

Complete Subdivision Algorithms, II: Isotopic Meshing of Singular Algebraic Curves

[Extended Abstract]

Michael Burr^{*}

Sung Woo Choi[†]

Benjamin Galehouse^{*}

Chee K. Yap^{*}

ABSTRACT

Given a real function $f(X, Y)$, a box region B and $\varepsilon > 0$, we want to compute an ε -isotopic polygonal approximation to the curve $C : f(X, Y) = 0$ within B . We focus on subdivision algorithms because of their adaptive complexity. Plantinga & Vegter (2004) gave a numerical subdivision algorithm that is exact when the curve C is non-singular. They used a computational model that relies only on function evaluation and interval arithmetic. We generalize their algorithm to any (possibly non-simply connected) region B that does not contain singularities of C . With this generalization as subroutine, we provide a method to detect isolated algebraic singularities and their branching degree. This appears to be the first complete *numerical* method to treat implicit algebraic curves with isolated singularities.

Categories and Subject Descriptors

F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems— *Geometrical Problems and Computations*

General Terms

Algorithms, Theory

Keywords

Meshing, Singularity, Root bound, Evaluation bound, Im-

^{*}Courant Institute, NYU, 251 Mercer Street, NY, NY 10012. Emails: {burr@cims, galehouse@cims, yap@cs}.nyu.edu. Yap's work is supported by NSF Grant CCF-0728977. Burr's work is partially supported by NSF Grant DMS-0701578. Galehouse's work is partially supported by the DoE under contract DE-FG02-00ER25053.

[†]Department of Mathematics, Duksung Women's University, Seoul, Korea. Email: swchoi@duksung.ac.kr

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISSAC'08, July 20–23, 2008, Hagenberg, Austria.

Copyright 2008 ACM 978-1-59593-904-3/08/07 ...\$5.00.

implicit algebraic curve, Complete numerical algorithm, Subdivision algorithm.

1. INTRODUCTION

Given $\varepsilon > 0$, a box region B and a real function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, we want to compute a polygonal approximation P to the implicit curve $C : f = 0$ restricted to B . The approximation P must be (1) “topologically correct” and (2) “ ε -close” to $C \cap B$. We use the standard interpretation of requirement (2), that $d(P, C \cap B) \leq \varepsilon$ where $d(\cdot, \cdot)$ is the Hausdorff distance on compact sets. In recent years, it has become accepted [2] to interpret requirement (1) to mean P is isotopic to $C \cap B$, which we denote by $P \approx C \cap B$. This means we not only require that P and $C \cap B$ to be homeomorphic, but they must be embedded in \mathbb{R}^2 “in the same way”. E.g., if $C \cap B$ consists of two ovals, these can be embedded in \mathbb{R}^2 as two ovals exterior to each other, or as two nested ovals. Isotopy, but not homeomorphism, requires P to respect this distinction. In this paper, we mainly focus on topological correctness, since achieving ε -closeness is not an issue for our subdivision approach (but cf. [2, pp. 213-4])

We may call the preceding the **2-D implicit meshing problem**. The term “meshing” comes from the corresponding problem in 3-D: given $\varepsilon > 0$ and an implicit surface $S : f(X, Y, Z) = 0$, we want to construct a triangular mesh M such that $d(M, S) \leq \varepsilon$ and $M \approx S$. It is interesting to identify the 1-D meshing with the well-known problem of real root isolation and refinement for a real function $f(X)$.

Among the approaches to most computational problems on curves and surfaces, the algebraic approaches and geometric/numerical approaches constitute two extremes of a spectrum. Algebraic methods can clearly solve most problems in this area, e.g., by an application of the general theory of cylindrical algebraic decomposition (CAD) [1]. Purely algebraic methods are generally not considered practical, even in the plane (e.g., [13, 21]). But efficient solutions have been achieved for special cases such as intersecting quadrics in 3-D [20]. At the other end of the spectrum, we have the geometric/numerical approaches that emphasize numerical approximation and iteration. An important class of such algorithms is the class of **subdivision algorithms** which can be viewed as generalized binary search. Such algorithms are practical in two senses: they are easy to implement and their complexity is more adaptive [26]. Another key feature of subdivision algorithms is “locality”, meaning that we can

restrict our computation to some region of interest.

Unfortunately, geometric/numerical methods seldom have global correctness guarantees. The most famous example is the Marching Cube algorithm (1987) of Lorensen & Cline. Many authors have tried to improve the correctness of subdivision algorithms (e.g., Stander & Hart [24]). So far, such efforts have succeeded under one of the following situations:

- (A0) Requiring niceness assumptions such as being non-singular or Morse.
- (A1) Invoking algebraic techniques such as resultant computation or manipulation of algebraic numbers.

It is clear that (A0) should be avoided. Generally, we call a method “complete” if the method is correct without any (A0) type restrictions. But many incomplete algorithms (e.g., Marching cube) are quite useful in practice. We want to avoid (A1) because algebraic manipulations are harder to implement and such techniques are relatively expensive and non-adaptive [26]. For instance, the subdivision meshing algorithm of Plantinga & Vegter [19, 18] requires an (A0) assumption, the non-singularity of surfaces. The subdivision algorithm of Seidel & Wolpert [21] requires¹ an (A1) technique, namely, the computation of resultants. We thus classify [21] as a hybrid approach that combines geometric/numerical with algebraic techniques. Prior to our work, we are not aware of any meshing algorithm that can handle singularities without resorting to resultant computation. In general, hybrid methods offer considerable promise (e.g., Hong [13]). This is part of a growing trend to employ numerical techniques for speeding up algebraic computation.

The recent collection [2, Chapter 5] reviews the current literature in meshing in 2- and 3-D: the subdivision algorithms of Snyder [23, 22] and also Plantinga & Vegter; the sampling approach of Boissonnat & Oudot [4] and Cheng, Dey, Ramos and Ray [8]; the Morse Theory approaches of Stander & Hart [24] and Boissonnat, Cohen-Steiner & Vegter [3]; and an algebraic sweepline approach of Mourrain & T  court [17]. The subdivision algorithm of Plantinga & Vegter is remarkable in the following sense: even though it is globally isotopic, it does not guarantee isotopy of the curve within each cell of the subdivision. In contrast, Snyder’s subdivision approach [23, 22] computes correct topology in each cell (the algorithm is currently incomplete [2, p. 195])

The basic idea of sampling approaches is to reduce meshing of a surface S to computing the Delaunay triangulation of a sufficiently dense set of sample points on S [2, p. 201–213] Cheng, Dey, Ramos and Ray [8] need the primitive operation of solving a system of equations involving f and its derivatives. Boissonnat and Oudot [4] need a primitive for intersecting the surface with a Voronoi edge. These primitives yield sample points on the surface. These points are algebraic, so implementing the primitives exactly would require strong algebraic techniques. For restrictions and open problems in sampling approaches, see [2, p. 227–229]. In sharp contrast, the computational primitives needed by the Plantinga & Vegter approach works directly with bigfloats, with modest requirements on f .

