Записки научных
семинаров ПОМИ
Том 316, 2004 г.

S. Basu[1], **R. Pollack**[2], **M.-F. Roy**

COMPUTING THE DIMENSION OF A SEMI-ALGEBRAIC SET

Abstract. In this paper, we consider the problem of computing the real dimension of a given semi-algebraic subset of R^k, where R is a real closed field. We prove that the dimension, k', of a semi-algebraic set described by s polynomials of degree d in k variables can be computed in time

$$\begin{cases} s^{(k-k')k'} d^{O(k'(k-k'))} & \text{if } k' \geqslant k/2 \\ s^{(k-k'+1)(k'+1)} d^{O(k'(k-k'))} & \text{if } k' < k/2. \end{cases}$$

This result improves slightly the result proved in [22], where an algorithm with complexity bound $(sd)^{O(k'(k-k'))}$ is described for the same problem. The complexity bound of the algorithm described in this paper has a better dependence on the number, s, of polynomials in the input.

1. Introduction

Let R be a real closed field. Given a finite set $\mathcal{P}$ of s polynomials in $\mathrm{R}[X_1, \dots, X_k]$, a subset S of R^k is $\mathcal{P}$-semi-algebraic if S is the realization of a quantifier free formula with atoms $P = 0$, $P > 0$ or $P < 0$ with $P \in \mathcal{P}$. We call such a formula a $\mathcal{P}$-formula.

Computing topological invariants of a $\mathcal{P}$-semi-algebraic set S, such as deciding emptiness, computing the number of connected components, computing the Euler–Poincaré characteristic of S etc., has attracted a lot of attention in recent times. Using the method of cylindrical algebraic decomposition [12], all of these problems can be solved in time $(sd)^{2^{O(k)}}$, which is doubly exponential in k. However, singly exponential time algorithms are known for all these problems using the critical point method [15, 6, 7, 14, 17, 21, 16, 9, 13, 8, 18, 19, 2, 4, 1].

Amongst the problems that can be solved in single exponential time, there are certain ones for which there exist algorithms with singly exponential complexity bounds, where the exponent appearing in the bound

[1] Supported in part by NSF grant CCR-0049070 and an NSF Career Award 0133597
[2] Supported in part by NSF grant CCR-0098246.

is $O(k)$. Examples of such problems are, deciding the emptiness of a given semi-algebraic set, and computing the Euler–Poincaré characteristic of such a set. In contrast, there are certain other problems for which, even though singly exponential algorithms are known, the exponent in the bounds on the complexities of the most efficient algorithms is $O(k^2)$. The main examples of such problems, are computing the number of connected components of a given semi-algebraic set, as well as the problem of computing the dimension of such a set. It is currently not known whether the exponent, $O(k^2)$, in the complexity bounds of algorithms for either of these two problems can be improved to $O(k)$.

Note that, in several algorithms for computing properties of a given semi-algebraic set S, the bound on the complexity has a dependence on the dimension of S (for instance [3], see [5] for other examples). It is often assumed in such cases, that the dimension of the set is part of the input and is not computed by the algorithm. Thus, being able to compute the real dimension of a semi-algebraic set efficiently is an important problem in algorithmic semi-algebraic geometry. Computing the real dimension of a semi-algebraic set has also being considered from a complexity theoretic point of view and in [20] it is shown to be $\mathrm{NP}_{\mathbb{R}}$-complete in the Blum–Shub–Smale model of computation.

It is easy to see that the dimension of a semi-algebraic set S can be computed using cylindrical algebraic decomposition, since the dimension of S is the dimension of a maximal dimensional cell in such a decomposition of S. However, it is well known that the real dimension of S can in fact be computed in time $(sd)^{O(k^2)}$ (see [20]). On the other hand the complexity of computing the dimension of constructible sets over algebraically closed fields is $(sd)^{O(k)}$ [10].

Improving the exponent from $O(k^2)$ to $O(k)$ in the real closed case seems to be a difficult problem. The paper [22] obtained a partial result in this direction and gave an algorithm computing the dimension k' of a semi-algebraic set described by s polynomials of degree d in k variables, in time $(sd)^{O(k'(k-k')}$. Note that this complexity is output sensitive. Thus, for fixed $\ell \geq 0$, it is possible to decide in time $(sd)^{O(k)}$, whether the dimension of a given $\mathcal{P}$-semi-algebraic set, S, is $\leq \ell$. This generalizes to $\ell > 0$, the algorithm for checking whether the dimension of S is zero, described in [2].

In this paper, we improve slightly the results of [22] by giving more precise complexity bounds with respect to the number s of polynomials.

