A Probabilistic Symbolic Algorithm to Find the Minimum of a Polynomial Function on a Basic Closed Semialgebraic Set

Gabriela Jeronimo · Daniel Perrucci

Received: 19 April 2013 / Revised: 8 May 2014 / Accepted: 14 July 2014 /

Published online: 14 August 2014

© Springer Science+Business Media New York 2014

Abstract We consider the problem of computing the minimum of a polynomial function g on a basic closed semialgebraic set $E \subset \mathbb{R}^n$. We present a probabilistic symbolic algorithm to find a finite set of sample points of the subset $E^{\min}$ of E where the minimum of g is attained, provided that $E^{\min}$ is non-empty and has at least one compact connected component.

Keywords Polynomial optimization · Complexity · Deformation techniques

1 Introduction

The minimization of polynomial functions over $\mathbb{R}^n$, unrestricted or subject to polynomial constraints, is a classical problem with a variety of applications. In the last years, it has been extensively studied in the algorithmic framework through numerical or symbolic-numerical methods based on certificates of positivity (see, for instance, [8, 10, 18–21, 25]).

From the symbolic computation perspective, a possible way to tackle the problem is to restate it as a quantifier elimination problem over the reals and to apply a symbolic algorithm to solve this more general task; however, better complexity bounds

G. Jeronimo · D. Perrucci

Departamento de Matemática, FCEN, Universidad de Buenos Aires, Buenos Aires, Argentina

e-mail: perrucci@dm.uba.ar

G. Jeronimo

IMAS, CONICET-UBA, Buenos Aires, Argentina

e-mail: jeronimo@dm.uba.ar

D. Perrucci

CONICET, Buenos Aires, Argentina

should be expected by means of algorithms particularly designed for optimization. For instance, a quantifier elimination based deterministic algorithm for optimization with lower complexities is given in [1, Section 14.2]. This approach enables one to deal with optimization problems on arbitrary semi-algebraic sets defined over any real closed field. In [23], a probabilistic algorithm for unconstrained global optimization of polynomial functions over $\mathbb{R}^n$ with better complexity estimates is presented. The better complexity is due to an alternative strategy relying on the computation of generalized critical values. Also, the problem of deciding algorithmically whether the global infimum of a polynomial function is attained is considered in [9].

Recently, in [16] (see also [6, 13]), another approach based on deformation techniques and resultants led to a lower bound for the minimum of a polynomial function on a basic closed semialgebraic set in $\mathbb{R}^n$, provided that the set where this minimum is attained is compact. Here, our aim is to obtain an algorithmic counterpart of this theoretical result. In order to obtain a more efficient algorithm, we replace the use of resultants by polynomial system solving techniques based on the Newton–Hensel lifting. To be able to do so, we need to apply a different deformation.

The precise formulation of the optimization problem we consider is the following. Let $\mathbf{K} \subset \mathbb{R}$ be an effective field. Let $f_1, \dots, f_m, g \in \mathbf{K}[x_1, \dots, x_n]$, with $n \geq 2$, and

$$E = \{x \in \mathbb{R}^n \mid f_1(x) = \dots = f_l(x) = 0, f_{l+1}(x) \geq 0, \dots, f_m(x) \geq 0\},$$

and suppose that g attains a minimum value $g_{\min}$ at E . We look for a symbolic algorithm to compute at least one point in

$$E^{\min} = \{x \in E \mid g(x) = g_{\min}\}.$$

In this paper we assume that $E^{\min}$ has at least one compact connected component, but we do not make any assumptions on the number of constraints or on the genericity of the polynomials giving the constraints. The required compactness assumption holds, for instance, in many families of instances of known optimization problems (see [16, Section 4]).

Our approach consists in finding a finite set of points containing at least one point in each compact connected component of $E^{\min}$. The natural tool to use when solving this problem is the Lagrange Multiplier Theorem; nevertheless, a direct application of this result may lead to a degenerate system or a system with infinitely many solutions. A rough application of existing methods to compute one point in each connected component of the solution set of the Lagrange system would lead to high complexities since the number of variables is increased due to the new variables for the multipliers. In order to overcome these difficulties, we apply deformation techniques as in [15] (see also [1, Chapter 13]), which enable us to deal with “nice” systems that, in the limit, define finite sets containing the required minimizing points. These sets are described by geometric resolutions, which are parametric representations where the parameter ranges over the set of roots of a univariate polynomial. Finally, we compare the values that the given polynomial function g takes at the computed points and obtain the Thom encodings characterizing the minimizers.