The complete removal of (A0) type restrictions is the major open problem faced by purely numerical approaches to meshing. Thus, Boissonnat et al [2, p. 187] states that

¹Their paper subtitle “Exploiting a little more Geometry and a little less Algebra” speaks to our concerns with (A1).

“*meshing in the vicinity of singularities is a difficult open problem and an active area of research*”. Most of the techniques described in their survey are unable to handle singularities. It should be evident that this open problem has an implicit requirement to avoid the use of (A1) techniques.

The present paper presents a purely numerical subdivision method for meshing algebraic curves with isolated singularities. In a certain sense, this is the most general geometric situation since reduced algebraic curves have only isolated singularities by Proposition 1. Our starting point is the algorithm of Plantinga & Vegter [19, 18] for implicit meshing of curves. It is important to understand the computational model of Plantinga & Vegter which is also used in this paper. Two capabilities are assumed with regards to $f(X, Y)$:

- (i) Sign evaluation of $f(p)$ at dyadic points p .
- (ii) f is C^1 and we can evaluate the interval analogues of f , $\frac{\partial f}{\partial X}$, $\frac{\partial f}{\partial Y}$ on dyadic intervals.

Note that Marching Cube only requires capability (i). Let the **class** PV denote the set of all real functions $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ for which capabilities (i) and (ii) are available. Many common functions of analysis (e.g., hypergeometric functions [10]) belong to PV . So the approach of Plantinga & Vegter admits a more general setting than algebraic curves.

Some of our recent work that addresses the above (A0)/(A1) concerns include [26] (Bezier curve intersection), [7] (solving triangular systems), [6] (numerical root isolation for multiple zeros) and [5] (integral analysis of real root isolation). The last two papers study the 1-D analogue of the Plantinga & Vegter Algorithm. The philosophy behind all these papers is the design and analysis of complete numerical methods based on approximations, iteration and adaptive methods. Topological exactness is achieved using suitable algebraic bounds, ranging from classic root separation bounds to evaluation bounds and geometric separation bounds. We emphasize that the worst-case complexity of adaptive algorithms (e.g., as determined by the worst case root bounds) ought not to be the chief criterion for evaluating the usefulness of these algorithms: for the majority of inputs, these algorithms terminate fast. Note that the zero bounds are only used as stopping criteria for iteration in the algorithms, and simple estimates can easily be computed. Such computations does not mean we compute resultants, even though their justification depend on resultant theory. The present paper continues this line of investigation.

Overview of Paper.

- In Section 3, we extend the Plantinga & Vegter algorithm to compute an isotopic approximation of the curve $C : f = 0$ restricted to a “nice region” that need not be simply connected. C may have singularities outside R and we only need $f \in PV$.
- In Section 4, we provide the algebraic evaluation bounds necessary for meshing singular curves.
- In Section 5, we provide a subdivision method to isolate all the singularities of a square-free integer polynomial $f(X, Y)$.
- In Section 6, given a box B containing an isolated singularity p , we provide a method to compute the branching degree of p .

- In Section 7, we finally present the overall algorithm to compute the isotopic polygonal approximation.
- We conclude in Section 8.

Basic Terminology. Let $\mathbb{F} := \mathbb{Z}[\frac{1}{2}] = \{m2^n : m, n \in \mathbb{Z}\}$ be the set of **dyadic numbers** or **bigfloats**. All our numerical computation are performed in \mathbb{F} . For $S \in \mathbb{R}$, let $\square S$ be the set of closed intervals $[a, b]$ with endpoints in S , $\{a, b\} \subseteq S$. We write $\square S^n$ for $(\square S)^n$. In particular, $\square \mathbb{F}$ is the set of dyadic intervals, and $\square \mathbb{R}^n$ is the set of n -boxes. The boundary of a set $S \subseteq \mathbb{R}$ is denoted ∂S . If $f : \mathbb{R}^n \rightarrow \mathbb{R}$, and $S \subseteq \mathbb{R}^n$, then $f(S) := \{f(x) : x \in S\}$. A function $\square f : \square \mathbb{F}^2 \rightarrow \square \mathbb{F}$ is a **box function** for f provided (i) $f(B) \subseteq \square f(B)$ and (ii) if $B_0 \supseteq B_1 \supseteq \dots$ and $\lim_i B_i \rightarrow p$ then $\lim_i \square f(B_i) \rightarrow f(p)$. We regard the limit of intervals in terms of the limits of their endpoints. We say $f \in PV$ if $f \in C^1$ (has continuous derivatives $\frac{\partial f}{\partial X}, \frac{\partial f}{\partial Y}$), there is an algorithm to determine $\text{sign}(f(p))$ for $p \in \mathbb{F}^2$ and $\square f, \square \frac{\partial f}{\partial X}, \square \frac{\partial f}{\partial Y}$ are computable.

The **size** of a box is the maximum length of one of its sides (all of our boxes will be square). We **split** a box B by subdividing it into 4 subboxes of equal size. These subboxes are the **children** of B , of half the size of B . Two boxes are **neighbors** if one box has an edge that overlaps an edge of the other box (the two boxes may have different sizes). Starting with B_0 , the child-parent relationships obtained by an arbitrary sequence of splits yields a **quadtrees** rooted at B_0 . The set of leaves in such a quadtree constitute a partition of B_0 . We only consider boxes of the form $B = I \times J$ where I, J are dyadic intervals. For simplicity, we assume B is square although it is possible to extend our algorithms to boxes with aspect ratio at most 2.

