

A singly exponential stratification scheme for real semi-algebraic varieties and its applications*

Bernard Chazelle

Princeton University

Herbert Edelsbrunner

University of Illinois at Urbana-Champaign

Leonidas J. Guibas

Stanford University and DEC/SRC

Micha Sharir

New York University and Tel Aviv University

Abstract

Chazelle, B., H. Edelsbrunner, L.J. Guibas and M. Sharir, A singly exponential stratification scheme for real semi-algebraic varieties and its applications, *Theoretical Computer Science* 84 (1991) 77-105.

This paper describes an effective procedure for stratifying a real semi-algebraic set into cells of constant description size. The attractive feature of our method is that the number of cells produced is singly exponential in the number of input variables. This compares favorably with the doubly exponential size of Collins' decomposition. Unlike Collins' construction, however, our scheme does not produce a cell complex but only a smooth stratification. Nevertheless, we are able to apply our results in interesting ways to problems of point location and geometric optimization.

1. Introduction

This paper studies techniques for building economical stratifications of real semi-algebraic sets. Let f_1, \dots, f_n be n d -variate polynomials with rational coefficients; we assume that the number of variables d as well as the maximum algebraic degree of the polynomials are independent of n . We seek a partition of

* The authors wish to thank DEC/Systems Research Center and DEC/Paris Research Laboratory, where part of this research was conducted. For individual support, Bernard Chazelle acknowledges the National Science Foundation for supporting this research in part under Grant CCR-8700917. Herbert Edelsbrunner acknowledges the support of the National Science Foundation under Grant CCR-8714565. Micha Sharir acknowledges the Office of Naval Research under Grant N00014-87-K-0129, the National Science Foundation under Grant No. NSF-DCR-83-20085, grants from the Digital Equipment Corporation and the IBM Corporation, and a research grant from the US-Israeli Binational Science Foundation.

\mathfrak{R}^d into “simply-shaped” cells, of dimensions ranging from 0 to d , so that each f_i has constant sign (0, positive, or negative) over each cell c in the decomposition. If, in addition, each cell is a smooth manifold, such a decomposition is then called a *sign-invariant stratification*. Our goals are (i) to keep the number of cells as small as possible, and (ii) to keep the “shape” of each cell as simple as possible (both topologically and combinatorially). Obviously, the number of cells cannot be smaller than the number of connected components into which the varieties $\bigvee f_i \equiv \{f_i = 0\}$ partition \mathfrak{R}^d . In the worst case this number is on the order $\Theta(n^d)$, as easily follows from Milnor’s Theorem [6, 7, 38]. Note that these components might be very complex and thus completely unsuitable for our purposes. In particular, the number of polynomials needed to define a single connected component (in the unquantified first-order theory of the reals) might be very large, not to mention its topology which can be also very complex. To enforce property (ii), and more specifically, to ensure that each cell can be described by a constant-size formula and is diffeomorphic to an open k -ball, for some $k \leq d$, we need to cut up each such component still further.

This problem has been studied extensively over the last 15 years. Collins’ landmark paper [22] yields a sign-invariant stratification with $O(n^{2^d-1})$ cells of simple shape. The resulting structure is powerful enough to decide the truth of any quantified formula in the first-order theory of reals, and in doing so, eliminates quantifiers from such formulae. In fact, quantifier elimination has been recently shown to be inherently doubly exponential in the number of variables [25]. Recent findings show, however, that many restricted problems related to the theory of reals can be solved in singly exponential time and storage. For example, deciding the existential theory of the reals [42], eliminating quantifiers from a formula with a bounded number of alternations between universal and existential quantifiers [9, 30], or deciding if two points lie in the same connected component [10]. Our paper can be regarded as another step in that direction.

Let us first motivate our study by its applications. A major one is the generalized point location problem discussed in [16] and its applications. Let f_1, \dots, f_n be n d -variate polynomials as above, and let x be a point in \mathfrak{R}^d : is x a zero of any f_i ? It is understood that the polynomials are given once and for all, but that the point x is a *query* which must be answered on-line. In many applications it is desirable to obtain more information than a simple yes-or-no answer, so we add the following requirements. If the answer is positive, the index i of some f_i for which x is a zero should be given. Otherwise, the point x falls in some connected component c of $\bigcap_{1 \leq i \leq n} \{y \in \mathfrak{R}^d \mid f_i(y) \neq 0\}$, and the output should return a pointer to some precomputed point in c , or more generally, some precomputed attribute associated with c . Often, it is useful to obtain information about the varieties at or right above the query point. For example, if $x = (x_1, \dots, x_d)$ is not a zero of any f_i , this might mean providing the index k of some f_k (if any) such that $f_k(x_1, \dots, x_{d-1}, z)$ has the smallest real root (in z) larger than x_d among all f_i ’s.

The motivation for studying this generalized form of point location is that its language is powerful enough to express *any* multidimensional searching problem

expressed as a first-order predicate in the theory of real-closed fields. A related application, which in fact is also used as a subroutine in the point location algorithm, is the following general paradigm: We are given the polynomials f_1, \dots, f_n as input data to some problem that needs to be solved over the entire space \mathfrak{R}^d . We would like to break the problem into independent subproblems, by decomposing \mathfrak{R}^d into a small number of cells and by obtaining in each cell c a subproblem that involves only the polynomials whose varieties $\bigvee f_i$ intersect c . If we can keep both the number of cells and the number of varieties crossing each cell small, then this divide-and-conquer scheme will be efficient. This paradigm has indeed been used for point location [18] (albeit only for hyperplanes), as well as for a miscellany of other algorithmic and combinatorial applications (see e.g., [4, 14, 17, 19, 20, 27, 32, 41]). With the exception of [20], however, these applications involve only linear features (points, lines, hyperplanes, etc.). Moreover, most of these studies involve planar decompositions, and only very few efficient decomposition techniques are known in three dimensions [4, 12, 15] or higher [22].

The extensive theory of *random sampling* that has been developed in the last few years (e.g., in [14, 17, 19, 20, 27, 32, 41]) provides a tool to implement this divide-and-conquer paradigm: Choose a random sample R of r of the varieties $\bigvee f_i$ and obtain a sign-invariant decomposition of \mathfrak{R}^d for R . The analysis of [14, 17, 19, 32] implies that if each cell c in the decomposition has a simple shape, then, with high probability, no cell meets more than $an(\log r)/r$ varieties (for some constant a that depends on the dimension d and the degree of the given polynomials). Chazelle and Friedman [14] provide a deterministic method for constructing such a decomposition. Thus the size of the decomposition is a crucial factor in the overall complexity of this divide-and-conquer technique.

This paper provides an efficient new technique for stratifying real semi-algebraic sets. Roughly speaking, we show how to partition d -space into cells of constant description-size, over which the signs of the f_i 's remain invariant. Each cell is a smooth connected manifold which admits a simple parametrization and can be fully specified as a semi-algebraic set over a constant number of polynomials. The number of cells is $O(n)$ in one dimension and $O(n^{2^{d-2}})$ in dimension $d > 1$. Actually, with a bit of extra work it is possible to lower the space requirement to $O(n^{2^{d-3}}\beta(n))$ for $d > 2$, where $\beta(n)$ is a very slow-growing function (so slow that its inverse is not even primitive-recursive); specifically, we have $\beta(n) = 2^{\alpha(n)^c}$, where c is a constant dependent only on the dimension d and the maximum degree of the input polynomials, and α is a functional inverse of Ackermann's function. This fairly minor improvement requires a lengthy analysis, so it will be omitted. The construction can be performed in time $O(n^{2^{d-1}} \log n)$. Within the same asymptotic time we can also compute an algebraic point in each cell of the decomposition.

As we mentioned earlier, our construction produces a number of cells which is singly exponential in the dimension (as a function of n), and is thus a noticeable improvement over the doubly exponential size of Collins' decomposition [22]. Of course, the purpose of Collins' construction is different from ours, since it is designed as a decision procedure for the first-order theory of real-closed fields. Incidentally,

our algorithm can decide the existential restriction of that theory, albeit not as fast as in [11, 42]. One drawback of our method is that, like Collins', it generates polynomials of degree doubly exponential in the dimension. Reducing this bound to singly exponential is a challenging open problem.

Applying this stratification technique in conjunction with the random sampling approach, we obtain an efficient point location algorithm that can answer any query in $O(\log n)$ time, using $O(n^{2d-2+\varepsilon})$ space, in dimension $d > 1$, for any fixed $\varepsilon > 0$. The preprocessing time is $O(n^{2d+1})$ time (deterministic) and $O(n^{2d-2+\varepsilon})$ (randomized). These bounds assume that the coefficients of each input polynomial f_i , as well as of certain auxiliary polynomials derived by the construction, can be stored in a single computer word and that arithmetic operations on word-size integers can be performed in constant time. To obtain an upper bound on the bit complexity of the algorithm we must multiply both preprocessing and query times by a polynomial in the maximum number of bits required to encode any coefficient in the f_i 's.

Our result is a substantial improvement over the best previous algorithm, which requires storage doubly exponential in the dimension; namely, $O(n^{2^d-1})$ [16]. Many algorithms have been given for searching among curves in two-dimensions [21, 28, 44]. See also [26, 39] for background information.

Point location among algebraic varieties is at the center of subquadratic algorithms for many optimization problems. By straight substitution of our techniques we improve upon all these algorithms at once. Here are a few examples among many others:

- (1) Computing the minimum vertical separation between two sets of line segments in 3-space [37].
- (2) Computing the longest line segment which fits inside a simple polygon [37].
- (3) Computing the time at which the convex hull of a set of points in (polynomial) motion enters its steady-state [5].
- (4) Given m red objects (algebraic curves, surface patches, etc.) and n blue objects, does any red object intersect any blue object? (A generalization of *Hopcroft's problem*).
- (5) Given m rays and n triangles in 3-space, find the first triangle hit by each of the rays, or alternatively, find the number of triangles stabbed by each ray [16].

This paper is organized as follows. In the next two sections we discuss our stratification technique and we introduce the key notion of a semi-cylindrical cell decomposition. We discuss point location in Section 4 and mention some applications of our techniques in Section 5. To preserve the flow of the presentation, all the proofs that are not essential for the understanding of the overall discussion have been relegated to an appendix.

2. Preliminaries

We recall some standard terminology and introduce some of the basic concepts to be used later. In particular, we define a sign-invariant stratification formally, and

we discuss the notion of a cylindrical cell and its upper boundary. Finally, we review the algebraic tools needed for eliminating variables from polynomials, and in particular, the fundamental theorem of subresultant theory.