More precisely, we show that the dimension k' can be computed in time

$$\begin{cases} s^{(k-k')k'}d^{O(k'(k-k'))} & \text{if } k' \geqslant k/2 \\ s^{(k-k'+1)(k'+1)}d^{O(k'(k-k'))} & \text{if } k' < k/2. \end{cases}$$

The rest of the paper is organized as follows. In Sec. 2, we give some geometric properties of the dimension in terms of well chosen projections. In Sec. 3, we describe several algorithmic prerequisites which will be necessary for our algorithm. The main technique used, similarly to [22], is that of block elimination, which makes possible the projection of one block of variables efficiently in a single step. Block elimination technique, introduced for the first time by Dima Grigoriev in [14], is a powerful tool in several recent algorithms for real algebraic geometry (see [5]). Finally, in Sec. 4, we describe our algorithm for computing the dimension.

2. Geometric properties of dimension

Let S be a $\mathcal{P}$-semi-algebraic set where $\mathcal{P}$ is a finite set of s polynomials in $\mathrm{R}[X_1,\ldots,X_k]$. A subset T of R^k is $\mathcal{P}$-invariant if every polynomial $P\in\mathcal{P}$ has constant sign (>0, <0, or $=0$) on T. We denote by $\mathrm{Sign}(\mathcal{P})\subset\{0,1,-1\}^{\mathcal{P}}$ the set of all realizable sign conditions for $\mathcal{P}$ over R^k, i.e., those $\sigma\in\{0,1,-1\}^{\mathcal{P}}$ such that

$$\mathcal{R}(\sigma)=\{x\in\mathrm{R}^k \mid \bigwedge_{P\in\mathcal{P}} \mathrm{sign}\,(P(x))=\sigma(P)\}\neq\emptyset.$$

We denote by $\mathrm{SSign}(\mathcal{P})$ the set of realizable strict sign conditions of $\mathcal{P}$, i.e., the realizable sign conditions $\sigma\in\{0,1,-1\}^{\mathcal{P}}$ such that for every $P\in\mathcal{P}$, $P\neq 0$, $\sigma(P)\neq 0$.

Proposition 1. *The dimension of S is k if and only if there exists $\sigma\in$ $\mathrm{SSign}(\mathcal{P})$ such that $\mathcal{R}(\sigma)\subset S$.*

Proof. The dimension of S is k if and only if there exists a point $x\in S$ and $r>0$ such that $B(x,r)\subset S$. The sign condition satisfied by $\mathcal{P}$ at such an x is necessarily strict. In the other direction, if the sign condition σ satisfied by $\mathcal{P}$ at such an x is strict, $\mathcal{R}(\sigma)$ is open, and contained in S since S is defined by a quantifier free $\mathcal{P}$-formula. □

It is reasonable to expect that the dimension of S is $\geqslant j$ if and only if there exists a linear surjection π of R^k to R^j such that the dimension of $\pi(S)$ is j. We are going to prove that a stronger statement is true, namely that it suffices to consider $j(k-j)+1$ well chosen linear surjections, rather than all linear surjections. □

Notation 1. We denote by $v_k(x)$ the Vandermonde vector $(1, x, \ldots, x^{k-1})$. For $\ell \in \mathbb{N}$, we denote by V_ℓ the vector subspace of R^k generated by

$$v_k(\ell), v_k(\ell+1), \ldots, v_k(\ell+k-k'-1).$$

It is clear that V_ℓ is of dimension $k-k'$ since the matrix of coordinates of the vectors

$$v_{k-k'}(\ell), v_{k-k'}(\ell+1), \ldots, v_{k-k'}(\ell+k-k'-1)$$

is an invertible $(k-k') \times (k-k')$ Vandermonde matrix. We now describe equations for V_ℓ. Let, for $k-k'+1 \leqslant j \leqslant k$,

$$\begin{aligned} \overline{X}_j &= (X_1, \ldots, X_{k-k'}, X_j), \\ v_{k,j}(\ell) &= (1, \ldots, \ell^{k-k'-1}, \ell^{j-1}) \\ f_{\ell,j} &= \det(v_{k,j}(\ell), \ldots, v_{k,j}(\ell+k-k'-1), \overline{X}_j), \\ L_{k',\ell}(X_1, \ldots, X_k) &= (X_1, \ldots, X_{k-k'}, f_{\ell,k-k'+1}, \ldots, f_{\ell,k}). \end{aligned}$$

Note that the zero set of the linear forms $f_{\ell,j}$, for $k-k'+1 \leqslant j \leqslant k$ is the vector space V_ℓ and that $L_{k',\ell}$ is a linear bijection such that $L_{k',\ell}(V_\ell)$ consists of vectors of R^k having their last k' coordinates equal to 0. We also denote by $M_{k',\ell} = (d_{k-k',\ell})^{k'} L_{k',\ell}^{-1}$, where

$$d_{k-k',\ell} = \det(v_{k-k'}(\ell), \ldots, v_{k-k'}(\ell+k-k'-1)).$$

Note that $M_{k',\ell}$ plays the same role as the inverse of $L_{k',\ell}$ but is with integer coordinates, since, for $k-k'+1 \leqslant j \leqslant k$, $d_{k-k',\ell}$ is the coefficient of X_j in $f_{\ell,j}$.