Basic Algebraic Facts. Let \mathbb{D} be a UFD and $f, g \in \mathbb{D}[\mathbf{X}] = \mathbb{D}[X_1, \dots, X_n]$ where $\mathbf{X} = (X_1, \dots, X_n)$. We say f, g are **similar** if there exists $a, b \in \mathbb{D} \setminus \{0\}$ such that $af = bg$, and write $f \sim g$. Otherwise, f and g are **dissimilar**. The **square-free part** of f is defined as

$$\text{SqFree}(f) := \frac{f}{\text{GCD}(f, \partial_1 f, \dots, \partial_n f)} \quad (1)$$

where $\partial_{X_i} = \partial_i$ indicates differentiation with respect to X_i . f is said to be **square-free** if $\text{SqFree}(f) = f$. From (1) we see that computing $\text{SqFree}(f)$ from f involves only rational operations of \mathbb{D} . As the gradient of f is $\nabla f = (\partial_1 f, \dots, \partial_n f)$, we may also write $\text{GCD}(f, \nabla f)$ for $\text{GCD}(f, \partial_1 f, \dots, \partial_n f)$. See [25, Chap. 2] for standard conventions concerning GCD .

Let k be an algebraically closed field. For $S \subseteq k[\mathbf{X}] = k[X_1, \dots, X_n]$ and $B \subseteq k^n$, let $\text{Zero}_B(S) := \{p \in B : f(p) = 0 \text{ for all } f \in S\}$ denote the **zero set of S relative to B** . If $B = k^n$, then we simply write $\text{Zero}(S)$.

In 1-dimension, a square-free polynomial $f \in \mathbb{Z}[X]$ has no singularities (i.e., multiple zeros). We now recall two generalizations of this result that will be necessary in the remainder of the paper. See [12, 9, 11] for similar results.

PROPOSITION 1 ([11, Ex.14.3]). *The singular points of any variety form a proper subvariety.*

The singular points of $\text{Zero}(f)$ are defined to be the points where $\nabla \text{SqFree}(f) = 0$. The above result is critical in our paper, because it implies that if $f \in \mathbb{R}[X, Y]$ is square-free, then the singular points are a proper subvariety of union of curves and hence must be a finite set of points.

PROPOSITION 2 (Sard's Theorem [11, Prop.14.4]). *Let $f : X \rightarrow Y$ be any surjective map of varieties defined over a field k of characteristic 0. Then there exists a nonempty open subset $U \subseteq Y$ such that for any smooth point $p \in f^{-1}(U) \cap X_{sm}$ in the inverse image of U , the differential df_p is surjective.*

Note that X_{sm} denotes the set of smooth points of variety X . The condition that the differential df_p is surjective is equivalent to insisting that $\nabla f(p) \neq 0$. The most important example that we consider is $f : \mathbb{R}^2 \rightarrow \mathbb{R}$. Every point in $\mathbb{R}^2 = X$ is smooth and $\mathbb{R} \setminus U$ is only a finite set. Hence, there are only a finite number of level sets parameterized by h where $\text{Zero}(f(X, Y) - h)$ has a singular point.

2. ALGORITHM OF Plantinga & Vegter

First we recall the Plantinga & Vegter algorithm. Given $\varepsilon > 0$, a box $B_0 \in \square \mathbb{F}^2$ and $f : \mathbb{R}^2 \rightarrow \mathbb{R}$, we want to compute a polygonal ε -approximation P of the restriction of the curve $C : f = 0$ to B_0 : $d(P, C \cap B_0) \leq \varepsilon$ and $P \approx C \cap B_0$. For simplicity, we focus on topological correctness: $P \approx C \cap B_0$, since it is easy to refine the subdivision to achieve $d(P, C \cap B_0) \leq \varepsilon$. The Plantinga & Vegter algorithm is based on two simple predicates on boxes B :

- Predicate $C_0(B)$ holds if $0 \notin \square f(B)$.
- Predicate $C_1(B)$ holds if $0 \notin (\square \frac{\partial f}{\partial X}(B))^2 + (\square \frac{\partial f}{\partial Y}(B))^2$.

These predicates are easily implemented if $f \in PV$, using interval arithmetic. Moreover, if $C_0(B)$ holds, then the curve C does not intersect B . Note that if B satisfies C_1 , then any child of B also satisfies C_1 .

The input box B_0 is a dyadic square, and output is an undirected graph $G = (V, E)$ where each vertex $v \in V$ is a dyadic point, $v \in \mathbb{F}^2$. G represents a polygonal approximation P of $C \cap B_0$.

The algorithm has 3 phases, where Phase i ($i = 1, 2, 3$) is associated with a queue Q_i containing boxes. Initially, $Q_1 = \{B_0\}$, and $Q_2 = Q_3 = \emptyset$. When Q_i is empty, we proceed to the Phase $i + 1$.

- **PHASE 1: SUBDIVISION.** While Q_1 is non-empty, remove some B from Q_1 , and perform the following: If $C_0(B)$ holds, B is discarded. If $C_1(B)$ holds, insert B into Q_2 . Otherwise, split B into four subboxes which are inserted into Q_1 .
- **PHASE 2: BALANCING.** This phase ‘‘balances’’ the subdivision; a subdivision is **balanced** the size of two neighboring boxes differ by at most a factor of 2. Queue Q_2 is a min-priority queue, where the size of a box serves as its priority. While Q_2 is non-empty, remove the min-element B from Q_2 , and perform the following: For each B -neighbor B' with size more than twice the size of B , remove B' from Q_2 and split B' . Insert each child B'' of B' into Q_2 provided $C_0(B'')$ does not hold. B'' might be a new neighbor of B and B'' might be split subsequently. When, finally, every neighbor of B is at most twice the size of B , we insert B into Q_3 .
- **PHASE 3: CONSTRUCTION.** This phase constructs the graph $G = (V, E)$. Initially, the boxes in Q_3 are unmarked. While Q_3 is non-empty, remove any B from Q_3 and mark it. Now construct a set $V(B)$ of vertices.

For each B -neighbor B' , if B' is unmarked, evaluate the sign of $f(p)f(q)$ where p, q are endpoints of the segment $B \cap B'$. If $f(p)f(q) < 0$, create a vertex $v = (p + q)/2$ for the graph G . Also put v into $V(B)$. NOTE: if $f(p) = 0$ for any endpoint p , arbitrarily perturb p so that $f(p) \neq 0$. If B' is marked, retrieve the vertex v (if any) on the edge $B \cap B'$, and put v into $V(B)$. It can be shown that $|V(B)| \in \{0, 2, 4\}$. If $|V(B)| = 2$, put the edge (p, q) into G to connect the vertices in $V(B)$. If $|V(B)| = 4$, there is a simple rule to insert two non-crossing edges into G (see [19, 18]).