Let $Q_d = Q[x_1, \dots, x_d]$ be the ring of polynomials with rational coefficients. A subset of \mathfrak{R}^d is a *semi-algebraic set* if it can be derived from sets of the form $\{x \in \mathfrak{R}^d \mid f(x) \geq 0\}$, where $f \in Q_d$, by union, intersection, and complementation. It is a classical result that any semi-algebraic set in \mathfrak{R}^d can be partitioned into manifolds of dimension between 0 and d [49]. Such a partition is called a *stratification*; its elements are called *strata*. It is immediate that the n -fold product of the stratification of \mathfrak{R} given by $(-\infty, 0)$, $\{0\}$, and $(0, +\infty)$ is itself a stratification of \mathfrak{R}^n : its strata are called *sign-sequences*. Given a polynomial map $F = (f_1, \dots, f_n): \mathfrak{R}^d \mapsto \mathfrak{R}^n$, where each $f_i \in Q_d$, the preimage $F^{-1}(\sigma)$ of a sign-sequence σ is called a *maximal sign-invariant set*. A stratification of \mathfrak{R}^d is *sign-invariant* for F if each stratum is a subset of a maximal sign-invariant set.

Let us make a few remarks to clarify these concepts. It should be clear that the collection S of maximal sign-invariant sets need not always be a stratification. For example, let $d = 2$ and $F = (f_1)$, with $f_1(x, y) = xy$. The variety $\{(x, y) \in \mathfrak{R}^2 \mid f_1(x, y) = 0\}$ belongs to S , but it contains the critical point $(0, 0)$ and thus fails to be a stratum. Interestingly, however, perturbing f_1 into $f_1^\varepsilon = f_1 + \varepsilon$, for almost any $\varepsilon \neq 0$, ensures that the variety $f_1^\varepsilon(x, y) = 0$ consists of regular points, and hence, is a 1-manifold¹. In general, it follows from Sard's Theorem [47] that the values of $F(x)$ are all regular, except for a zero-measure subset of \mathfrak{R}^n . Consequently, for almost any change of F into $F + \varepsilon$, where $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n) \in \mathfrak{R}^n$, the perturbed variety

$$\{x \in \mathfrak{R}^d \mid f_{i_1}(x) + \varepsilon_{i_1} = 0, \dots, f_{i_k}(x) + \varepsilon_{i_k} = 0\},$$

for any $k \leq d$, is a $(d - k)$ -manifold. (This means, for example, that a randomly perturbed polynomial curve in \mathfrak{R}^2 does not self-intersect.) It follows trivially that each maximal sign-invariant set is now a manifold. Thus, if S is not a stratification to begin with, almost any perturbation in the constant terms of the n coordinate polynomials of F will make it into one. Although not essential for our theory, this might be a useful tool in practice.

The main tool behind our data structure for point location is a new constructive proof that semi-algebraic sets admit sign-invariant stratifications. A crucial feature of the construction is that each stratum is a semi-algebraic set which can be defined by a constant number of polynomials of Q_d . We call such a set a *Tarski cell*. This can be regarded as a first step towards triangulating real-algebraic varieties. What will be lacking in our construction, however, is that our Tarski cells do not "glue" properly to one another to form a cell complex [45].

A *cylindrical cell* of \mathfrak{R} is either a singleton $\{a\}$, where a is real-algebraic, or an open interval (a, b) , where a and b are real-algebraic or $\pm\infty$. The *upper boundary* of the cell c , abbreviated $ubd(c)$, is $\{a\}$ in the first case and $\{b\}$ in the second case.

¹ Throughout this paper, unless specified otherwise, the term manifold will refer to a smooth manifold without boundary.

If $b = +\infty$, however, the upper boundary of c is not defined. Given $x = (x_1, \dots, x_{d-1}) \in \mathfrak{R}^{d-1}$ and $Y \subseteq \mathfrak{R}$, the set $\{(x_1, \dots, x_{d-1}, y) \mid y \in Y\}$ is denoted $x \otimes Y$. If $k > 1$, a cylindrical cell of \mathfrak{R}^k falls in one of the five categories below, where c' is a cylindrical cell of \mathfrak{R}^{k-1} , and f, g are real-valued smooth (i.e., infinitely differentiable) functions over c' :

- (i) $c = \bigcup \{x \otimes (f(x), g(x)) \mid x \in c'\}$, where c' is a cylindrical cell of \mathfrak{R}^{k-1} , and $f(x) < g(x)$ for all $x \in c'$. The upper boundary of c is $\bigcup \{x \otimes \{g(x)\} \mid x \in c'\}$.
- (ii) $c = \bigcup \{x \otimes (-\infty, g(x)) \mid x \in c'\}$ and $ubd(c) = \bigcup \{x \otimes \{g(x)\} \mid x \in c'\}$.
- (iii) $c = \bigcup \{x \otimes (f(x), +\infty) \mid x \in c'\}$; its upper boundary is not defined.
- (iv) $c = \bigcup \{x \otimes \mathfrak{R} \mid x \in c'\}$; its upper boundary is not defined.
- (v) $c = \bigcup \{x \otimes \{f(x)\} \mid x \in c'\}$ and $ubd(c) = \{c\}$.

The smoothness of f and g ensures that cylindrical cells and their upper boundaries (when defined) are connected smooth manifolds which admit single-chart bases [47] (meaning that they can be described by a single local parametrization). In the following the dimension of a cell will refer to the dimension of the corresponding manifold.

Lemma 2.1. *A cylindrical cell of \mathfrak{R}^d is a k -manifold ($k \leq d$) which can be parametrized by a single smooth diffeomorphism mapping the open unit ball U^k to the cell.*

Proof. See Appendix. \square

Lemma 2.2. *Whenever defined, the upper boundary of a cylindrical cell of dimension k (as a manifold) is a cylindrical cell of dimension k or $k-1$.*

Proof. Straightforward induction. \square

The notion of *upper boundary* allows us to define cell decompositions in a two-stage process: First, we pack \mathfrak{R}^d with cylindrical cells whose closures cover \mathfrak{R}^d ; then we complete the packing into a covering by adding on appropriate upper boundaries. We develop this idea in detail in the next Section.

We close these preliminaries with a short review of subresultant theory. Let $A(x) = \sum_{0 \leq i \leq a} \alpha_i x^i$ and $B(x) = \sum_{0 \leq i \leq b} \beta_i x^i$ be two polynomials with coefficients in Q or Q_d , (or actually in any integral unique-factorization domain with identity [48]), where $\alpha_a, \beta_b \neq 0$. From the unique factorization Theorem we easily find that $A(x)$ and $B(x)$ have at least one common divisor if and only if there exist two polynomials $U(x)$ and $V(x)$ of degree $b-1$ and $a-1$ respectively, which do not vanish identically, such that

$$U(x)A(x) = V(x)B(x). \quad (2.1)$$

Indeed, if the identity above is true then all the irreducible factors of $U(x)A(x)$ divide $V(x)B(x)$. But V is of degree too small to contain all the factors of A with

\mathfrak{R}^d , referred to as *K-decompositions*. Next, we examine every cell of every *K*-decomposition in turn, and keep only those that are free of intersections with any variety $\bigvee f_k$. These candidate cells might still be intersecting, so we add one final selection criterion based on the indices of their defining polynomials. This gives us a collection of mutually disjoint Tarski cells which, together with their upper boundaries, constitute the desired sign-invariant stratification of \mathfrak{R}^d .

The resulting stratification, denoted $\mathcal{S}_d(F)$, is called a *semi-cylindrical cell decomposition*. If $d = 1$, we have Collins' decomposition: The union of the n varieties $\bigvee f_i$ is a discrete set of real-algebraic numbers $\xi_1 < \xi_2 < \dots < \xi_k$, and $\mathcal{S}_1(F)$ consists of the cylindrical cells

$$(-\infty, \xi_1), \{\xi_1\}, (\xi_1, \xi_2), \dots, \{\xi_k\}, (\xi_k, +\infty).$$

To treat the general case we must define the intermediate *K*-decomposition $K(\varphi, \psi)$, where φ and ψ are two polynomials of Q_d . As we just outlined, the master plan is to identify the building blocks of $\mathcal{S}_d(F)$ among the cylindrical cells of $K(f_i, f_j)$, for all i, j ($i \leq j$).

Let's look at an example. Consider the four bivariate polynomials f_i ($1 \leq i \leq 4$) whose varieties, A, B, C, D , are shown in Fig. 1. Pairing A and B , we obtain the decomposition of \mathfrak{R} corresponding to the sequence of points and horizontal segments in Fig. 2. Lifting this decomposition in the vertical direction gives us our first *K*-decomposition (one should ignore the dashed curves in the figure). It consists of a collection of cylindrical cells. Let us restrict our attention to the two-dimensional cells that do not intersect any of the varieties A, B, C, D (dotted and hashed regions in Fig. 2). Some of the cells (the hashed regions) will be rediscovered during the pairings (A, A) and (B, B) , and are best ignored for the time being. The three dotted




Fig. 1.




Fig. 2.

regions are the two-dimensional cells which we keep once and for all as part of our final decomposition. We also add on their upper boundaries. The remaining cells are obtained by repeating the argument with the nine other pairings of A, B, C, D (Fig. 3). The labels of the regions indicate the pairings at which they are selected. Note that because of the junctions at a and b the final decomposition does not form a cell complex. These “faulty” junctions always occur at the bottom of vertical segments and not at the top because of our rule of adding upper and not lower boundaries. Of course, this problem is easy to fix in two dimensions, but it appears much more formidable in higher dimensions.

3.1. The K -decomposition

Let A be a polynomial in Q_d . Regarding A as a univariate polynomial with coefficients in the ring Q_{d-1} , we can write

$$A(x_1, \dots, x_d) = \sum_{0 \leq i \leq a} A_i(x_1, \dots, x_{d-1})x_d^i,$$

where A_a is not identically null. Following Collins' notation [22] we define $\deg(A) = a$ and $\text{ldcf}(A) = A_a(x_1, \dots, x_{d-1})$. For any k ($0 \leq k \leq a$) we also need the k th *reductum*

$$\text{red}^k(A) = \sum_{0 \leq i \leq a-k} A_i(x_1, \dots, x_{d-1})x_d^i.$$

Let G be the polynomial map whose coordinate functions are the nonzero polynomials in $\bigcup_{3 \leq i \leq 5} G_i$, where

- (i) $G_1 = \{\text{red}^k(g) \mid k \geq 0 \text{ and } \deg(\text{red}^k(g)) \geq 1 \text{ and } g \in \{\varphi, \psi, f_1, \dots, f_n\}\}$,
- (ii) $G_2 = \{\text{red}^k(g) \mid k \geq 0 \text{ and } \deg(\text{red}^k(g)) \geq 1 \text{ and } g \in \{\varphi, \psi\}\}$,




Fig. 3.