For a family of polynomials $\mathcal{P} = \{P_1, \ldots, P_s\} \subset \mathrm{R}[X_1, \ldots, X_k]$, and a $k \times k$ matrix M, we denote by $\mathcal{P}(M) = \{P_1(M \cdot X), \ldots, P_s(M \cdot X)\}$

The following proposition is proved in [11].

Proposition 2. *Any linear subspace T of R^k of dimension $j \geqslant k'$ is such that there exists $0 \leqslant \ell \leqslant k'(k-k')$ such that V_ℓ and T span R^k.*

We denote by π_j the canonical projection of R^k to R^j forgetting the last $k-j$ coordinates.

Proposition 3. *Let $0 \leqslant j \leqslant k$. The dimension of S is $\geqslant j$ if and only if there exists $0 \leqslant i \leqslant j(k-j)$ such that the dimension of $\pi_j(L_{j,i}(S))$ is j.*

Proof. It is clear that if the dimension of $\pi_j(L_{j,i}(S))$ is j,the dimension of S is $\geqslant j$. In the other direction, if the dimension of S is $k' \geqslant j$, there exists a smooth point x of S of dimension k' with tangent space denoted by T (see [5, Proposition 5.54]). By Proposition 2, there exists $0 \leqslant i \leqslant j(k-j)$, such that V_i and T span R^k. Since $L_{j,i}(V_i)$ consists of vectors of R^k having their last j coordinates equal to 0, and $L_{j,i}(V_i)$ and $L_{j,i}(T)$ span R^k, $\pi_j(L_{j,i}(T))$ is R^j. Then the dimension of $\pi_j(L_{j,i}(S))$ is j. □

3. Algorithmic prerequisites

All the computations take place in an ordered domain D contained in a real closed field R such that D contains the coefficients of the elements of $\mathcal{P}$. We denote by K the quotient field of D and by $\mathrm{C} = \mathrm{R}[i]$ the algebraic closure of R.

The following notations and algorithms are used in the rest of the paper. We state only the inputs, outputs and the complexities of these algorithms, and refer to [5] for a precise description.

For $\mathcal{Q} \subset \mathrm{R}[X_1, \ldots, X_k]$ we denote the set of zeros of $\mathcal{Q}$ in R^k by $\mathrm{Z}(\mathcal{Q}, \mathrm{R}^k) = \{x \in \mathrm{R}^k \mid Q(x) = 0, Q \in \mathcal{Q}\}$.

A k-univariate representation is a $k+2$-tuple of polynomials of $\mathrm{K}[T]$,

$$(f(T), g_0(T), g_1(T), \ldots, g_k(T)),$$

such that f and g_0 are coprime. Note that $g_0(t) \neq 0$ if $t \in \mathrm{C}$ is a root of $f(T)$. The points associated to a univariate representation are the points

$$\left(\frac{g_1(t)}{g_0(t)}, \ldots, \frac{g_k(t)}{g_0(t)}\right) \in \mathrm{C}^k,$$

where $t \in \mathrm{C}$ is a root of $f(T)$. A real k-univariate representation is a pair u, σ where u is a k-univariate representation and σ is the Thom encoding of a root of f, $t_\sigma \in \mathrm{R}$. The Thom encoding of a root $t_\sigma \in \mathrm{R}$ is the list of signs of the derivatives of f at t_σ and it characterizes the root t_σ (see [5, Proposition 10.62]). The point associated to the real univariate representation u, σ is the point

$$\left(\frac{g_1(t_\sigma)}{g_0(t_\sigma)}, \ldots, \frac{g_k(t_\sigma)}{g_0(t_\sigma)}\right) \in \mathrm{R}^k.$$

Notation 2. Let $u = (f, g_0, \ldots, g_k) \subset \mathrm{K}[T]^{k+2}$ be a k-univariate representation and $P \in \mathrm{K}[X_1, \ldots, X_k]$. Set

$$P_u = g_0^e P\left(\frac{g_k}{g_0}, \ldots, \frac{g_k}{g_0}\right), \tag{1}$$

where e is the least even number not less than the degree of P.