The output graph $G = (V, E)$ is a collection $P = P(G)$ of closed polygons or polygonal lines with endpoints in ∂B_0 .




Figure 1: (a) incursion, (b) excursion, (c) boundary boxes and their complements

In what sense is P the correct output? Intuitively, P should be isotopic to $\{f = 0\} \cap B_0$. We certainly cannot handle the curve C having tangential but non-crossing intersection [26] with ∂B_0 . Assuming only transversal intersections, we still face two problems: if the curve C (locally) enters and exits ∂B_0 by visiting only one box $B \subseteq B_0$, the above algorithm would fail to detect this small component. See Figure 1(a). Conversely, the curve C might escape undetected from B_0 locally at a box B (Figure 1(b)). If we choose B_0 large enough, such errors cannot arise; but this is wasteful if we are only interested in a local region. If C has singularities, this is not even an option.

In this paper, we avoid any “largeness” assumption on B_0 . We next extend the Plantinga & Vegter algorithm to arbitrary B_0 so that a suitable correctness statement can be made about the output polygonal approximation P . In fact, B_0 need neither be a box nor be simply-connected.

3. EXTENSION OF Plantinga & Vegter

A major limitation Plantinga & Vegter that we seek to address is that small incursions and excursions, as displayed in Figure 1 might not be represented. The problem arises because Plantinga & Vegter constructs an isotopy to pull any B -excursion into B , or B -incursion into the neighboring B' ; but this can change the desired topology when B is a boundary box. Hence, we could eliminate this problem by ensuring that the curve passes through each boundary edge at most once. A test is for this can be done by ensuring that $0 \notin \square f_x(H)$ for horizontal boundary edges H , and similarly $0 \notin \square f_y(V)$ for vertical boundary edges V . This clearly yields a polygonal approximation P that satisfies $P \approx C \cap B_0$. This approach requires knowing the exact topology on the boundary of B_0 and resembles Snyder’s approach [23]; in higher dimensions, we need to recursively solve the problem in lower dimensions (on ∂B_0). This recursive solution can become expensive in higher dimensions.

In the spirit of the Plantinga & Vegter algorithm, we now provide an alternative solution that avoids exact boundary

topology. The idea is to slightly enlarge B_0 so that incursions/excursions can be removed by isotopy. This leads to a weaker correctness statement (Theorem 3). Compared to the exact (recursive) approach, we may split less often.

The basic idea is that, in addition to subdividing B_0 we find a slightly larger region B_0^+ which includes a collar of squares around B_0 . We ensure that at least one of predicates C_0, C_1 holds on each of these squares. Such a collar rules out excessive excursions. We then do some additional checks to ensure that any incursion is detected.

Call a box $B \subseteq B_0$ a **boundary box** if ∂B intersects ∂B_0 . Let B be a such a box. If B is not a corner box, it has a unique **complementary box** B' such that $\partial B' \cap \partial B_0 = \partial B \cap \partial B_0$ and the interiors of B' and B_0 are disjoint. Say B, B' are **partners** of each other. If B is a corner box, it determines two complementary boxes B', B'' . See Figure 1(c). Among complementary boxes that satisfy C_1 but not C_0 , we classify them into **transient** and **non-transient**, according to the sign pattern of f at their vertices, Figure 2. Intuitively, transient boxes are inconclusive for detecting incursions and need to be split. Eventually the split boxes are discarded or non-transient.




Figure 2: Classification of complementary boxes according to the sign of its vertices (up to reflection): (b), (c) are transient, but (a), (d), (e) are non-transient. Case (f) is excluded by C_1 . NOTE: The white and black vertices have opposite signs: if white is + then black is -, and vice-versa.

We now present the Extended Plantinga & Vegter algorithm. It has 3 Phases that parallel the algorithm in Section 2. Phase i (for $i = 1, 2, 3$) works off queues Q_i and Q'_i , transferring boxes into Q_{i+1} and Q'_{i+1} .

- PHASE 1: SUBDIVISION. While Q_1 is non-empty, remove some B from Q_1 , and perform the following: If $C_0(B)$ holds B is discarded. If $C_1(B)$ holds, and also $C_1(B')$ or $C_0(B')$ holds for every complementary box B' of B , put B into Q_2 and place its complementary boxes B' into Q'_2 . Otherwise, split B into four subboxes which are inserted into Q_1 .

We define B_0^+ to be the union of B_0 and all of the boxes B' which were placed into Q'_2 . Now, any excursion from boundary box B through ∂B_0 must remain within some complementary $B' \subseteq B_0^+$. So excursions are accounted for using deformations within B_0^+ . But we must detect incursions.

- PHASE 2: BALANCING with INCURSION CHECK. The balancing of boxes in Q_2 is done as in Phase 2 of Section 2. Next, we perform an analogous while-loop on Q'_2 : While Q'_2 is non-empty, remove any B' from Q'_2 . If the partner of B' had been split during balancing, we **half-split** B' (this means we split it into its four children and put the two children that intersect ∂B_0 into Q'_3). Otherwise, we place B' into Q'_3 .

For each B' that we put into Q'_3 , if B' is transient, we do an additional incursion check. Note that B' is transient means that the two endpoints of $B' \cap \partial B_0$ have a common sign $\sigma(B') \in \{+1, -1\}$. An incursion is evidenced by discovering any point in $B' \cap \partial B_0$ whose sign is different from $\sigma(B')$. If there is an incursion, we will place another copy of B' into an incursion queue Q_I . Correctness of this process (detailed next) will be demonstrated in full paper.

Let $Q_{B'}$ be a working queue, initialized to store B' . While $Q_{B'}$ is non-empty, remove some B'' from $Q_{B'}$. If an endpoint of $B'' \cap \partial B_0$ has sign different than $\sigma(B')$, insert B'' into Q_I and terminate this while-loop. If B'' is C_0 or non-transient, discard B'' . Otherwise, half-split B'' and put its two children into $Q_{B'}$. Termination of this loop is assured.

- PHASE 3: CONSTRUCTION. First, perform the Phase 3 of Section 2 which constructs a graph $G = (V, E)$. Next we augment this graph by adding a small incursion from each B' in Q_I into B_0 . More precisely, if B is the partner of B' , we insert two points u, v from the interior of the edge $B' \cap \partial B_0$ into the vertex set V . Also insert the edge (u, v) into the edge set E . This edge will be homotopic to a suitably defined incursion component.