(iii) $G_3 = \{lcf(g) \mid g \in G_1\}$,

(iv) $G_4 = \{psc^k(g, \partial g / \partial x_d) \mid g \in G_1 \text{ and } 0 \leq k < deg(\partial g / \partial x_d)\}$,

(v) $G_5 = \{psc^k(f, g) \mid f \in G_2 \text{ and } g \in G_1 \text{ and } 0 \leq k < \min(deg(f), deg(g))\}$.

The reader familiar with Collins decomposition will recognize similarities in the variable elimination procedure. One crucial difference, however, is that all pairings here involve either φ or ψ , and are therefore considerably fewer. Regard each g as a univariate polynomial in x_d , so its coefficient domain is parametrized by a point in \Re^{d-1} . Roughly, (iv) delimits the regions of \Re^{d-1} where the number of real roots of each $g \in G_1$ changes, while (v) keeps track of where φ and ψ (and their reductae) acquire or lose common roots with each g . The reason for including (iii) is that changes in the number of roots might occur simply because of changes in the degree of g . (Actually, this slight annoyance can be avoided by applying a normalization procedure described in [40] for Collins' decomposition: The idea is to change coordinates so that each g receives a constant nonzero leading coefficient.)

We are now ready to construct $\mathcal{S}_{d-1}(G)$ recursively. At this point we must mention an assumption which we wish to make for the sake of convenience: Every polynomial $g(x_1, \dots, x_d)$ should be *well-based* [46], meaning that g , as a univariate polynomial in x_d , should never vanish identically. In other words, its coefficients in Q_{d-1} should never be all 0 simultaneously. Furthermore, this should also be true in all the recursive calls made by the algorithm. As it turns out, a random rotation in the coordinate axes ensures well-basedness with probability 1. We shall not elaborate on this issue, which is thoroughly discussed in [46].

Note that, unlike the base decomposition used in Collins' construction, $\mathcal{S}_{d-1}(G)$ is much too coarse to delineate the x_d -roots of each f_i . Still, since $\mathcal{S}_{d-1}(G)$ is sign-invariant for G , we have a form of partial delineation. To elaborate on this point, we need a few definitions. Let $x \in \mathfrak{R}^{d-1}$ and let $p(x; z) \in Q_{d-1}[z]$. Given a connected manifold $S \subseteq \mathfrak{R}^{d-1}$ we say that the functions $\{\zeta_i : S \rightarrow \mathfrak{R} \mid 1 \leq i \leq l\}$ delineate p over S if

- (i) each ζ_i is smooth over S ;
- (ii) for each $x \in S$, we have $\zeta_1(x) < \zeta_2(x) < \cdots < \zeta_l(x)$;
- (iii) for each $k = 1, \dots, l$, there is an integer m_k such that, for each $x \in S$, $\zeta_k(x)$ is the k th largest *distinct* real root of p , and this root has multiplicity m_k ;
- (iv) for each $x \in S$, p has exactly l distinct real roots.

Note that the domain of ζ_i need not extend beyond S and that the functions trace only *distinct* real roots. Our definition of delineation differs from the standard one [22] in two minor aspects: ignoring complex-valued roots and requiring smoothness. Let us now substantiate our previous claim about partial delineation. We will eventually prove that the cells of $\mathcal{S}_d(F)$ are manifolds which are diffeomorphic to the k -dimensional unit ball U^k , so let us assume inductively that this is true of $\mathcal{S}_{d-1}(G)$ (the basis case being obvious), and that $\mathcal{S}_{d-1}(G)$ is a sign-invariant stratification for G . We also assume that $d > 1$.

Lemma 3.1. *The functions $\varphi, \psi, f_1, \dots, f_n$ can all be delineated over each cell of $\mathcal{S}_{d-1}(G)$.*

Proof. See Appendix. \square

Let g be the product of two polynomials r and s , where $r \in \{\varphi, \psi\}$ and $s \in \{\varphi, \psi, f_1, \dots, f_n\}$. We will now show that g is delineated over each cell c of $\mathcal{S}_{d-1}(G)$. From the proof of Lemma 3.1, it suffices to show that for g as a univariate polynomial in x_d , the number of distinct roots of $g(x_1, \dots, x_d)$ remains constant for each $(x_1, \dots, x_{d-1}) \in c$. We have $\deg(g) = \deg(r) + \deg(s)$, so G_2 ensures that the degree of g is invariant over c . Now what about root multiplicities? Since both r and s can be delineated over c the only thing to check is that the degree of the greatest common divisor of r and s (again as polynomials in x_d) is constant over c . But this is precisely what G_5 is there to ensure.

For a given $x \in c$ and polynomial $g(x, z)$, form the list of distinct real roots of g and merge together these lists for all g in $\{\varphi(x, z), \psi(x, z), f_1(x, z), \dots, f_n(x, z)\}$. We obtain a list of smooth functions $\rho_1(x) \leq \cdots \leq \rho_l(x)$. Since c delineates $\varphi \times g$ for any $g \in \{\psi, f_1, \dots, f_n\}$, the real-root functions associated with φ are strictly ordered among the others: This means that if ρ_i is associated with φ , then for all j , we have $\rho_i < \rho_j$, or $\rho_i = \rho_j$, or $\rho_i > \rho_j$ over the entire domain c . We refer to this property as *partial delineation*. Of course, the same applies to ψ . Now let $\bar{\rho}_1(x) < \cdots < \bar{\rho}_k(x)$

($k \leq l$) be the (distinct) real-root functions associated with $\varphi \times \psi$. Since these functions are smooth we can build a stack of cylindrical cells:

- (i) $\bigcup \{x \otimes (-\infty, \bar{\rho}_1(x)) \mid x \in c\}$,
- (ii) $\bigcup \{x \otimes (\bar{\rho}_k(x), +\infty) \mid x \in c\}$,
- (iii) $\bigcup \{x \otimes (\bar{\rho}_i(x), \bar{\rho}_{i+1}(x)) \mid x \in c\}$ ($1 \leq i < k$),
- (iv) $\bigcup \{x \otimes \{\bar{\rho}_i(x)\} \mid x \in c\}$ ($1 \leq i \leq k$).

Cells of type (i)–(iii) are called *layer cells*, whereas cells of type (iv) are called *section cells* (think of a birthday cake). Note that these notions are well defined because the polynomials are well-based. A remark which will have its importance later is that each section cell is the upper boundary of a unique layer cell. Collecting cells for all $c \in \mathcal{S}_{d-1}(G)$ forms the desired decomposition $K(\varphi, \psi)$. In light of Lemmas 2.1 and 3.1 it follows by induction that $K(\varphi, \psi)$ is a stratification of \mathfrak{R}^d into cylindrical cells. The lemma that follows describes the most useful property of layer cells, for our purposes. It is an immediate corollary of partial delineation. Roughly speaking, the lemma says that if we can poke a layer cell from floor to ceiling with a vertical segment that intersects none of the varieties in the middle, then the whole cell is itself free of intersections with the varieties.

Lemma 3.2. *Suppose that a layer cell $c \in K(\varphi, \psi)$ contains a point $x = (x_1, \dots, x_d)$ such that $g(x_1, \dots, x_{d-1}, z) \neq 0$, for any $g \in \{\varphi, \psi, f_1, \dots, f_n\}$ and any $z \in \mathfrak{R}$ satisfying $(x_1, \dots, x_{d-1}, z) \in c$. Then the same is true of any $x \in c$. Furthermore, given $x = (x_1, \dots, x_d) \in c$, the subset of functions g in $\{\varphi, \psi, f_1, \dots, f_n\}$ which contribute the next real root (as polynomials of $Q_{d-1}[z]$) larger (or smaller) than x_d is invariant for all $x \in c$.*

Let us now show that these cells are Tarski cells (i.e., admit constant size representations) and that an algebraic point can be computed for each of them. Again, we proceed by induction on the dimension d . Regarding the representation issue, it follows from the four cases listed above that all we need to show is that being the k th largest distinct real root of $\varphi_x(z) \times \psi_x(z)$ can be expressed by a quantifier-free formula involving only a constant (dependent on d) number of polynomials and Boolean connectives. This is quite obvious if we allow quantifiers [2] which is fine since we can use Collins' method afterwards to eliminate all the quantifiers. To compute an algebraic sample point in each cell is straightforward. As in [22], we lift an algebraic point $x \in c \in \mathcal{S}_{d-1}(G)$ into \mathfrak{R}^d by assigning to it the following sequence of x_d -coordinates:

$$\bar{\rho}_1(x) - 1, \bar{\rho}_1(x), \frac{\bar{\rho}_1(x) + \bar{\rho}_2(x)}{2}, \dots, \bar{\rho}_{k-1}(x), \frac{\bar{\rho}_{k-1}(x) + \bar{\rho}_k(x)}{2}, \bar{\rho}_k(x), \bar{\rho}_k(x) + 1.$$

Of course, the difficulty is to compare and do arithmetic with (recursively represented) real-algebraic numbers, [9, 23, 24, 29, 30, 35, 36, 43] for a discussion of this and related issues. A very short primer on real-algebraic numbers is given in the Appendix. We will have to come back to the subject later when we analyze the complexity of the algorithm.

3.2. The semi-cylindrical cell decomposition

We are now ready to assemble our semi-cylindrical cell decomposition. Given $F = (f_1, \dots, f_n) \in Q_d^n$, we begin by computing $K(f_i, f_j)$ for each pair i, j such that $1 \leq i \leq j \leq n$. Then we argue that the $\binom{n+1}{2}$ K -decompositions contain all the cells necessary to form $\mathcal{S}_d(F)$. The only problem is finding the right cells. Intuitively, we would like to include only layer cells that are not crossed by any variety; collecting such cells over all pairs f_i, f_j will give us only “empty” layer cells, which, put together and glued to their upper boundaries, will yield the overall desired decomposition. Some caution must be used, however, to avoid accepting the same cell several times. This selection process is now described in detail. Let c be a layer cell of $K(f_i, f_j)$ and let $\alpha = (\alpha_1, \dots, \alpha_d) \in c$ be its precomputed sample point. Should c be accepted into $\mathcal{S}_d(F)$? To decide, we compute three sets of indices $L(\alpha), M(\alpha), U(\alpha)$. Let $z_1 \leq \dots \leq z_l$ be the real roots of the univariate polynomials $f_k(\alpha_1, \dots, \alpha_{d-1}, z)$ ($1 \leq k \leq n$), where each z_i is associated with a unique f_k . We partition the sequence of roots into blocks B_1, B_2, \dots of equal value. Thus, all z_k 's in B_1 are equal and strictly less than the roots in B_2 , etc. Now let B_m be the block (if one exists) whose corresponding root value is precisely α_d . We define $M(\alpha)$ (resp. $L(\alpha)$ and $U(\alpha)$) as the set of indices associated with B_m (resp. B_{m-1} and B_{m+1}). If there is no such block B_m , then $M(\alpha)$ is empty and $L(\alpha)$ (resp. $U(\alpha)$) is the set of indices associated with the block whose root value is the one immediately smaller (resp. larger) than α_d . Note that any one of $L(\alpha), M(\alpha)$, or $U(\alpha)$ might be empty. Assume that all three sets have been computed. With the convention that $\min \emptyset = 1$, the inclusion rule for c is particularly simple: Accept c if and only if

$$M(\alpha) = \emptyset \quad \text{and} \quad \{i, j\} = \{\min L(\alpha), \min U(\alpha)\}. \quad (3.1)$$

(The minimization is used to ensure that no cell is accepted more than once.) To complete the construction of $\mathcal{S}_d(F)$, we simply throw in the upper boundaries of each layer cell accepted. This asymmetry justifies the name semi-cylindrical cell decomposition. The following falls straight out of Lemma 3.2.