The following algorithm enables us to compute a set of real univariate representations whose associated points meet the realizations of each realizable sign condition of a family of polynomials in $\mathrm{R}[X_1, \ldots, X_k]$.

Algorithm 1 (Sampling ([5], Algorithm 13.11))

Input: *a set of s polynomials,*

$$\mathcal{P} = \{P_1, \ldots, P_s\} \subset \mathrm{D}[X_1, \ldots, X_k],$$

each of degree at most d.

Output: *a set $\mathcal{U}$ of real univariate representations in $\mathrm{D}[T]^{k+2}$ such that the associated points form a set of sample points for $\mathcal{P}$ in R^k, meeting every semi-algebraically connected component of $\mathcal{R}(\sigma)$ for every $\sigma \in \mathrm{Sign}(\mathcal{P})$, and the signs of the elements of $\mathcal{P}$ at these points.*

Complexity: *The number of arithmetic operations in D is bounded by*

$$s \sum_{j \leqslant k} 4^j \binom{s}{k} d^{O(k)} = s^{k+1} d^{O(k)}.$$

However, the number of points actually constructed is only

$$\sum_{j \leqslant k} 4^j \binom{s}{j} O(d)^k.$$

If $\mathrm{D} = \mathbb{Z}$ and the bitsize of the coefficients of the input polynomials is τ, the size of the integer coefficients of the univariate representations in $\mathcal{U}$ are bounded by $\tau d^{O(k)}$.

We denote by $\mathrm{R}\langle\varepsilon\rangle$ the real closed field of algebraic Puiseux series where ε is a variable and by $\mathrm{R}\langle\varepsilon, \delta\rangle$ the field $\mathrm{R}\langle\varepsilon\rangle\langle\delta\rangle$. A parametrized univariate representation with parameters Y is a $k+2$-tuple

$$u(Y) = (f(Y,T), g_0(Y,T), \ldots, g_k(Y,T)) \in \mathrm{D}[Y][T]^{k+2}.$$

The following algorithm for efficiently eliminating a block of variables plays a key role in our algorithm.

Algorithm 2 (Block Elimination ([5], Algorithm 14.6))

Input: *a set of s polynomials $\mathcal{P} \subset \mathrm{D}[Y,X]$, with $Y = Y_1,\dots,Y_\ell$ and $X = X_1,\dots,X_k$ each of degree at most d.*

Output: *a set $\mathrm{BElim}_X(\mathcal{P}) \subset \mathrm{D}[Y_1,\dots,Y_\ell]$ such that $\mathrm{Sign}(\mathcal{P}(y, X_1,\dots, X_k))$ is fixed as y varies over a semi-algebraically connected component of a realizable sign condition of $\mathrm{BElim}_X(\mathcal{P})$.*

a set $\mathrm{U}_X(\mathcal{P})$ of parametrized univariate representations of the form

$$u(Y,\varepsilon,\delta) = (f, g_0, \dots, g_k),$$

where $f, g_i \in \mathrm{D}[Y,\varepsilon,\delta][T]$. The set $\mathrm{U}_X(\mathcal{P})$ has the property that for any point $y \in \mathrm{R}^\ell$, denoting by $\mathrm{U}_X(\mathcal{P})(y)$ the subset of $\mathrm{U}_X(\mathcal{P})$ such that $f(y,T)$ and $g_0(y,T)$ are coprime, the points associated to the univariate representations $u(y,\varepsilon,\delta)$ in $\mathrm{U}_X(\mathcal{P})(y)$ intersect every semi-algebraically connected component of every realizable sign condition of the set $\mathcal{P}(x)$ in $\mathrm{R}\langle\varepsilon,\delta\rangle^k$.

Complexity: *The number of arithmetic operations in D is*

$$s \sum_{i \leqslant j} 4^i \binom{s}{i} d^{O(k)} = s^{j+1} d^{O(k)}.$$

If $\mathrm{D} = \mathbb{Z}$, and the bitsizes of the coefficients of the polynomials are bounded by τ, then the bitsizes of the integers appearing in the intermediate computations and the output are bounded by $\tau d^{O(k)}$.