Recall B_0^+ is B_0 augmented by a set of complementary boxes. The graph $G = (V, E)$ constructed by our algorithm represents a polygonal approximation $P \subseteq B_0$ comprising of polygonal paths and closed polygons.

THEOREM 3 (Weak Correctness). *Let $C = \{p \in \mathbb{R}^2 \mid f(p) = 0\}$ be non-singular in the (original) box B_0 . Let P be the polygonal approximation from the Extended Plantinga & Vegter Algorithm. If C only meets ∂B_0 transversally, then:*

- (1) The above procedure always halts.
- (2) There exists a region \mathcal{B} isotopic to B_0 , which satisfies $B_0 \subset \mathcal{B} \subset B_0^+$, such that $P \approx C \cap \mathcal{B}$.
- (3) Every component of P contained in ∂B_0 corresponds to at least one incursion.

Under condition (2) of this theorem, we call P a **weak isotopic approximation** to C within the region B_0 .

3.1 Extension to Nice Regions

It is essential in our applications later to extend the above refinements to non-simply connected regions. For this purpose, we define a **nice region** R_0 (relative to a square B_0) is the union of any collection of leaves taken from a quadtree rooted at B_0 . Thus, $R_0 \subseteq B_0$. To extend Theorem 3 to nice regions, we note two simple modifications:

- (a) A complementary box B' of a boundary box $B \subseteq R_0$ may intersect the interior of R_0 or other complimentary boxes.

Thus, Phase 1 must split such boundary boxes B sufficiently. (b) The region R_0 can have concave corners. We can classify the complementary boxes at concave corners into transient and non-transient ones as shown in Figure 3.




Figure 3: Classification of complementary boxes (up to reflection) at a concave corner: the case (b) is transient, and (a), (c), (d) are non-transient. Case (e) is excluded by the C_1 assumption.

4. EVALUATION BOUNDS

For any function f , define its **evaluation bound** to be

$$EV(f) := \inf\{|f(p)| : f(p) \neq 0, \nabla f(p) = 0\} \quad (2)$$

Such bounds were used in [7, 5]. From Proposition 2, we see that $\{f(p) : p \in \text{ZERO}(f), \nabla f(p) = 0\}$ is a finite set and therefore $EV(f) > 0$. However there is no explicit bound readily available. We provide such a bound:

LEMMA 4. *If $f \in \mathbb{Z}[X, Y]$ has degree d and $\|f\| < 2^L$ then $-\lg EV(f) = O(d^2(L + d))$. More precisely,*

$$EV(f)^{-1} \leq \max \left\{ \left[d^{d+5} 2^{L+2d+10} \right]^{d^2-1}, \left[d^{2d+6} 2^{3L+5d} \right]^d \right\}$$

Let f_x, f_y denote the derivatives of f . We may write

$$\text{ZERO}(f_x, f_y) = \bigcup_i U_i \cup \bigcup_j V_j$$

where U_i are 1-dim and V_j are 0-dim irreducible components. On each component U_i or V_j , one can show that the function f is constant. E.g. $f = (xy + 1)^2 - 1$, $f_x = 2(xy + 1)y$ and $f_y = 2(xy + 1)x$. Then $U_1 = \{xy + 1 = 0\}$ and $V_1 = \{(0, 0)\}$. The function f is equal to 1 on U_1 and equal to 0 on V_1 .

Let $g := \text{GCD}(f_x, f_y)$, and also

$$g_x := f_x/g, \quad g_y := f_y/g.$$

Clearly, we have

$$\text{ZERO}(f_x, f_y) = \text{ZERO}(g) \cup \text{ZERO}(g_x, g_y).$$

Since $\text{GCD}(g_x, g_y) = 1$, we conclude that $\text{ZERO}(g_x, g_y)$ has no 1-dimensional component. Conversely, the hyper-surface $\text{ZERO}(g)$ has no 0-dimensional component. This proves:

LEMMA 5.

$$\text{ZERO}(g) = \bigcup_i U_i, \quad \text{ZERO}(g_x, g_y) = \bigcup_j V_j.$$

We now provide some bounds. Let $\|f\|_k$ will denote the k -norm of f , where we use $k = 1, 2, \infty$. We just write $\|f\|$ for $\|f\|_\infty$, denoting the height of f . As parameters, we use d and L where $\deg f \leq d$ and $\|f\| < 2^L$.

We now view the ring $\mathbb{Z}[X, Y] \simeq \mathbb{Z}[X][Y] \simeq \mathbb{Z}[Y][X]$ in three alternative ways. A bivariate polynomial f in X and Y can be written as $f = f(X, Y)$, $f = f(X; Y)$ or $f = f(Y; X)$ to indicate these three views. As a member of $\mathbb{Z}[X, Y]$, the coefficients of $f(X, Y)$ are elements of \mathbb{Z} . But $f = f(X; Y)$ is a member of $\mathbb{Z}[Y][X]$ whose coefficients are polynomials in $\mathbb{Z}[Y]$. Similarly for $f = f(Y; X)$. The leading coefficient and degree of f are likewise affected by these views: $lc(f(X; Y)) \in \mathbb{Z}[X]$ but $lc(f(X, Y)) \in \mathbb{Z}$, $d = \deg(f(X, Y))$ is the total degree of f while $\deg(f(X; Y))$ is the largest power of Y occurring in f .

We use Mahler's basic inequality ([25, p. 351]) that if $g \in \mathbb{Z}[X, Y]$ and $g|f$ then

$$\|g(X, Y)\|_1 \leq 2^D \|f(X, Y)\|_1 \quad (3)$$

where $D = \deg(f(X; Y)) + \deg(f(Y; X))$. This implies:

$$\|g(X, Y)\|_1 \leq 4^{d-1} d^2 2^L, \quad \|g_x(X, Y)\|_1 \leq 4^{d-1} d^2 2^L. \quad (4)$$

It suffices to show the bound for $\|g\|_1$: note that $g|f_x$ and $\|f_x\|_1 \leq d^2 2^L$, $\deg(f_x(X; Y)) + \deg(f_x(Y; X)) \leq 2d - 2$. The bound then follows from (3).