Lemma 3.3. *Given a polynomial map $F = (f_1, \dots, f_n) \in Q_d^n$, the set $\mathcal{S}_d(F)$ is a sign-invariant stratification of \mathfrak{R}^d into Tarski cylindrical cells.*

Proof. It suffices to show that given $x \in \mathfrak{R}^d$ there is a unique cell c in $\mathcal{S}_d(F)$ that contains x . We begin with the case where $M(x) = \emptyset$. The key observation comes from Lemma 3.2: Given $x \in c$, the set $\{\min L(x), \min U(x)\}$ is invariant over c . This implies that the unique cell of $K(f_i, f_j)$ containing x , where $i = \min\{\min L(x), \min U(x)\}$ and $j = \max\{\min L(x), \min U(x)\}$, is also the unique cell of $\mathcal{S}_d(F)$ that contains x . Suppose now that $M(x) \neq \emptyset$. Then because of well-basedness, the point $y = x + (0, \dots, 0, -\varepsilon)$ satisfies $M(y) = \emptyset$, for any positive ε small enough. Therefore, it lies in a unique layer cell of $\mathcal{S}_d(F)$. The upper boundary of that cell is the unique cell of $\mathcal{S}_d(F)$ that contains x . \square

The reader is invited to check that Fig. 3 is indeed the decomposition resulting from the curves of Fig. 1. A few observations are in order. Why are not the two regions labelled AA at the bottom left merged together? The reason is that during the pairing (A, A) , the silhouette of B (and C for that matter) is projected in the vertical direction and causes this apparently useless separation. To avoid it might be tricky because silhouettes are sometimes needed for delineation, as shown in Fig. 4. Note that these two regions are discovered during the pairing (A, A) , where they are included in the final decomposition, but also during the pairings (A, B) , (A, C) , and (A, D) . Finally, the reader should pay particular attention to the “faulty” junctions a and b . What happens there is that the two-dimensional region labelled AB , incident upon these points, forces its upper boundary into the decomposition, but this clashes with the lower boundaries of the regions right above.




Fig. 4.

3.3. Trimming the stratifications in lower dimensions

Semi-cylindrical cell decompositions often contain many superfluous features: certain cells could be merged together and we would still have a sign-invariant stratification. As we already saw, Fig. 3 displays several examples of that. This is a phenomenon which seems difficult to avoid. As we will show in Section 3.4, our construction yields $O(n^{2d-2})$ cells, which is still far from the Thom–Milnor bound of $O(n^d)$ on the maximum number of sign-invariant components. It is possible to trim down the decomposition in two and three dimensions. The three-dimensional case is quite complicated, however, and yields only modest savings, so we will only discuss the trimming process in two dimensions.

We begin with a brief review of Collins’ decomposition in two dimensions. Let $F = (f_1, \dots, f_n)$ be a polynomial map in Q_2^n and let (x, y) be a Cartesian system of coordinates. A *cylindrical algebraic decomposition* for the polynomials f_1, \dots, f_n , or *cad* for short [22], is defined by considering the *projection set* $C = \bigcup_{2 \leq i \leq n} C_i$, where

- (i) $C_1 = \{\text{red}^k(f_i) \mid k \geq 0 \text{ and } \deg(\text{red}^k(f_i)) \geq 1 \text{ and } 1 \leq i \leq n\}$,
- (ii) $C_2 = \{\text{ldef}(g) \mid g \in C_1\}$,
- (iii) $C_3 = \{\text{psc}^k(g, \partial g / \partial y) \mid g \in C_1 \text{ and } 0 \leq k < \deg(\partial g / \partial y)\}$,
- (iv) $C_4 = \{\text{psc}^k(f, g) \mid f, g \in C_1 \text{ and } 0 \leq k < \min(\deg(f), \deg(g))\}$.

This is the projection set in its most general form, so that generalizing it to higher dimensions is just a matter of substituting the right variables. As it turns out, reductae are not necessary in two dimensions, as our previous discussion on delineation should make clear. Indeed, each polynomial $ldcf(f_i)$ is univariate and therefore has a finite number of roots. Since delineation among these roots will be ensured, anyway, the reductae become irrelevant. (We shall leave them in, however, for the sake of simplicity).

A *cad* of the real line for polynomials g_1, \dots, g_u is defined just like a semi-cylindrical cell decomposition for the polynomial map (g_1, \dots, g_u) . To define the *cad* for F (in two dimensions), we begin by computing a *cad* of \mathfrak{R} for C , which we will use as a *base decomposition*. Then we build cylindrical cells by lifting the cells of the one-dimensional decomposition, using the f_i 's to create sections. The process is exactly the same as if we tried to define a K -decomposition with respect to f_1, \dots, f_n using the one-dimensional *cad* as a base decomposition. We do not elaborate on Collins' construction any further and refer the reader to [22] for details. However, let us mention the useful fact, proven in [46], that because of well-basedness, a *cad* is a cell complex.

We define a *vertical edge* to be any one-dimensional layer cell. Similarly, a *vertex* is a 0-dimensional cell. Next, we set out to remove extraneous vertical edges. To do so, we need adjacency information about the *cad*, which we obtain by computing all cell incidences. There are several ways to do that. For example, Schwartz and Sharir [46] give a method for determining into how many real roots a given root function splits, as we move from a cell to one next to it (which is the key question for determining incidences among the cells of a Collins decomposition). Given a real root ρ of $\varphi(x, z) \in Q_1[z]$, what happens to it as x moves to $x + \varepsilon v$, where v is a vector pointing towards the next cell, and ρ splits into several roots? For each new root z , we can express $z - \rho$ by a fractional power series in ε . A method is then needed to assess how many terms must be computed to be able to count the number of splits. This leads to a polynomial-time algorithm for computing incidences between cells of codimension 0 and 1. Using a different approach based on certain gap Theorems for real-algebraic numbers, Prill [40] gives a general polynomial-time algorithm for computing cell incidences. The rough idea is to compute approximate sample points for the cells and test incidence between two cells by checking how close their sample points are. The key here is to prove that points need not be too close and that fairly coarse approximations can be used. In our case, however, we can avoid many of these difficulties by using a simple procedure from [3] which is tailored for two dimensions and relies only on root isolation. The gist of the method is to enclose each critical point in a box small enough so that all the branches at that point cross the same vertical side. See also [33]. Other techniques for analyzing the topology of real-algebraic curves (which is what the discussion above is all about) are given in [24, 29, 43].

We now return to our main objective, which is to characterize the necessary vertices and edges and set out to eliminate all the others. We must assume that all

the cell incidences of the *cad* have been computed. We say that a point $(a, b) \in \mathbb{R}^2$ is *proper* if

- (1) $f_i(a, b) = \text{ldef}(\text{red}^k(f_i))(a) = 0$, for some i ($1 \leq i \leq n$) and some $k \geq 0$ such that $\text{deg}(\text{red}^k(f_i)) \geq 1$, or
- (2) $f_i(a, b) = \text{psc}^l(g, \partial g / \partial y)(a) = 0$ and $g = \text{red}^k(f_i)$, for some i, k, l such that $1 \leq i \leq n, k \geq 0$, and $0 \leq l < \text{deg}(\partial g / \partial y)$, or
- (3) $f_i(a, b) = f_j(a, b) = 0$, for some i, j ($1 \leq i, j \leq n$).

We extend case (1) to the points at infinity along asymptotic branches. Figure 5 depicts proper vertices of all three types. Case (1) shows two proper points of type (1), one of which is at infinity. We shall now remove every vertical edge of the *cad* that is not incident upon at least one proper vertex (possibly at infinity). An example is given in Fig. 6. Because the edges removed do not delineate any function locally, the natural variant of Lemma 3.2 still holds. That is, given any layer cell c of the new *cad*, the functions f_1, \dots, f_n which contribute the next real root larger (or smaller) than y are the same for all $(x, y) \in c$. Similarly, any point in a given section cell is the zero of the same subset of f_i 's. Note that the order of removal does not matter. (One might also observe that this cleanup will not always produce a minimal set of vertical edges: Indeed, edges might still remain which play no role in the delineation process.) Identifying edges to be removed can be done directly on the basis of the information provided by the cell incidence algorithms mentioned earlier. Similarly, repairing the decomposition (e.g., merging edges adjacent to a removed edge) involves only straightforward local surgery, once incidences are known. It is a simple exercise to show that the edge removal keeps all the cells cylindrical and, in particular, maintains their smooth differential structure. This completes our discussion of the sign-invariant semi-cylindrical cell decomposition of \mathbb{R}^2 for F , or




Fig. 5.




Fig. 6.

2-*scd* for short. Since the number of proper vertices is $O(n^2)$, a simple planarity argument shows that the number of resulting cells is $O(n^2)$ as well.

3.4. Complexity analysis

The combinatorial complexity of $\mathcal{S}_d(F)$ obeys a simple recurrence relation. Let $c(d, b, n)$ be the maximum number of cells in $\mathcal{S}_d(F)$, given that $F = (f_1, \dots, f_n)$ and each $f_i \in Q_d$ has degree at most b . The size of $\bigcup_{3 \leq i \leq 5} G_i$ does not exceed $2b^3n + b^2n$ and, because the subresultants we use are determinants of size at most $2b$ by $2b$, their maximum degree is at most $2b^2$. Consequently, we have $c(1, b, n) = O(n)$ and

$$c(d, b, n) \leq (4b+1) \binom{n+1}{2} c(d-1, 2b^2, (2b+1)b^2n), \quad \text{for } d > 1. \quad (3.2)$$

This recurrence is very conservative, so let us look more closely at the case $d = 2$. In particular, let us estimate the number E of edges in $\mathcal{S}_2(F)$ when b is considered a constant. This will give us an asymptotic upper bound on the total number of cells. We have $E = E_s + E_v$, where E_s counts the section edges and E_v the vertical edges. The closure of every vertical edge contains at least one proper point and there are $O(n^2)$ proper points, so $E_v = O(n^2)$. Since, obviously, $E_s = E + O(n^2)$, we derive $c(2, b, n) = O(n^2)$, in the case where b is a constant.