The following notations are adapted from that used in [5]. Here, we need them for the special case, when there are only two blocks of variables. Let $\mathcal{P}$ be a set of s polynomials in k variables $X_1,\dots,X_k$, and let Π denote a partition of the list of variables $X_1,\dots,X_k$ into two blocks, $X_{[1]}$, $X_{[2]}$. Let $\mathrm{R}^{[2]} = \mathrm{R}^k$, $\mathrm{R}^{[1]} = \mathrm{R}^j$, and let $\pi_{[1]}$ be the projection from $\mathrm{R}^{[2]} = \mathrm{R}^k =$ to $\mathrm{R}^{[1]} = \mathrm{R}^j$ forgetting the last $k-j$ coordinates. We define the tree of realizable sign conditions of $\mathcal{P}$ with respect to Π. For $z \in \mathrm{R}^k$, let $\mathrm{sign}(\mathcal{P})(z)$, be the sign condition on $\mathcal{P}$ mapping $P \in \mathcal{P}$ to $\mathrm{sign}(P)(z) \in \{0,1,-1\}$. Let

$$\mathrm{Sign}_{\Pi,1}(\mathcal{P})(y) = \{\mathrm{sign}(\mathcal{P})(z) | z \in \mathrm{R}^k, \pi_{[1]}(z) = y\},$$

and

$$\mathrm{Sign}_\Pi(\mathcal{P}) = \{\mathrm{Sign}_{\Pi,1}(\mathcal{P})(y) | y \in \mathrm{R}^{[1]}\}.$$

Note that $\mathrm{Sign}_\Pi(\mathcal{P})$ is naturally equipped with a tree structure. We call $\mathrm{Sign}_\Pi(\mathcal{P})$ the tree of realizable sign conditions of $\mathcal{P}$ with respect to Π. A Π-set $\mathcal{A} = \mathcal{A}_1, \mathcal{A}_2$ is a list of two finite sets such that $\mathcal{A}_1 \subset \mathrm{R}^{[1]}$, $\mathcal{A}_2 \subset \mathrm{R}^{[2]}$, and $\pi_{[1]}(\mathcal{A}_2) = \mathcal{A}_2$ Given a Π-set $\mathcal{A} = \mathcal{A}_1, \mathcal{A}_2$, let for $y \in \mathcal{A}_1$

$$\mathrm{Sign}_{\Pi,1}(\mathcal{P}, \mathcal{A})(y) = \{\mathrm{sign}(\mathcal{P})(z) | z \in \mathcal{A}_2, \pi_{[1]}(z) = y\},$$

and

$$\mathrm{Sign}_\Pi(\mathcal{P}, \mathcal{A}) = \{\mathrm{Sign}_{\Pi,1}(\mathcal{P})(z) | z \in \mathcal{A}_1\}.$$

A Π-set $\mathcal{A} = \mathcal{A}_1, \mathcal{A}_2$ is a set of Π-sample points for $\mathcal{P}$ if

$$\mathrm{Sign}_\Pi(\mathcal{P}, \mathcal{A}) = \mathrm{Sign}_\Pi(\mathcal{P}).$$

Notation 3. Let $\mathrm{B}_{\Pi,1}(\mathcal{P}) = \mathrm{BElim}_{X_{[2]}}(\mathcal{P})$ and

$$\mathrm{U}_{\Pi,1}(\mathcal{P}) = \mathrm{U}_{X_{[2]}}(\mathcal{P}),$$
$$\mathrm{U}_{\Pi,0}(\mathcal{P}) = \mathrm{U}_{X_{[1]}}(\mathrm{B}_{\Pi,1}(\mathcal{P})).$$

The elements of $\mathrm{U}_{\Pi,1}(\mathcal{P})$ are parametrized univariate representations in the variable T_2 contained in $\mathrm{D}[X_{[1]}, \varepsilon_2, \delta_2][T_2]^{k-j+2}$. The elements of $\mathrm{U}_{\Pi,0}(\mathcal{P})$ are univariate representations in the variable T_1 contained in $\mathrm{D}[\varepsilon_1, \delta_1][T_1]^{j+2}$. Let

$$u = (u_0, u_1) \in \mathcal{U} = \mathrm{U}_{\Pi,0}(\mathcal{P}) \times \mathrm{U}_{\Pi,1}(\mathcal{P}),$$

with

$$u_1 = (f^{[2]}, g_0^{[2]}, g_1^{[2]}, \ldots, g_{k-j}^{[2]}),$$
$$u_0 = (f^{[1]}, g_0^{[1]}, g_1^{[1]}, \ldots, g_j^{[1]}).$$

For a polynomial $P(X_{[1]}, X_{[2]})$, let $P_u(T_1, T_2)$ denote the polynomial obtained by replacing the block of variables $X_{[1]}$, with the rational fractions associated with the tuple u_1 then the block of variables $X_{[2]}$, with the rational fractions associated with the tuple u_2 (as in Notation 2). Define

$$\mathcal{T}_u = (f^{[1]}(T_1), {f^{[2]}}_u(T_1, T_2)).$$

Also let

$$\overline{u}_1 = ({f^{[2]}}_u, {g_0^{[2]}}_u, {g_1^{[2]}}_u, \ldots, {g_{k-j}^{[2]}}_u) \in \mathrm{D}[T_1, \varepsilon_1, \delta_1, \varepsilon_2, \delta_2][T_2]^{k-j+2}.$$