Let $h(X)$ be the leading coefficient of $g(X; Y)$. Since $h(X)$ has degree $\leq d - 1$, there is an integer $x_0 \in \{0, 1, \dots, d - 1\}$ such that $h(x_0) \neq 0$. Intersect $\text{ZERO}(g)$ with the line $X = x_0$. CLAIM: This line cuts each non-vertical component U_i in a finite but non-zero number of points. In proof, let $g = \prod_i g_i$ where $\text{ZERO}(g_i) = U_i$. Setting $d_i := \deg g_i(X; Y)$, we see that the vertical components correspond to $d_i = 0$. Then $lc(g(X; Y)) = \prod_i lc(g_i(X; Y))$ and $lc(g(x_0; Y)) \neq 0$ iff for all i , $lc(g_i(x_0; Y)) \neq 0$. But $g_i(x_0; Y)$ is a polynomial of degree d_i in $\mathbb{Z}[Y]$, and has exactly d_i solutions in \mathbb{C} .

Write $f_0(Y) := f(x_0, Y)$ and $g_0(Y) = g(x_0, Y)$. From (4):

$$\|g_0\|_1 \leq d^d \|g(X, Y)\|_1 \leq 4^{d-1} d^{d+2} 2^L. \quad (5)$$

It is also easy to see that

$$\|f_0\| \leq d^{d+1} 2^L. \quad (6)$$

Suppose $\beta \in \text{ZERO}(g_0) \setminus \text{ZERO}(f_0)$. We want a lower bound on $|f_0(\beta)|$. For this purpose, we use an evaluation bound from [5, Theorem 13(b)]:

PROPOSITION 6 (Evaluation Bound [5]). *Let $\phi(x), \eta(x) \in \mathbb{C}[x]$ be complex polynomials of degrees m and n . Let β_1, \dots, β_n be all the zeros of $\eta(x)$. Suppose there exists relatively prime $F, H \in \mathbb{Z}[x]$ such that $F = \phi \bar{\phi}, H = \eta \bar{\eta}$ for some $\bar{\phi}, \bar{\eta} \in \mathbb{C}[x]$. If the degrees of $\bar{\phi}$ and $\bar{\eta}$ are \bar{m} and \bar{n} , then*

$$\prod_{i=1}^n |\phi(\beta_i)| \geq \frac{1}{lc(\bar{\eta})^m \cdot ((m+1)\|\phi\|)^{\bar{m}} M(\bar{\eta})^m \cdot ((\bar{m}+1)\|\bar{\phi}\|)^{n+\bar{n}} M(H)^{\bar{m}}}. \quad (7)$$

We shall choose the variables in Proposition 6 as follows:

$$\phi := f_0, \quad H = \frac{g_0}{\text{GCD}(f_0, g_0)}.$$

Moreover, let $\bar{\phi} := 1$, $\eta(X) := X - \beta$ and $\bar{\eta} := H/\eta$. Hence

$$m \leq d, \quad n \leq d, \quad \bar{m} = 0, \quad \bar{n} \leq d - 1.$$

Also

$$lc(\bar{\eta}) = lc(H) = lc(g_0) \leq \|g_0\| \leq \|g_0\|_1. \quad (8)$$

Further,

$$M(\bar{\eta}) \leq M(H) \leq \|H\|_1 \leq 2^d \cdot \|g_0\|_1.$$

Finally, an application of Proposition 6 gives

$$\begin{aligned} |f_0(\beta)|^{-1} &\leq lc(\bar{\eta})^d \cdot ((d+1)\|f_0\|)^{d-1} \cdot M(\bar{\eta})^d \\ &< [lc(\bar{\eta}) \cdot (d+1)\|f_0\| \cdot M(\bar{\eta})]^d \\ &\leq \left[\|g_0\|_1 \cdot (d+1)2^L \cdot 2^d \|g_0\|_1 \right]^d \\ &\leq \left[d^{2d+6} 2^{3L+5d} \right]^d. \end{aligned} \quad (9)$$

(9) is a lower bound on $|f(p)|$ where p lies in a non-vertical component U_i . By considering $g(Y; X)$, the same bound applies for $|f(p)|$ when p lies in a vertical component U_i .

We obtain a lower bound for $f(p)$ with $p \in \text{ZERO}(g_x, g_y)$. Consider the system $\Sigma \subseteq \mathbb{Z}[X, Y, Y]$ where

$$\Sigma = \{Z - f(X, Y), g_x(X, Y), g_y(X, Y)\}$$

The zeros $(\xi_1, \xi_2, \xi_3) = (\xi_1, y, f(x, y)) \in \mathbb{C}^3$ of Σ satisfy $\xi_3 = f(\xi_1, \xi_2)$. Since Σ is a zero dimensional system, we may apply the multivariate zero bound in [25, p. 350]. This bound says that

$$|\xi_3|^{-1} < (2^{3/2} NK)^D 2^{8(d-1)} \quad (10)$$

where $N = \binom{1+2(d-1)}{3}$, $D = d^2 - 1$ and

$$K = \max\{\sqrt{3}, \|g_x\|_2, \|g_y\|_2, \|Z - f(X, Y)\|_2\}.$$

We have $\|Z - f(X, Y)\|_2 \leq 1 + (d+1)2^L$. From (4), we see that $K \leq 4^{d-1} d^{d+2} 2^L$. Using the bound $N < 2d^3$, we obtain

$$|\xi_3|^{-1} < [4^{d+5} d^{d+5} 2^L]^{d^2-1}. \quad (11)$$

Now Lemma 4 easily follows from (9) and (11).

5. ISOLATING SINGULAR POINTS

In the rest of this paper, we assume that $f \in \mathbb{Z}[X, Y]$, and the curve $C : f = 0$ intersects ∂B_0 tangentially. We would like to use the Extended Plantinga & Vegter algorithm to compute an isotopic approximation to $\text{Zero}(f)$ when f has only isolated singularities. Since the Plantinga & Vegter algorithm does not terminate near singular points, it is necessary to isolate the singular points from the rest of B_0 .

We use the auxiliary function $F = f^2 + f_X^2 + f_Y^2$. Finding the singular points of $f = 0$ amounts to locating and isolating the zeros of this non-negative function. We use a simple mountain pass theorem [14] adapted to B_0 to ensure our algorithm isolates the zeros.

THEOREM 7. *Suppose that $F \geq 0$ on B_0 , and that $F > 0$ on ∂B_0 . Then there is path $\gamma : [0, 1] \rightarrow B_0$ connecting the two distinct roots of F which minimizes $M_\gamma = \max_{x \in [0, 1]} F(\gamma(x))$ and it either contains a point y where $\nabla F(\gamma(y)) = 0$ or a point y where $\gamma(y) \in \partial B$.*

This can be proved using path deformation and the compactness of B_0 , or it can be seen as a simple application of the topological mountain pass theorem presented in [14]. Because of this theorem, distinct zeros of F within B_0 are separated by barriers of height $\epsilon = \min(\text{EV}(F), \min F(\partial B_0))$. This leads us to the following multistep process to localize these zeros. The goal is to find a small rectangle with diameter less than some δ around each zero.