Resolving the recurrence in (3.2) we find that for any $d \geq 2$, $c(d, b, n) = O(n^{2d-2})$. Note that if $b = 1$ (the linear case) then we can use simpler and more efficient methods (e.g., Clarkson [17], Edelsbrunner [26]), which produce only $O(n^d)$ cells. Let l be the maximum *norm-length* of the f_i 's, that is,

$$l = \max_{1 \leq i \leq n} [\log(w(f_i) + 1)].$$

It follows from Collins' analysis that the norm-length of any intermediate polynomial is at most $O(1)$, if we take b to be a constant and assume that a computer word is at least l bits long. Similarly, encoding the sample points will require $O(1)$ words per point. An important remark is that although we can assume that b and d are fixed constants, we cannot extend this to l . Indeed, treating l as a constant would limit the maximum number of distinct polynomials to a constant: not a very wise thing to do!

The preprocessing time $t(d, b, n)$ follows a recurrence similar to (3.2). Up to within a constant factor, we have

$$\begin{aligned} t(d, b, n) \leq & \binom{n+1}{2} t(d-1, 2b^2, (2b+1)b^2n) \\ & + (4b+1) \binom{n+1}{2} c(d-1, 2b^2, (2b+1)b^2n) h(d, b, n), \end{aligned}$$

where $h(d, b, n)$ is the time for checking whether a cell of a K -decomposition of \mathfrak{R}^d should be accepted in the semi-cylindrical cell decomposition. For simplicity,

we will only count the number of word operations. Since the norm-length of all intermediate polynomials remains linear in the maximum norm-length l of the input polynomials, the bit complexity of the preprocessing will differ from our measure by at most a polynomial in l . As usual, we assume that b and d are constants. There are two (related) points to be discussed: (i) computing sample points and (ii) testing acceptance of a cell into $\mathcal{S}_d(F)$.

Recall that the data structure must provide a precomputed algebraic point in each cell of the semi-cylindrical cell decomposition. We have already seen how to specify these points, but we have not said anything about representation. The obvious solution is to use a recursive specification of real-algebraic numbers. One problem with that approach, however, is that an operation as simple as comparing two algebraic reals becomes a major challenge. Instead, we follow the approach of Collins [22] which is intimately based on Rubald's methods for computing in algebraic extension fields without requiring minimum defining polynomials. Collins' approach works fine when computing samples, but it does not fare nearly as well when testing cell acceptance. The reason is that it tends to make the asymptotic cost too heavily dependent on n , as opposed to the other parameters b, d (which we like to regard as constants). Fortunately, it is not too difficult to fix these problems.

Without loss of generality, we will consider the representation of a sample point $(\alpha_1, \dots, \alpha_d)$ of $K(f_1, f_2)$. The point is specified by lifting into \mathfrak{R}^d the (recursively computed) algebraic point $(\alpha_1, \dots, \alpha_{d-1})$, which itself has been computed recursively from some other K -decomposition of lesser dimension. From now on, we say that a real-algebraic number is *isolated* if it is expressed as the unique distinct real root in a rational interval of some primitive squarefree integral polynomial². We assume that α_1 has been isolated. Let $Q(\alpha_1, \dots, \alpha_i)$ denote the multiple real-algebraic extension field obtained by adjoining $\alpha_1, \dots, \alpha_i$ to Q . We shall inductively assume that $Q(\alpha_1, \dots, \alpha_{d-1})$ has been reduced to a simple extension field $Q(\hat{\alpha})$ and that $\hat{\alpha}$ has been isolated. We also assume that each α_i ($1 \leq i < d$) is expressed as $A_i(\hat{\alpha})$, where A_i is an integral polynomial.

For each $i = 1, 2$, let $\varphi_i(z)$ be the univariate polynomial $f_i(\alpha_1, \dots, \alpha_{d-1}, z)$ with coefficients in $Q(\hat{\alpha})$. First, we compute a coarsest square-free basis $\Psi = \{\psi_1, \dots, \psi_m\}$ for $\{\varphi_1, \varphi_2\}$. Next, we compute a list of distinct open rational intervals I_1, \dots, I_ν , along with a list of indices μ_1, \dots, μ_ν , such that (i) $I_1 < \dots < I_\nu$, (ii) each I_j contains one real root of ψ_{μ_j} and (iii) each distinct real root of $\prod_{1 \leq i \leq m} \psi_i$ lies in a distinct I_j . After this root isolation process, we must redefine the real roots by means of

² We recall some standard terminology. An integral univariate polynomial p is *primitive* if its coefficients are relatively prime. If they are not, their greatest common divisor is called the *contents* of the polynomial; factoring out the contents from each coefficient gives the *primitive part* of p . These notions generalize trivially to any unique-factorization domain. Given a set P of primitive polynomials, a *basis* B for P is a set of primitive polynomials of positive degree, pairwise relatively prime, such that (i) any $b \in B$ divides at least one polynomial of P and (ii) any $p \in P$ can be expressed as a product of polynomials in B . If P is arbitrary, then its basis consists of the contents of its polynomials along with the basis of their primitive parts. Finally, it is well-known that P always admits a *coarsest basis* B^+ , in the sense that any element of any basis for P divides some element of B^+ .

polynomials with *integral* coefficients. For each ψ_i , retrieve the intervals I_{h_1}, \dots, I_{h_v} which isolate its own real roots and compute a nonzero primitive square-free integral polynomial $\hat{\psi}_i$ as well as a sequence of nonoverlapping intervals $\hat{I}_{h_1}, \dots, \hat{I}_{h_v}$ such that, for each j ($1 \leq j \leq v$) $\hat{I}_{h_j} \subseteq I_{h_j}$ and the unique root of ψ_i in I_{h_j} is also the unique root of $\hat{\psi}_i$ in \hat{I}_{h_j} . Finally, once we have merged all the intervals \hat{I} 's, it becomes trivial to express the d th coordinates of all the sample points lifted from $(\alpha_1, \dots, \alpha_{d-1})$ in $K(f_1, f_2)$. The sample points in the section cells are already fully specified. The other sample points (the midpoints in layer cells) follow readily; we omit the details. To maintain the induction invariant, we must now compute and isolate a new number $\hat{\alpha}$ for each sample point which we just computed. This is a case of reducing a real-algebraic extension field $Q(a, b)$ to a simple one $Q(c)$.

Collins [22] shows how to carry out each of the steps described above in time polynomial in the number ($=2$) of functions involved in the lifting and in the number and degrees of all the other polynomials. The latter quantities depend only on b and d , and therefore are $O(1)$ for our purposes. The function $h(d, b, n)$ measures the worst-case time complexity of the following problem. Given an algebraic point $(\alpha_1, \dots, \alpha_d)$, let $\varphi_i(z)$ be the univariate polynomial $f_i(\alpha_1, \dots, \alpha_{d-1}, z)$ ($1 \leq i \leq n$) and let $\rho_1 < \dots < \rho_u$ be the distinct real roots of all the φ_i 's in increasing order: find which φ_i 's (if any) contribute ρ_k , where $\rho_{k-1} < \alpha_d \leq \rho_k$. Clearly, we can extend the previous technique to solve this problem, by simply substituting $\{f_1, \dots, f_n\}$ for $\{f_1, f_2\}$. The running time of this method would not be linear in n , however, so we slightly modify it. From our previous discussion we know that we can isolate (and thus compare) the real roots of any two polynomials φ_i and φ_j . Similarly, we can compare α_d against the real roots of any φ_i . Since any of these tests requires constant time it is immediate that $h(d, b, n) = O(n)$.

Let us now return to $t(d, b, n)$. We claim that $t(1, b, n) = O(n \log n)$. In $O(n)$ time we can certainly isolate the real roots of each f_i individually. Our claim will now follow readily if we can prove that comparing the r th real root of f_i against the s th real root of f_j can be done in constant time. But this is clear, since we can isolate the roots of $f_i \times f_j$ in constant time. Thus we obtain the following recurrence: $t(1, b, n) = O(n \log n)$, and for $d > 1$,

$$t(d, b, n) \leq \binom{n+1}{2} t(d-1, 2b^2, (2b+1)b^2n) \\ + (4b+1)n \binom{n+1}{2} c(d-1, 2b^2, (2b+1)b^2n),$$

where $t(d, b, n)$ is measured up to within a constant factor. This gives us $t(d, b, n) = O(n^{2d-1} \log n)$.

Theorem 3.4. *Let $F = (f_1, \dots, f_n)$ be a polynomial map from \mathbb{R}^d to \mathbb{R}^n . Suppose that each f_i is a polynomial of degree at most b in $Q[x_1, \dots, x_d]$ (whose norm-length does not exceed the size of a computer word). It is possible to construct a sign-invariant*

stratification of \mathbb{R}^d for F consisting of $O(n^{2d-2})$ cylindrical Tarski cells, if $d \geq 2$. If $d = 1$ the number of cells is respectively $O(n)$ and $O(n^2)$. In all cases, the construction can be done in time $O(n^{2d-1} \log n)$. Within the same asymptotic cost we can also compute an algebraic point in each cell of the decomposition.

4. Point location among real-algebraic varieties

We are now ready to attack the problem of preprocessing the set of varieties $\bigvee f_1, \dots, \bigvee f_n$ to support fast point location. We use probabilistic divide-and-conquer in the sense of Clarkson [18]: We choose a small random sample of varieties and compute a semi-cylindrical cell decomposition compatible with them. Next, we recurse in each cell c , passing only the varieties that intersect c down the recursion. To locate a point, we perform an exhaustive search in the top cell decomposition and iterate this process in the cell that contains the query point. The success of this method depends on how evenly the n varieties intersect the cells of the decomposition. We can show that uniform random sampling ensures success with high probability. To make the construction deterministic we use the general derandomization technique of Chazelle and Friedman [14]. This requires a certain amount of formalism which we discuss below.

4.1. Geometric divide-and-conquer

Let r be a fixed integer parameter between 1 and n . Our first task is to show how to select r varieties among $\bigvee f_1, \dots, \bigvee f_n$ and set the ground for divide-and-conquer. To do so we must recall some terminology [14]. Let $H = (V, E)$ be a multi-hypergraph (E is a multiset of edges in 2^V) and let $\varphi: 2^V \mapsto 2^E$ be a map such that (i) $\varphi(V) = E$ and (ii) $W' \subseteq W \subseteq V$ implies $\varphi(W') \subseteq \varphi(W)$. The pair $(H; \varphi)$ is called a *frame*. It is said to be of dimension δ if δ is the smallest positive (constant) real such that, for each $W \subseteq V$, the size of $\{W \cap e \mid e \in \varphi(W)\}$ is at most $c|W|^\delta$, for some constant c . The ratio $\min\{|e|/|V| : e \in E\}$ is called the *threshold* of the frame. Finally, a subset R of r vertices is called an *r -cover* if it has a nonempty intersection with every edge of $\varphi(R)$.