For $u \in \mathcal{U}$ and $t_\sigma \in \mathrm{Z}(\mathcal{T}_u, \mathrm{R}\langle\varepsilon_1, \delta_1, \varepsilon_2, \delta_2\rangle)$, with Thom encoding σ, let $x_{u,\sigma,0} \in \mathrm{R}\langle\varepsilon_1, \delta_1, \varepsilon_2, \delta_2\rangle^{[j]}$ be the point obtained by substituting the first coordinate of t_σ in the rational functions associated to u_1, and similarly let $x_{u,\sigma,1} \in \mathrm{R}\langle\varepsilon_1, \delta_1, \varepsilon_2, \delta_2\rangle^{[k]}$ be the point obtained by substituting t_σ in the rational functions associated to u_1 and $\overline{u}_2$. Let $\mathcal{A}_1$ be the set of points $x_{u,\sigma,0}$ obtained by considering all $u \in \mathcal{U}$ and $t_\sigma \in \mathrm{Z}(\mathcal{T}_u, \mathrm{R}\langle\varepsilon_1, \delta_1, \varepsilon_2, \delta_2\rangle)$. Similarly let $\mathcal{A}_2$ be the set of points $x_{u,\sigma,1}$ obtained by considering all $u \in \mathcal{U}$ and $t_\sigma \in \mathrm{Z}(\mathcal{T}_u, \mathrm{R}\langle\varepsilon_1, \delta_1, \varepsilon_2, \delta_2\rangle^2)$. Then $\mathcal{A} = \mathcal{A}_1, \mathcal{A}_2$ is a Π-set, specified by $\mathcal{V}$ where the elements of $\mathcal{V}$ are pairs of an element $u \in \mathcal{U}$ and a Thom encoding σ of an element of $\mathrm{Z}(\mathcal{T}_u, \mathrm{R}\langle\varepsilon_1, \delta_1, \varepsilon_2, \delta_2\rangle)$.

The following algorithm is a special case of Algorithm 14.13 in [5], when the number of blocks of variables is equal to 2.

Algorithm 3 (Block Structured Signs)

Input: *a set of s polynomials $\mathcal{P} \subset \mathrm{R}[X_1, \dots, X_k]$, and a partition, Π, of the variables $X_1, \dots, X_k$ into two blocks, $X_1, \dots, X_j$ and $X_{j+1}, \dots, X_k$.*

Output: *the tree $\mathrm{Sign}_\Pi(\mathcal{P})$ of realizable sign conditions of $\mathcal{P}$ with respect to Π and the specification $\mathcal{V}$ of a Π-set $\mathcal{A}$ of sample points for $\mathcal{P}$.*

Complexity: *The number of arithmetic operations in* D *is*

$$s^{(j+1)(k-j+1)} d^{O(j(k-j))}.$$

If $\mathrm{D} = \mathbb{Z}$, and the bitsizes of the coefficients of the polynomials are bounded by τ, then the bitsizes of the integers appearing in the intermediate computations and the output are bounded by $\tau d^{O(j)O(k-j)}$.

4. Computing dimension

The main idea behind computing the dimension of S is quite simple and is based on the geometric characterization of dimension described in Sec. 2. We first check whether the dimension of S is k making use of Proposition 1 or -1 (i.e., is empty). If it is not the case, try $k-1$ or 0, then $k-2$ or 1, etc. making use of Proposition 3.

Algorithm 4 (Dimension)

Input: *a finite subset $\mathcal{P} \subset \mathrm{D}[X_1, \dots, X_K]$ and a semi-algebraic set S described by a quantifier free $\mathcal{P}$-formula $\Phi(X)$.*

Output: *the dimension of* S.

Procedure: *Initialize* $j := 0$.

($\star$) *Consider the block structure* Π_{k-j} *with two blocks of variables:* $X_{j+1}, \dots, X_k$ *and* $X_1, \dots, X_j$.