STEP 0: DETERMINING ϵ . Push B_0 into a queue of squares Q_1 . While there is an S in Q_1 remove it and evaluate $\square F(S)$. If $\square F(S) > 0$ we push S into the queue Q_∂ . If $0 \in \square F(S)$, subdivide S and push the children of S which intersect ∂B_0 into Q_1 and the others into Q_{int} . Once this terminates, $Q_1 = \emptyset$ and we have a collection of final squares Q_∂ , which contains all of ∂B_0 . For each of these S we actually can find an ϵ_S with $\square F(S) > \epsilon_S > 0$. We take ϵ to be the minimum of all these ϵ_S and $\text{EV}(F)$

STEP 1: INITIAL SUBDIVISION. Initialize queue Q_2 with the union of Q_∂ and Q_{int} . Initialize Q_3 to be empty. Reusing the initial subdivision is only an optimization. While there is an S in Q_2 remove it and evaluate $\square F(S)$. If $\square F(S) > \epsilon/3$, discard S . Else if $F(S) < 2\epsilon/3$, place S into Q_3 . Else subdivide S and push its children into Q_2 .

Once Q_2 is empty, group the elements of Q_3 into connected regions A_i ($i \in I$). Each A_i contains at most one root, since otherwise, there would be a path connecting the roots within A_i . The value of F along this path would be bounded above by $2\epsilon/3$ contradicting the mountain pass theorem. For later reference, let C be the region $B_0 \setminus \cup_i A_i$. F is greater than $\epsilon/3$ on C and that $\partial B_0 \subset C$ by Step 0.

STEP 2: REFINEMENT. For each A_i ($i \in I$), initialize queue $Q_{4,i}$ with all squares $S \in A_i$. So long as neither terminating condition 1 nor 2 (below) hold, we perform the following: For each S in $Q_{4,i}$, if $0 \in \square F(S)$, subdivide S and push its children into $Q_{4,i}$. If $0 \notin \square F(S)$, discard S . We terminate when either of the following two conditions are met:

1. $Q_{4,i}$ is empty, in which case there isn't a zero in A_i .
2. A'_i , the contents of $Q_{4,i}$ satisfy all of the following:
 - (a) $\square F(S) < \epsilon/3$ for some $S \in A'_i$
 - (b) R_i , the smallest rectangle containing A'_i , lies within the region covered by the original A_i .
 - (c) The diameter of R_i is less than δ .

We claim that each R_i contains exactly one root. In Step 1, we showed that A_i contains at most one root. To see that R_i contains a root, take a point of A'_i where $F < \epsilon/3$, then follow the path of steepest descent to reach a zero of F . Because F is less than $\epsilon/3$ on this curve, the curve cannot pass through the region C to reach any other R_j or to leave B_0 . Therefore there must be a zero within A_i . It is in R_i because our conditions ensure that F is positive on $A_i \setminus R_i$.

6. DETERMINING THE DEGREE OF SINGULAR POINTS

The following standard result from [15, 16] shows that the global structure of zero sets:

PROPOSITION 8 (Zero Structure). *Let f be real analytic. Then $\text{Zero}(f)$ can be decomposed into a finite union of pieces homeomorphic to $(0, 1)$, pieces homeomorphic to S^1 , and singular points.*

Viewing $\text{Zero}(f)$ as a multigraph G , the **degree** of a singular point is its degree as a vertex of G . We now determine such degrees. Let δ_3 be a separation bound between singular points, so if p and q are two distinct singular points of $\text{Zero}(f)$, then the distance between p and q is at least δ_3 . Let δ_4 be a separation bound so that if r is a point on $\text{Zero}(f)$

such that $\nabla f(r)$ is in the same direction as the line between p and r , then the distance between p and r is at least δ_4 . If s is on $\text{Zero}(f)$ so that s is closer to the singular point p than either δ_3 or δ_4 , then by following the paths $\text{Zero}(f)$ away from s , one of the paths strictly monotonically approaches p until it reaches p and the other path locally strictly monotonically recedes from p . See [26] for explicit bounds on δ_3 as a function of degree and height of $f(X, Y)$. We can similarly derive explicit bounds on δ_4 .

To find the degree of a singular point, assume that we have two boxes $B_1 \supseteq B_2$ where the diameter of B_1 is less than both δ_3 and δ_4 , B_2 contains a singular point of f and there is some radius $r > 0$ such that a circle of radius r centered at any point p inside B_2 must lie entirely within the annulus $B_1 \setminus B_2$. See Figure 4. Furthermore, to apply our extended Plantinga & Vegter algorithm of Section 4, we can ensure that $B_1 \setminus B_2$ is a nice region.




Figure 4: Annular region $B_1 \setminus B_2$ with singularity p and the three types (1), (2), (3) of components.

Now, there are 3 types of components in $\text{Zero}(f) \cap (B_1 \setminus \text{int}(B_2))$: (1) images of $[0, 1]$ both of whose endpoints are on ∂B_1 , (2) images of $[0, 1]$ both of whose endpoints are on ∂B_2 , and (3) images of $[0, 1]$ with one endpoint on each of ∂B_1 and ∂B_2 . These three types are illustrated in Figure 4. Let s be a point on any of these components, then traveling along $\text{Zero}(f)$ in one direction must lead to the singular point and the other direction must leave the neighborhood (be further than $\min\{\delta_3, \delta_4\}$) of the singular point.

LEMMA 9. *The degree of the singular point in B_2 is the number of components of type 3.*

Any component accumulating on a singular point exits the neighborhood of the singular point and the only way to leave the neighborhood is by way of a type 3 component.

7. OVERALL ALGORITHM

We now put all the above elements together to find a weak isotopic approximation to the algebraic curve $C : f = 0$ within a nice region R_0 where $f(X, Y) \in \mathbb{Z}[X, Y]$ has only isolated singularities. For simplicity, we assume that ∂R_0 intersects the curve $C : f = 0$ transversally. We first find the singularities of the curve $C : f = 0$ in R_0 . Using the technique of Section 5, we can isolate the singularities p_i ($i = 1, 2, \dots$) into pairwise disjoint boxes B_i . We may assume the size of the B_i 's is at most $\min\{\delta_3, \delta_4\}/6$. Let B'_i be the box of size 5 times the size of B_i , and concentric with B_i ; we may further assume $B'_i \subseteq R_0$. Now we proceed to run the extended Plantinga & Vegter algorithm on the nice region $R^* := R_0 \setminus \cup_i B_i$, yielding a polygonal approximation P . We directly incorporate the technique of Section 6 into

the following argument. If p_i is the singular point in B_i , then the degree of p_i is equal to the number of type 3 components in $P \cap (B'_i \setminus B_i)$. We directly connect these components directly to p_i , and discard any type 2 components. This produces the desired isotopic approximation.