Theorem 4.1 (Chazelle–Friedman [14]). *Consider a frame of dimension δ with n vertices and let $r \leq n$ be any integer larger than some fixed constant. If the threshold of the frame is at least $a(\log r)/r$, for some appropriate constant a , then it is possible to find an r -cover for the frame in $O(rn^{\delta+1})$ (deterministic) time. A random subset of r vertices (under the hypergeometric distribution) is an r -cover with probability larger than some constant.*

We will now establish the relationship between frames and the problem at hand. The basic idea is to construct a frame where the vertices are the varieties and the edges represent all possible cells of the K -decompositions used in the construction of $\mathcal{S}_d(F)$. The vertices contained in an edge denote the varieties interfering with

its associated cell. In this way, a cell is accepted into the semicylindrical cell decomposition if and only if its corresponding edge is empty. This will allow us to prove the following important fact.

Theorem 4.2. *Consider n real-algebraic varieties in \mathfrak{R}^d of degree at most b and assume that $d > 1$. Given any integer $r \leq n$ large enough, there exists a semi-cylindrical cell decomposition of size $O(r^{2d-2})$, each of whose cells intersects $O(n(\log r)/r)$ varieties. The preprocessing requires $O(rn^{2d+1})$ deterministic time or $O(nr^{2d-2} + r^{2d-1} \log r)$ expected (randomized) time.*

Let f_1, \dots, f_n be n polynomials of Q_d of degree at most b . Our first task is to define the notion of an *abstract cylindrical cell*. The idea is to take the recursive definition of a cell of $\mathcal{S}_d(F)$ and remove all acceptance tests from it. Let us consider a cell c of $\mathcal{S}_d(F)$ and retrace its recursive definition. To begin with, we define the cell c in reference to some $K(f_i, f_j)$ by lifting a cell $c' \subseteq \mathfrak{R}^{d-1}$ into d -space (and perhaps taking its upper boundary). The lifting can be entirely specified by indicating its level l_1 (i.e., as a real-root rank), which is an integer between 0 and $2b$. We can define c' similarly, except that the varieties have changed. Now, a variety can be specified by a polynomial of the form $ldcf(g)$, $psc^{l_2}(g, \partial g / \partial x_d)$, or $psc^{l_2}(f, g)$, where $f = red^{l_3}(f_i)$ or $red^{l_3}(f_j)$, and $g = red^{l_4}(f_k)$; each of the l_i 's is bounded by b , the maximum degree of the polynomials. By agreeing once and for all on a certain syntax, we can therefore specify the variety by means of the sequence (i, j, k) , called its *multi-index*, followed by $O(\log(b+1))$ *parameter bits*. Note that, strictly speaking, i and j are not both needed: they are included as a reminder of the “genesis” of the variety. In a similar manner, we can specify any variety at any level of the recursion by a multi-index consisting of up to $2d$ integers between 1 and n , followed by $O(\log \delta)$ parameter bits, where δ is the maximum degree of specified polynomials. Since the degree of any intermediate variety is bounded above by $b^{O(2^d)}$, we can similarly specify any cell c of $\mathcal{S}_d(F)$ combinatorially by providing a multi-index of size $2d$, followed by $O(2^d \log(b+1))$ parameter bits. Any cell used in the intermediate decompositions (of type K or semi-cylindrical) at any level of the recursion can be expressed in a similar manner. This set-up allows us to define *abstract cylindrical cells* by first-order sentences. To be accepted into $\mathcal{S}_d(F)$, such an abstract cell must pass two different types of tests: (i) it must specify a nonempty cylindrical cell, and (ii) it must pass the acceptance test at each level of the recursion, meaning that it must pass, its base cell must pass, the base cell of its base cell must pass, etc.

Let us follow the chronological sequence of tests (3.1) which an abstract cylindrical cell c with multi-index S has to pass in order to make it into $\mathcal{S}_d(F)$. Suppose that the k th test (which takes place in \mathfrak{R}^k) is the first one which fails. There are two ways of failing. One is an unconditional failure caused by S itself, meaning that even if the varieties specified in S were the only ones considered the cell would still fail. In that case we say that every variety $\bigvee f_1, \dots, \bigvee f_n$ is a *witness*. What may happen, however, is that the k th test fails because of varieties not specified by S .

In that case, the witness set consists of the minimal subset of varieties $\bigvee f_i$'s whose removal would let the cell pass all the tests and make it into $\mathcal{S}_d(F)$. To make this definition sound we must prove that such a set is unique.

Let c be an abstract cylindrical cell with multi-index S and let c_1, c_2, \dots, c_d be the sequence of cells leading to $c = c_d$ by successive lifting $\mathfrak{R} \mapsto \mathfrak{R}^2 \mapsto \dots \mapsto \mathfrak{R}^d$. At the k th test, let E_k be the set of varieties in \mathfrak{R}^k which cause c_k to fail. We easily argue that if Σ is the set of multi-indices of the varieties in E_1, \dots, E_d , then the witness set of c is precisely $\bigcup \{\sigma \setminus S \mid \sigma \in \Sigma\}$. Therefore, the witness set of an abstract cylindrical cell is uniquely defined.

Our next task is to construct an appropriate frame $\mathcal{F} = (H; \varphi)$, with $H = (V, E)$. We define V by putting the vertices in bijection with the n input varieties. Given a subset $S \subseteq V$ of size $2d$, let $\kappa(S)$ be the set of all abstract cylindrical cells with multi-index S . For any $W \subseteq V$, let $\varphi(W)$ be the set

$$\bigcup \{\kappa(S) \mid S \subseteq W \text{ and } |S| = 2d\}.$$

We define the edge set E by putting it in bijection with $\varphi(V)$ and making each edge consist exactly of its witness set. From now on, we will not distinguish between edges and abstract cells, or between vertices and varieties. We easily check that \mathcal{F} is a frame. As we observed earlier, an abstract cell can be specified combinatorially by its multi-index and $O(2^d \log(b+1))$ bits. This means that $|\kappa(S)|$ is at most on the order of b^{2^d} . We derive that the frame \mathcal{F} is of dimension $2d$, since given any $W \subseteq V$,

$$|\{W \cap e \mid e \in \varphi(W)\}| \leq |\varphi(W)| = O(|W|^{2^d}).$$

Let us remove all edges of H of size at most $an(\log r)/r$, for the value of a required for the application of Theorem 4.1. We are now ready to compute an r -cover for the frame, which we can do in deterministic time $O(rn^{2^d+1})$. Let R be the polynomial map in Q_d^r formed by the defining polynomials of the varieties in the r -cover, and let c be a cell of $\mathcal{S}_d(R)$. Obviously, the cell c has an edge $e \in E$ associated with it. We will now show that the size of e cannot exceed $an(\log r)/r$. If it did, indeed, there would be a variety f_i in both e and the r -cover. This would mean that f_i is in the witness set of c , when regarded as an abstract cylindrical cell defined with respect to R . But this would deny its membership in $\mathcal{S}_d(R)$, which is a contradiction. We have not mentioned the fact that the decomposition algorithm is different in two dimensions. It is easy to show that our claims still hold true, however. Computing $\mathcal{S}_d(R)$ takes $O(r^{2^d-1} \log r)$ deterministic time. If we pick the r varieties at random, it takes us $O(r^{2^d-1} \log r)$ to construct the semi-cylindrical cell decomposition and $O(nr^{2^d-2})$ time to check that it satisfies the desired properties. The proof of Theorem 4.2 is now complete.

4.2. Point location

We follow the approach which Clarkson used in the linear case [18] and bring in the new machinery we just built. Applying Theorem 4.2 for a fixed (but large)

value of r gives us a semi-cylindrical cell decomposition of size $O(r^{2d-2})$. For each cell $c \in \mathcal{S}_d(\mathcal{R})$, identify the subset $V(c) \subseteq \{\sqrt{f_1}, \dots, \sqrt{f_n}\}$ of varieties that intersect c . Each variety in $V(c)$ is a witness of c , therefore $|V(c)| \leq an(\log r)/r$. Now recurse with respect to each $V(c)$. (Do not try to clip the resulting decompositions within c .) Here is how a point location query is answered. First, locate the point among the cells of $\mathcal{S}_d(\mathcal{R})$ by exhaustive search. If the point is found to lie in one of the varieties specified by \mathcal{R} then we can stop. Otherwise, we recurse in the data structure associated with the cell containing the query point.

In light of the previous Section, it is easy to argue that the query-answering terminates after $O(\log n)$ word operations. A multiplicative factor polynomial in the norm-length of the input polynomials must be added to get the bit complexity. Assume that $d \geq 2$; the storage requirement $s(n)$ follows the recurrence $s(O(1)) = O(1)$ and

$$s(n) \leq cr^{2d-2}s(\lceil an(\log r)/r \rceil),$$

which gives

$$\log s(n) \leq \frac{(2d-2)\log r + O(1)}{\log r - \log(a \log r)} \log n,$$

or $s(n) = O(n^{2d-2+\epsilon})$, for any fixed $\epsilon > 0$. Similarly, the preprocessing time can be estimated at $O(n^{2d+1})$ (deterministic) and $O(n^{2d-2+\epsilon})$ (randomized).

Theorem 4.3. *Consider n real-algebraic varieties in \mathbb{R}^d ($d > 1$) of degree at most b . It is possible to perform point location among the varieties in $O(\log n)$ query time, using $O(n^{2d-2+\epsilon})$ space, for any fixed $\epsilon > 0$. The data structure can be constructed deterministically in $O(n^{2d+1})$ time, or by using a Las Vegas algorithm, in $O(n^{2d-2+\epsilon})$ expected time. These bounds assume that the coefficients of the polynomials defining the varieties are rationals that can be stored in a single computer word and that arithmetic operations on word-size integers can be performed in constant time. To obtain an upper bound on the bit complexity of the algorithm we must multiply both preprocessing and query times by a polynomial factor in the maximum number of bits required to encode any coefficient in the defining polynomials.*

5. Concluding remarks

Our point location method allows us to improve upon the solutions currently known for a wide variety of optimization problems. Some of these problems have been studied in Chazelle and Sharir [16] and we direct the reader to this reference for details. Examples of these problems are:

- (1) Computing the minimum vertical separation between two sets of line segments in 3-space.
- (2) Computing the longest line segment which fits inside a simple polygon.