For every $i = 0, \dots, j(k-j)$ *let* $\mathcal{P}_{k-j,i} = \mathcal{P}(M_{k-j,i})$ *(using Notation* 1*)) and let*

$$S_{k-j,i} = \{x \in \mathrm{R}^k \mid \Phi(M_{k-j,i} \cdot x)\}.$$

Compute $\mathrm{Sign}_{\Pi_{k-j}}(\mathcal{P}_{k-j,i})$ *using Algorithm* 3 *(Block Structured Signs).*

Defining $X_{\leqslant j} = X_1 \dots, X_j$, *compute*

$$\mathrm{SSign}(\mathrm{BElim}_{X_{\leqslant j}}(\mathcal{P}_{k-j,i}))$$

using Algorithm 1 *(Sampling). Note that every sample point output by Algorithm* 3 *(Block Structured Signs) is above a sample point for* $\mathrm{BElim}_{X_{\leqslant j}}(\mathcal{P}_{k-j,i})$ *output by Algorithm* 1 *(Sampling).*

Check whether one of the strict sign conditions in

$$\mathrm{SSign}(\mathrm{BElim}_{X_{\leqslant j}}(\mathcal{P}_{k-j,i}))$$

is satisfied at some point of $\pi_{k-j}(S_{k-j,i})$.

If one of the strict sign conditions in

$$\mathrm{SSign}(\mathrm{BElim}_{X_{\leqslant j}}(\mathcal{P}_{k-j,i}))$$

is satisfied at some point of $\pi_{k-j}(S_{k-j,i})$, *output* $k-j$.

Consider the block structure Π_j *with two blocks of variables:* $X_{k-j+1}, \dots, X_k$ *and* $X_1, \dots, X_{k-j}$.

For every $i = 0, \dots, j(k-j)$ *let* $\mathcal{P}_{j,i} = \mathcal{P}(M_{j,i})$ *(again using Notation* 1*)) and let*

$$S_{j,i} = \{x \in \mathrm{R}^k \mid \Phi(M_{j,i} \cdot x)\}.$$

Compute $\mathrm{Sign}_{\Pi_j}(\mathcal{P}_{j,i})$ *using Algorithm* 3 *(Block Structured Signs).*

Defining $X_{\leqslant kj} = X_1 \dots, X_{k-j}$, *compute*

$$\mathrm{SSign}(\mathrm{BElim}_{X_{\leqslant k-j}}(\mathcal{P}_{j,i}))$$

using Algorithm 1 *(Sampling). Note that every sample point output by Algorithm* 3 *(Block Structured Signs) is above a sample point for* $\mathrm{BElim}_{X_{\leqslant k-j}}(\mathcal{P}_{j,i})$ *output by Algorithm* 1 *(Sampling).*

Check whether one of the strict sign conditions in

$$\mathrm{SSign}(\mathrm{BElim}_{X_{\leqslant k-j}}(\mathcal{P}_{j,i}))$$

is satisfied at some point of $\pi_j(S_{j,i})$*.*

If for every $i = 0 \ldots j(k-j)$ *none of the strict sign conditions in*

$$\mathrm{SSign}(\mathrm{BElim}_{X_{\leqslant k-j}}(\mathcal{P}_{j,i}))$$

is satisfied at some point of $\pi_j(S_{j,i})$*, output* $j-1$*.*

Otherwise define $j := j+1$ *and go to* $(\star)$

Proof of correctness. Follows clearly from Proposition 1, Proposition 3, the correctness of of Algorithm 2 (Block Elimination), Algorithm 1 (Sampling) and Algorithm 3 (Block Structured Signs). □

Complexity analysis. Let s be a bound on the number of elements of $\mathcal{P}$ and d abound on their degrees. There are at most $(k+1)/2$ values of j considered in the algorithm. For a given j, the complexity of the call to Algorithm 3 (Block Structured Signs) performed is $s^{(j+1)(k-j+1)}d^{O(j(k-j))}$, using the complexity analysis of Algorithm 3 (Block Structured Signs) The call to Algorithm 3 (Block Structured Signs) for $\mathrm{BElim}_{X_{>j}}(\mathcal{Q})$, $\mathcal{Q} \in \mathcal{L}_j$ has complexity $s^{(j+1)(k-j+1)}d^{O(j(k-j))}$, using the complexity analysis of Algorithm 3 (Block elimination) and Algorithm 3 (Block Structured Signs), since the number of polynomials is $s^{k-j+1}d^{O(k-j)}$, their degrees are $d^{O(k-j)}$ and their number of variables is j. Finally the total cost of the algorithm is

$$\begin{cases} s^{(k-k')k'}d^{O(k'(k-k'))} & \text{if } k' \geqslant k/2 \\ s^{(k-k'+1)(k'+1)}d^{O(k'(k-k'))} & \text{if } k' < k/2. \end{cases}$$

If $\mathrm{D} = \mathbb{Z}$, and the bitsizes of the coefficients of the polynomials are bounded by τ, then the bitsizes of the integers appearing in the intermediate computations and the output are bounded by $\tau d^{O(k'(k-k'))}$. □

References

1. S. Basu, *On Bounding the Betti Numbers and Computing the Euler Characteristics of Semi-algebraic Sets.* — Discrete and Computational Geometry **22** (1999), 1–18.