Remarks: (1) We could avoid the assumption that C and ∂R_0 intersect transversally provided R_0 is a nice region relative to a box B_0 whose corners have integer or algebraic coordinates. Using the geometric separation bounds in [26] we can detect an actual transversal intersection.

(2) We have not discussed ε -approximation because this is relatively easy to achieve in the Plantinga & Vegter approach. We only have to make sure that each subdivision box that contains a portion of the polygonal approximation P has size at most $\varepsilon/4$.

8. CONCLUSION

This paper presents the first complete numerical subdivision algorithm for meshing an implicit algebraic curve that has only isolated singularities. This solves an open problem in the exact numerical approaches to meshing in 2-D [2, p. 187]. We pose three challenges:

(a) An worst case complexity bound for our procedure is possible. But this may not be the best way to measure adaptive algorithms. We would like to provide are adaptive bounds, similar to the integral analysis in [5].

(b) In 3-D, a square-free integer polynomial $f(X, Y, Z)$ could have a 1-dimensional singularities. We pose the problem of designing a purely numerical subdivision algorithm to handle 1-dimensional singularities.

(c) The practical implementation of an adaptive algorithm handling singularities, even based on our outline, must handle many important details. Computational experience is invaluable for future research into singularity computation.

9. REFERENCES

- [1] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*. Algorithms and Computation in Mathematics. Springer, 2003.
- [2] J.-D. Boissonnat, D. Cohen-Steiner, B. Mourrain, G. Rote, and G. Vegter. Meshing of surfaces. In J.-D. Boissonnat and M. Teillaud, editors, *Effective Computational Geometry for Curves and Surfaces*. Springer, 2006. Chapter 5.
- [3] J.-D. Boissonnat, D. Cohen-Steiner, and G. Vegter. Isotopic implicit surfaces meshing. In *ACM Symp. on Theory of Computing*, pages 301–309, 2004.
- [4] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of Lipschitz surfaces. In *Proc. 22nd ACM Symp. on Comp. Geometry*, pages 337–346, 2006. Sedona, Arizona.
- [5] M. Burr, F. Krahmer, and C. Yap. Integral analysis of evaluation-based real root isolation, Mar. 2008. Submitted, FOCS 2008.
- [6] M. Burr, V. Sharma, and C. Yap. Evaluation-based root isolation, Nov. 2007. In preparation.
- [7] J.-S. Cheng, X.-S. Gao, and C. K. Yap. Complete numerical isolation of real zeros in general triangular systems. In *Proc. ISSAC'07*, pages 92–99, 2007.
- [8] S.-W. Cheng, T. Dey, E. Ramos, and T. Ray. Sampling and meshing a surface with guaranteed topology and geometry. In *Proc. 20th ACM Symp. on Comp. Geometry*, pages 280–289, 2004.
- [9] D. Cox, J. Little, and D. O’Shea. *Ideals, Varieties and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*. Springer-Verlag, New York, 1992.
- [10] Z. Du and C. Yap. Uniform complexity of approximating hypergeometric functions with absolute error. In S. il Pae and H. Park, editors, *Proc. 7th Asian Symp. on Computer Math. (ASCM 2005)*, pages 246–249, 2006.
- [11] J. Harris. *Algebraic Geometry*. Springer-Verlag, 1992.
- [12] R. Hartshorne. *Algebraic Geometry*. Springer-Verlag, New York, 1977.
- [13] H. Hong. An efficient method for analyzing the topology of plane real algebraic curves. *Mathematics and Computers in Simulation*, 42:571–582, 1996.
- [14] Y. Jabri’. *The Mountain Pass Theorem: Variants, Generalizations and Some Applications*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2003.
- [15] S. G. Krantz and H. R. Parks. *A primer of real analytic functions*. Birkhäuser Verlag, Basel, 1992.
- [16] S. Lojasiewicz. *Introduction to complex analytic geometry*. Birkhäuser Verlag, Basel, 1991. Translated from the Polish by Maciej Klimek.
- [17] B. Mourrain and J.-P. Tércourt. Isotopic meshing of a real algebraic surface. Technical Report RR-5508, INRIA, Sophia-Antipolis, France, Feb. 2005. Also, electronic proceedings, MEGA 2005.
- [18] S. Plantinga. *Certified Algorithms for Implicit Surfaces*. Ph.D. thesis, Groningen University, Institute for Mathematics and Computing Science, Groningen, Netherlands, Dec. 2006.
- [19] S. Plantinga and G. Vegter. Isotopic approximation of implicit curves and surfaces. In *Proc. Eurographics Symposium on Geometry Processing*, pages 245–254, New York, 2004. ACM Press.
- [20] E. Schoemer and N. Wolpert. An exact and efficient approach for computing a cell in an arrangement of quadrics. *Comput. Geometry: Theory and Appl.*, 33:65–97, 2006.
- [21] R. Seidel and N. Wolpert. On the exact computation of the topology of real algebraic curves. In *Proc. 21st ACM Symp. on Comp. Geometry*, pages 107–116, 2005. Pisa, Italy.
- [22] J. M. Snyder. *Generative Modeling for Computer Graphics and CAD: Symbolic Shape Design using Interval Analysis*. Academic Press, 1992.
- [23] J. M. Snyder. Interval analysis for computer graphics. *SIGGRAPH Comput. Graphics*, 26(2):121–130, 1992.
- [24] B. T. Stander and J. C. Hart. Guaranteeing the topology of an implicit surface polygonalization for interactive meshing. In *Proc. 24th Computer Graphics and Interactive Techniques*, pages 279–286, 1997.
- [25] C. K. Yap. *Fundamental Problems of Algorithmic Algebra*. Oxford University Press, 2000.
- [26] C. K. Yap. Complete subdivision algorithms, I: Intersection of Bezier curves. In *22nd ACM Symp. on Comp. Geometry*, pages 217–226, July 2006.