- (3) Computing the time at which the convex hull of a set of points in (polynomial) motion enters its steady-state.
- (4) Given m red objects (algebraic curves, surface patches, etc.) and n blue objects, does any red object intersect any blue object?
- (5) Given m rays and n triangles in 3-space, find the first triangle hit by each of the rays, or alternatively, find the number of triangles stabbed by each ray.

In one way or the other all these problems can be reduced to a generic problem of the following kind. Given a collection of n blue “objects” (point, line, polygon, curve, algebraic surface, etc.) and n red objects, does some blue–red pair of objects interact in some predetermined manner? Each object is specified by a vector with a constant number of real coordinates and the interaction predicate is a constant-size formula in the unquantified first-order of the reals. If r is the maximum length of any vector then the problem can be solved in time at most proportional to $n^{2-1/O(2^r)}$. This assumes that point location among n varieties in d -space can be done in logarithmic time and $n^{O(2^r)}$ preprocessing. Plugging in our new point location result yields a slightly better subquadratic complexity, namely, $O(n^{2-1/O(r)})$.

This work leaves open three major problems: The first one is to obtain a triangulation and not a stratification of the manifolds. The second problem is to lower the space requirement to the Thom–Milnor bound of $O(n^d)$. Finally, it would be nice to be able to carry out the computations without generating polynomials whose degrees are doubly exponential in the number of variables.

Appendix

Lemma 2.1. *A cylindrical cell of \mathfrak{R}^d is a k -manifold ($k \leq d$) which can be parametrized by a single smooth diffeomorphism mapping the open unit ball U^k to the cell.*

Proof. We proceed by induction on the dimension of the ambient space. The one-dimensional case is trivial, so assume that $d > 1$. Of the five types of cells introduced in the definition it suffices to consider types (i) and (v). Assume that the cell c is of the form $\bigcup \{x \otimes (f(x), g(x)) \mid x \in c'\}$ (type (i)). By induction hypothesis, c' is a k -manifold, for some $k \leq d - 1$, and we assume that there is a smooth diffeomorphism $\varphi: U^{d-1} \mapsto \mathfrak{R}^{d-1}$, whose restriction to U^k parametrizes c' . Now, given $\tilde{u} = (u, \alpha) \in U^d$, with $u \in U^{d-1}$, let

$$\psi(\tilde{u}) = \left(\varphi(u), \frac{1}{2} \left(1 - \frac{\alpha}{\sqrt{1-|u|^2}} \right) f(\varphi(u)) + \frac{1}{2} \left(1 + \frac{\alpha}{\sqrt{1-|u|^2}} \right) g(\varphi(u)) \right).$$

We easily check that the Jacobian determinant $\Delta\psi$ is equal to

$$\left(\frac{g(\varphi(u)) - f(\varphi(u))}{2\sqrt{1-|u|^2}} \right) \Delta\varphi \neq 0.$$

From the Inverse Function Theorem, we derive that ψ is a smooth local diffeomorphism. Actually, it is now immediate that ψ globally immerses U^{k+1} into \mathfrak{R}^d . Its restriction to U^{k+1} parametrizes c (which is therefore a $(k+1)$ -manifold).

Consider now the case of the cell $c = \bigcup \{x \otimes \{f(x)\} \mid x \in c'\}$. As before, let $\varphi: U^{d-1} \mapsto \mathfrak{R}^{d-1}$, be a smooth diffeomorphism whose restriction to U^k parametrizes c' , for some $k \leq d-1$. Consider the map

$$\psi(\tilde{u}) = (\varphi(u), \alpha) + (0, f(\varphi(u))),$$

where $\tilde{u} = (u, \alpha) \in U^d$ and $u \in U^{d-1}$. We have $\Delta\psi_{\tilde{u}} = \Delta\varphi_u \neq 0$, so again by the Inverse Function Theorem, ψ is a smooth diffeomorphism whose restriction to U^k parametrizes c and c is a k -manifold. \square

Lemma 3.1. *The functions $\varphi, \psi, f_1, \dots, f_n$ can all be delineated over each cell of $\mathcal{S}_{d-1}(G)$.*

Proof. For definiteness, we will deal with φ only, but everything we will say applies to the other functions as well. Once again, we regard φ as a univariate polynomial $\varphi_x(x_d)$ in $Q_{d-1}[x_d]$. As we shall see we only need to look at a subset $H = H_2 \cup H_3$ of G 's coordinate functions, where

- (i) $H_1 = \{\text{red}^k(\varphi) \mid k \geq 0 \text{ and } \text{deg}(\text{red}^k(\varphi)) \geq 1\}$,
- (ii) $H_2 = \{\text{ldef}(g) \mid g \in H_1\}$,
- (iii) $H_3 = \{\text{psc}^k(g, \partial g / \partial x_d) \mid g \in H_1 \text{ and } 0 \leq k < \text{deg}(\partial g / \partial x_d)\}$.

We will repeatedly use the fact that $\mathcal{S}_{d-1}(G)$ is sign-invariant for the polynomial map induced by H . Let $c \in \mathcal{S}_{d-1}(G)$; because of the sign-invariance with respect to H_2 , $\text{deg}(\varphi)$ remains constant over c . Then H_1 contains a restriction g of φ to c , whose leading coefficient does not vanish anywhere in c . From the Fundamental Theorem of Algebra, it trivially follows that the number of distinct (real and complex) roots of g (as a polynomial in x_d) is equal to

$$\text{deg}(g) - \text{deg}(\text{GCD}(g, \partial g / \partial x_d)).$$

Consequently, the sign-invariance with respect to H_3 , combined with Lemma 2.3 proves that the number of distinct roots of $\varphi_x(x_d)$ is invariant over c .

Borrowing a technique from Schwartz and Sharir [46] we can establish the continuity of the roots of $\varphi_x(x_d)$ by expressing each of its roots as a ratio of line integrals. For completeness, let us rederive this result. Because of well-basedness, $\varphi_x(z)$ is not identically zero, so it can be written as $(z - z_0)^k \gamma_x(z)$, where z_0 is a root of $\varphi_x(z)$ of multiplicity k . Let us now regard z as a variable in the complex plane and let us choose a small circle Γ which encloses z_0 but no other root. Since z_0 is not a pole of $\gamma_x^{-1}(z)$, given any complex polynomial $w(z)$, we have

$$\begin{aligned} \int_{\Gamma} \frac{w(z)\varphi'_x(z)}{\varphi_x(z)} dz &= \int_{\Gamma} \frac{k w(z)}{z - z_0} dz + \int_{\Gamma} \frac{w(z)\gamma'_x(z)}{\gamma_x(z)} dz \\ &= \int_{\Gamma} \frac{k w(z)}{z - z_0} dz = 2\pi k w(z_0) i. \end{aligned}$$

Setting $w(z) \stackrel{\text{def}}{=} z$ and $w(z) \stackrel{\text{def}}{=} 1$ successively, we derive

$$z_0 = \int_{\Gamma} \frac{z\varphi'_x(z)}{\varphi_x(z)} dz \Big/ \int_{\Gamma} \frac{\varphi'_x(z)}{\varphi_x(z)} dz,$$

which immediately establishes the continuity of z_0 as a function of x . Let us now show that the number of distinct *real* roots is also invariant over c . To see this, place small disjoint disks centered at each root of $\varphi_x(z)$. Note that because of disjointness the disks centered at the real roots are the only ones to intersect the real axis, the reason being that complex roots occur in conjugate pairs. For that same reason, a root cannot wander in and out of the real axis without changing the total count of distinct roots, therefore every real root x of c has a neighborhood in c composed entirely of real roots. Since c is connected the number of distinct real roots must therefore remain constant for all $x \in c$. To appreciate the importance of connectivity in this argument, consider the case $\varphi_x(z) = z^2 - x$, where $x \in \Re$, and assume that $c = (-1, 0) \cup (0, 1)$. Then $\varphi_x(z)$ always keeps two distinct roots over c , but both roots are real for $x = \frac{1}{2}$ and imaginary for $x = -\frac{1}{2}$. Of course, our algorithm would not allow such a cell c , since G would include the polynomial $g(x) = x$ as a coordinate function.

Returning now to our general discussion, we have established all the conditions for the delineation of φ_x , except for the smoothness of the real-root functions $\zeta_1(x) < \dots < \zeta_l(x)$. Before we do so we should note that, again because c is connected, the sign of $\varphi_x(z)$, for any z between $\zeta_j(x)$ and $\zeta_{j+1}(x)$, does not depend on x . To prove that each ζ_j is smooth, we will forsake Cauchy integrals and use a more general argument. Let (α, U) be a (smooth) coordinate chart around some arbitrary point of c . Given $u \in \alpha(c)$ and $z \in \Re$, let $\hat{\varphi}(u, z) = \varphi(\alpha^{-1}(u), z)$. Fix j ($1 \leq j \leq l$) once and for all and put $v = (u, \zeta_j(\alpha^{-1}(u)))$; by definition we know that $\hat{\varphi}(v) = 0$. Now let

$$m(u) = \max \left\{ k \geq 0 \mid \frac{\partial^k \hat{\varphi}}{\partial z^k}(v) = 0 \text{ and } \frac{\partial^{k+1} \hat{\varphi}}{\partial z^{k+1}}(v) \neq 0 \right\},$$

where ∂^0/∂ is the identity operator. Note that $m(u)$ is well defined unless $\varphi(\alpha^{-1}(u), z) = 0$, for all z . But this cannot happen because the input polynomials are well-based. Now, since

$$\frac{\partial^k \hat{\varphi}}{\partial z^k} = \frac{\partial^k \varphi}{\partial z^k},$$

we derive that $m(u) + 1$ is the multiplicity of the j th largest real root of $\varphi_{\alpha^{-1}(u)}(z)$, which we know remains constant over c . Let

$$w(u, z) = \frac{\partial^{m(u)} \hat{\varphi}}{\partial z^{m(u)}}(u, z).$$

Here is what we know about w : (i) it is smooth, (ii) $w(v) = 0$, and (iii) $(\partial w/\partial z)(v) \neq 0$. Then by the Implicit Function Theorem it follows that locally around v the equation

$w(u, z) = 0$ can be traced by a smooth function $z = z(u)$. The key observation now is that this function also traces $\varphi_{\alpha^{-1}(u)}(z) = 0$ around the point $(\alpha^{-1}(u), \zeta_j(\alpha^{-1}(u)))$. Consequently, this function is precisely $\zeta_j(\alpha^{-1}(u))$ and the j th largest distinct real root of $\varphi_{\alpha^{-1}(u)}(z)$ is a smooth function of u . \square

Remarks. The main motivation for proving that $\zeta_j(x)$ is smooth over c is to endow the cells of $\mathcal{S}_d(F)$ with a C^∞ differential structure (via Lemma 2.1). Note that although $\zeta_j(x)$ can be extended outside of c into a continuous function, it might not be possible to make this extension differentiable (let alone smooth) over the closure of c . For example, consider the torus $(\sqrt{x^2 + y^2} - 2)^2 + z^2 = 1$, whose polynomial equation is

$$\varphi_{(x,y)}(z) = (x^2 + y^2 + z^2 + 3)^2 - 16(x^2 + y^2) = 0.$$

The surface is obtained by revolving a vertical unit-circle centered at $(2, 0, 0)$ around the z -axis. The set

$$c = \{(x, y) \mid 1 < x < 3 \text{ and } 1 < x^2 + y^2 < 9\}$$

is an algebraic cell over which $\varphi_{(x,y)}(z)$ has two real roots. Now the reader should appreciate the difficulties in trying to extend, say, the second root

$$\zeta_2(x, y) = \sqrt{1 - (\sqrt{x^2 + y^2} - 2)^2}$$

smoothly to the closure of c . Note that the function does not have a partial in x at $(3, 0)$.