2. S. Basu, R. Pollack, M.-F. Roy, *On the Combinatorial and Algebraic Complexity of Quantifier Elimination.* — J. of the ACM **43** (1996), 1002–1045.
3. S. Basu, R. Pollack, M.-F. Roy, *On Computing a Set of Points meeting every Semi-algebraically Connected Component of a Family of Polynomials on a Variety.* — J. of Complexity **13**, No. 1 (1997), 28–37.
4. S. Basu, R. Pollack, M.-F. Roy, *Computing Roadmaps of Semi-algebraic Sets on a Variety.* — J. of the AMS **3**, No. 1 (1999), 55–82.
5. S. Basu, R. M.-F. Roy, *Algorithms in Real Algebraic Geometry.* — Springer-Verlag (2003).
6. J. Canny, *The Complexity of Robot Motion Planning.* — MIT Press (1987).
7. J. Canny, *Some Algebraic and Geometric Computations in PSPACE.* — in: Proceedings of the Twentieth ACM Symposium on Theory of Computing (1988), pp. 460-467.
8. J. Canny, *Computing road maps in general semi-algebraic sets.* — The Computer Journal **36** (1993), 504–514.
9. J. Canny, D. Grigor'ev, N. Vorobjov, *Finding connected components of a semi-algebraic set in subexponential time.* — Appl. Algebra Eng. Commun. Comput. **2**, No. 4 (1992), 217–238.
10. A. Chistov, *Polynomial time computation of the dimension of algebraic varieties in zero characteristic.* — J. of Symbolic Computation **22** (1996), 1–25.
11. A. Chistov, H. Fournier, L. Gurvits, P. Koiran, *Vandermonde Matrices, NP-Completeness and Trasversal Subspaces.* — Foundations of Computational Mathematics **3**(4) (2003), 421–427.
12. G. Collins, *Quantifier elimination for real closed fields by cylindric algebraic decomposition.* — In: Second GI Conference on Automata Theory and Formal Languages. Lect. Notes Comp. Sci. **33**, Springer-Verlag, Berlin (1975), pp. 134–183.
13. L. Gournay, J. J. Risler, *Construction of roadmaps of semi-algebraic sets.* — Appl. Algebra Eng. Commun. Comput. **4**, No. 4 (1993), 239–252.
14. D. Grigor'ev, *The Complexity of deciding Tarski algebra,* — J. of Symbolic Computation **5** (1988), 65–108.
15. D. Grigor'ev, N. Vorobjov, *Solving Systems of Polynomial Inequalities in Subexponential Time.* — J. of Symbolic Computation **5** (1988), 37–64.
16. D. Grigor'ev, N. Vorobjov, *Counting connected components of a semi-algebraic set in subexponential time.* — Comput. Complexity **2**, No. 2 (1992), 133–186.
17. J. Heintz, M.-F. Roy, P. Solernò, *Sur la complexité du principe de Tarski-Seidenberg.* — Bull. Soc. Math. France **118** (1990), 101–126.
18. J. Heintz, M.-F. Roy, P. Solernò, *Single exponential path finding in semi-algebraic sets II : The general case.* — Bajaj, Chandrajit L. (ed.), Algebraic geometry and its applications. Collections of papers from Shreeram S. Abhyankar's 60th birthday conference held at Purdue University, West Lafayette, IN, USA, June 1–4, 1990. New York: Springer-Verlag (1994), pp. 449-465.
19. J. Heintz, M.-F. Roy, P. Solernò, *Description of the Connected Components of a Semialgebraic Set in Single Exponential Time.* — Discrete and Computational Geometry **11** (1994), 121–140.
20. P. Koiran, *The real dimension problem is* $\mathrm{NP}_{\mathbb{R}}$*-complete.* — J. of Complexity **15** (1997), 227–238.

21. J. Renegar, *On the computational complexity and geometry of the first order theory of the reals.* — J. of Symbolic Comput. **13** (1992), 255–352.

22. N. Vorobjov, *Complexity of computing the dimension of a semi-algebraic set.* — J. of Symbolic Comput. **27** (1999), 565–579.

Поступило 17 октября 2004 г.

School of Mathematics,
Georgia Institute of Technology,
Atlanta, GA 30332, U.S.A.

E-mail: saugata@math.gatech.edu

Courant Institute of Mathematical Sciences,
New York University, New York, NY 10012, U.S.A.

E-mail: pollack@cims.nyu.edu

IRMAR (URA CNRS 305),
Université de Rennes,
Campus de Beaulieu 35042 Rennes cedex FRANCE

E-mail: marie-francoise.roy@univ-rennes1.fr