination and Cylindrical Algebraic Decomposition, Texts and Monographs in Symbolic Computation*. Springer, 1998.

9. J. Cheng, S. Lazard, L. Penaranda, M. Pouget, F. Rouillier, E. Tsigaridas: "On the topology of planar algebraic curves". In: *SCG '09: Proc. of the 25th Annual Symposium on Computational Geometry*. ACM, New York, NY, USA, 2009 361–370.

10. D. Cox, J. Little, D. O'Shea: *Using Algebraic Geometry*. Springer, New-York, 1998.

11. D. I. Diochnos, I. Z. Emiris, E. P. Tsigaridas: "On the Asymptotic and Practical Complexity of Solving Bivariate Systems over the Reals". *J. Symb. Comp.* **44** (2009) 818–835.

12. A. Eigenwillig, M. Kerber, N. Wolpert: "Fast and Exact Geometric Analysis of Real Algebraic Plane Curves". In: *Proc. of the 2007 International Symposium on Symbolic and Algebraic Computation (ISSAC 2007)*, 2007 151–158.

13. W. Fulton: *Intersection Theory*. Springer, New-York, 1998.

14. M. Kerber: *Geometric Algorithms for Algebraic Curves and Surfaces*. Ph.D. thesis, Universität des Saarlandes, Germany, 2009.

15. S. A. Lavrenchenko: “Irreducible triangulations of the torus”. *J. Math. Sci.* **51** (1990) 2537–2543.

16. A. Nakamoto, K. Ota: “Note on irreducible triangulations of surfaces”. *J. of Graph Theory* **20** (1995) 227–233.

Received: October 13, 2010
Final version received: October 13, 2010