Algebraic Numbers. A standard representation of a real-algebraic number α consists of a pair $(P, [a, b])$, where P is a square-free polynomial with integer coordinates and $[a, b] \subseteq Q$ isolates α from the other real roots of P . Often we might be dealing with numbers in the extension field of α , which can then be expressed as quotients $A(\alpha)/B(\alpha)$, with $A, B \in Q_1$. Let us show briefly how the k th real root of P can be isolated in time polynomial in the degree of P and the logarithm of its weight. (The weight $w(P)$ of P is the sum of the magnitudes of its coefficients.) First, we can use Sturm sequences to compute the number of real roots in any interval $[a, b]$. This involves applying a straightforward variant of Euclid's GCD algorithm to the pair (P, P') and counting the sign changes in the resulting polynomial remainder sequences (evaluated at a and b). With this tool in hand, we can isolate the k th real root of P by binary search, starting with a large interval enclosing all the real roots, say $[-w(P), w(P)]$ and ending with an interval which is too small to enclose two distinct roots. A classical result of Mahler [36] says any two distinct real roots of P must be apart by at least $b^{-(b+2)/2} w(P)^{1-b}$. Consequently, the binary search will involve $O(b \log b + \log w(P) + 1)$ GCD computations, which proves that root isolation is polynomial. Collins and Loos [23] describe an efficient method for root isolation, whose bit complexity is $O(b^{10} + b^7 \log^3 w(P))$. Note that this discussion concerns only simple representations of real-algebraic numbers. For our purposes,

we must deal with algebraic numbers which are represented as roots of polynomials whose coefficients themselves are algebraic numbers represented recursively in the same manner [9, 23, 24, 29, 30, 35, 36, 43].

References

- [1] P. Agarwal, M. Sharir and P. Shor, Sharp upper and lower bounds on the length of general Davenport-Schinzel sequences, manuscript, 1988.
- [2] D.S. Arnon, Algorithms for the geometry of semi-algebraic sets, Tech. Rep. 436, Computer Science Dept., University of Wisconsin, Madison, 1981.
- [3] D.S. Arnon, G.E. Collins and S. McCallum, Cylindrical algebraic decomposition II: an adjacency algorithm for the plane, *SIAM J. Comput.* **13** (1984) 878–889.
- [4] B. Aronov and M. Sharir, Triangles in space, or building and analyzing castles in the air, in: *Proc. 4th Ann. ACM Sympos. Computational Geom.* (1988) 381–391.
- [5] M.J. Atallah, Dynamic computational geometry, *Comput. Math. Appl.* **11** (1985) 1171–1181.
- [6] R. Benedetti and J.J. Risler, On the number of connected components of a real algebraic set, Tech. Rept. LMENS-88-11, Ecole Normale Supérieure, Sept. 1988.
- [7] J. Bochnak, M. Coste and M.F. Roy, *Géométrie Algébrique Réelle* (Springer, Berlin, 1987).
- [8] W. Brown and J.F. Traub, On Euclid’s algorithm and the theory of subresultants, *J. ACM* **18** (1971) 505–514.
- [9] L. Caniglia, A. Galligo and J. Heintz, Some new effectivity bounds in computational geometry, in: *Proc. 6th Internat. Conf. on Applied Algebra, Algorithmic and Error Correcting Codes*, Rome (1988).
- [10] J.F. Canny, A new algebraic method for motion planning and real geometry, *Proc. 28th Ann. IEEE Symp. on Foundations of Computer Science* (1987) 39–48.
- [11] J.F. Canny, Some algebraic and geometric computations in PSPACE, *Proc. 20th Ann. ACM Symp. on Theory of Computability* (1988) 460–467.
- [12] B. Chazelle, Convex partitions of polyhedra: a lower bound and worst-case optimal algorithm, *SIAM J. Comput.* **13** (1984) 488–507.
- [13] B. Chazelle, Some techniques for geometric searching with implicit set representations, *Acta Inform.* **24** (1987) 565–582.
- [14] B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica* **10** (3) (1990) 229–249.
- [15] B. Chazelle and L. Palios, Triangulating a nonconvex polytope, *Discrete and Computational Geometry* **5** (1990) 505–526.
- [16] B. Chazelle and M. Sharir, An algorithm for generalized point location and its applications, *J. Symbolic Comput.* **10** (1990) 281–309.
- [17] K.L. Clarkson, A randomized algorithm for closest-point queries, *SIAM J. Comput.* **17** (1988) 830–847.
- [18] K.L. Clarkson, New applications of random sampling in computational geometry, *Discrete Comput. Geom.* **2** (1987) 195–222.
- [19] K.L. Clarkson, Applications of random sampling in computational geometry, II, in: *Proc. 4th Ann. ACM Sympos. Computational Geometry* (1988) 1–11.
- [20] K.L. Clarkson, H. Edelsbrunner, L.J. Guibas, M. Sharir and M. Welzl, Combinatorial complexity bounds for arrangements of curves and surfaces, in *Proc. 29th Ann. IEEE Symp. on Foundations of Computer Science* (1988) 568–579.
- [21] R. Cole, Searching and storing similar lists, *J. Algorithms* **7** (1986) 111–119.
- [22] G.E. Collins, Quantifier elimination for real closed fields by cylindrical algebraic decomposition, in: *Proc. 2nd GI Conf. on Automata Theory and Formal Languages*, Lecture Notes in Computer Science **33** (Springer, Berlin, 1975) 134–183.
- [23] G.E. Collins and R. Loos, Polynomial real root isolation by differentiation, in: *Proc. ACM Symp. on Symbolic and Algebraic Computations*, Yorktown Heights, NY (1976) 15–25.

- [24] 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. Symbolic Comput.* **5** (1988) 121–129.
- [25] J. Davenport and J. Heintz, Real quantifier elimination is doubly exponential, *J. Symbolic Comput.* **5** (1988) 29–35.
- [26] H. Edelsbrunner, *Algorithms in Combinatorial Geometry* (Springer, Heidelberg, Germany, 1987).
- [27] H. Edelsbrunner, L.J. Guibas and M. Sharir, The complexity of many faces in arrangements of lines and of segments, in: *Proc. 4th Ann. ACM Sympos. Computational Geometry* (1988) 44–55.
- [28] H. Edelsbrunner, L.J. Guibas and J. Stolfi, Optimal point location in a monotone subdivision, *SIAM J. Comput.* **15** (1986) 317–340.
- [29] P. Gianni and C. Traverso, Shape determination for real curves and surfaces, *Ann. Univ. Ferrara Sez. VII N.S.* **29** (1983) 87–109.
- [30] D. Grigor'ev and N. Vorobjov, Solving systems of polynomial inequalities in subexponential time, *J. Symbolic Comput.* **5** (1988) 37–64.
- [31] S. Hart and M. Sharir, Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes, *Combinatorica* **6** (1986) 151–177.
- [32] D. Haussler and E. Welzl, Epsilon-nets and simplex range queries, *Discrete Comput. Geom.* **2** (1987) 127–151.
- [33] D. Kozen and C. Yap, Algebraic cell decomposition in NC, in: *Proc. 26th Ann. IEEE Symp. on Foundations of Computer Science* (1985) 515–521.
- [34] R. Loos, Generalized polynomial remainder sequences, in: B. Buchberger, G. Collins, R. Loos, R. Albrecht, eds., *Computer Algebra: Symbolic and Algebraic Computation* (Springer, Berlin 1983).
- [35] R. Loos, Computing in algebraic extensions, in: B. Buchberger, G. Collins, R. Loos, R. Albrecht, eds., *Computer Algebra: Symbolic and Algebraic Computation* (Springer, Berlin, 1983).
- [36] K. Mahler, An inequality for the discriminant of a polynomial, *Michigan Math. J.* **11** (1964) 257–262.
- [37] M. McKenna, The biggest stick problem, in: *First Computational Geometry Day*, New York Univ., September 1986.
- [38] J. Milnor, On the Betti numbers of real varieties, *Proc. Amer. Math. Soc.* **15** (1964).
- [39] F.P. Preparata and M.I. Shamos, *Computational geometry: an introduction* (Springer, New York, 1985).
- [40] D. Prill, On approximations and incidence in cylindrical algebraic decompositions, *SIAM J. Comput.* **15** (1986) 972–993.
- [41] J.H. Reif and S. Sen, Optimal randomized parallel algorithms for computational geometry, in: *Proc. 16th Internat. Conf. Parallel Processing*, St. Charles, IL (1987; full version, Duke Univ., Tech. Rept. CS-88-01, 1988).
- [42] J. Renegar, A faster PSPACE algorithm for deciding the existential theory of the reals, in: *Proc. 29th Ann. IEEE Symp. on Foundations of Computer Science* (1988) 291–295.
- [43] M.F. Roy, Computation of the topology of a real algebraic curve, to appear in: *Proc. Congress on Computational topology and geometry*, Sevilla (1987).
- [44] N. Sarnak and R.E. Tarjan, Planar point location using persistent search trees, *Comm. ACM* **29** (1986) 669–679.
- [45] J.T. Schwartz, *Differential geometry and topology* (Gordon and Breach, London, 1968).
- [46] J.T. Schwartz and M. Sharir, On the “piano movers” problem. II: General techniques for computing topological properties of real algebraic manifolds, *Adv. in Appl. Math.* **4** (1983) 298–351.
- [47] M. Spivak, *A Comprehensive introduction to differential geometry, Vol. 1* (Publish or Perish, Berkeley).
- [48] B.L. van der Waerden, *Modern Algebra* (Ungar, New York, 1950).
- [49] H. Whitney, Elementary structure of real algebraic varieties, *Ann. of Math.* **66** (1957).
- [50] A.C. Yao, On constructing minimum spanning tree in k-dimensional space and related problems, *SIAM J. Comput.* **11** (1982) 721–736.