The main result of the paper is the following:

Theorem 1 Let $E = \{x \in \mathbb{R}^n \mid f_1(x) = \dots = f_l(x) = 0, f_{l+1}(x) \geq 0, \dots, f_m(x) \geq 0\}$ be defined by polynomials $f_1, \dots, f_m \in \mathbf{K}[x_1, \dots, x_n]$ with $n \geq 2$ and degrees bounded by an even integer d . Let $g \in \mathbf{K}[x_1, \dots, x_n]$ be a polynomial of degree at most d that attains a minimum value $g_{\min}$ at E in a non-empty set $E^{\min}$ with at least one compact connected component. Algorithm *FindingMinimum* (see Sect. 5) is a probabilistic procedure that, taking as input the integer d and the polynomials $f_1, \dots, f_m, g$ encoded by a straight-line program of length L , computes a family

$$\{(p_i, v_{i,1}, \dots, v_{i,n}), \tau_i\}_{i \in \mathcal{I}},$$

where $\mathcal{I}$ is a finite set and for every $i \in \mathcal{I}$, $(p_i, v_{i,1}, \dots, v_{i,n})$ is a geometric resolution in $\mathbf{K}[u]$ and $\tau_i \in \{-1, 0, 1\}^{\deg p_i}$ is the Thom encoding of a real root ξ_i of p_i such that the set

$$\{(v_{i,1}(\xi_i), \dots, v_{i,n}(\xi_i))\}_{i \in \mathcal{I}}$$

is included in $E^{\min}$ and intersects all its compact connected components. The complexity of the algorithm is

$$O((n^3(L + dn + n^{\Omega-1})D^2 \log^2(D) \log \log^2(D) + (m + D)D^2 \log^3(D))\Upsilon),$$

where

- $D = \max_{0 \leq s \leq \min\{n, m\}} \binom{n}{s} d^s (d-1)^{n-s}$,
- $\Upsilon = \sum_{0 \leq s \leq \min\{n, m\}} \sum_{s_1+s_2=s} \sum_{0 \leq s_1 \leq l, 0 \leq s_2 \leq m-l} \binom{l}{s_1} \binom{m-l}{s_2} 2^{s_1} \leq \sum_{0 \leq s \leq \min\{n, m\}} \binom{m}{s} 2^s$.

In the above statement, Ω denotes a positive real number such that for any ring R , addition, multiplication and the computation of determinant and adjoint of matrices in $R^{k \times k}$ can be performed within $O(k^\Omega)$ operations in R . We may assume $\Omega \leq 4$ (see [3]) and, in order to simplify complexity estimations, we also assume that $\Omega \geq 3$.

For the optimization problem over basic closed semialgebraic sets we consider, the complexity of our probabilistic algorithm improves the one of the best known quantifier elimination based deterministic procedure (see [1, Section 14.2]). In addition, for the particular case of unconstrained global optimization, our complexity also improves asymptotically the one of the probabilistic algorithm from [23], but this algorithm works even if the infimum or supremum is not attained.

The paper is organized as follows. In Sect. 2, we introduce the basic notation and state some previous results we use throughout the paper. In Sect. 3, we present the deformation we apply and we prove some of its geometric properties. Section 4 is devoted to showing how the deformation leads to a geometric resolution of the finite set we look for. Finally, in Sect. 5 we show the algorithmic counterparts of the previous theoretical results, proving Theorem 1.

2 Notation and Preliminaries

Throughout the paper, we denote $\mathbb{N}$ the set of positive integers. For a field K , we write $\bar{K}$ for an algebraic closure of K , $K(t)$ for the field of rational functions in a single variable t and $K[[t]]$ for the set of formal power series in t .

For a given $n \in \mathbb{N}$, the n -dimensional affine and projective spaces over an algebraically closed field K are denoted by $\mathbb{A}_K^n$ and $\mathbb{P}_K^n$ respectively. When the base field is $K = \mathbb{C}$, we write simply $\mathbb{A}^n$ and $\mathbb{P}^n$.

2.1 Algorithms and Complexity

The algorithms we consider in this paper are described over an effective field $\mathbf{K} \subset \mathbb{R}$. The notion of *complexity* of an algorithm we consider is the number of operations and comparisons over $\mathbf{K}$ that the execution of the algorithm requires. In this definition of complexity, accessing, reading and writing pre-computed objects is cost free.

Our algorithms are probabilistic in the sense that they make random choices of the values of certain parameters during their execution. However, on every input, a generic choice of these values ensures that the algorithm produces the correct output.

The objects we deal with are polynomials with coefficients in $\mathbf{K}$. In our algorithms we represent each polynomial either as the array of all its coefficients in a pre-fixed order of its monomials (*dense form*) or by a *straight-line program*. Roughly speaking, a straight-line program (or slp, for short) over $\mathbf{K}$ encoding a list of polynomials in $\mathbf{K}[x_1, \dots, x_n]$ is a program without branches (an arithmetic circuit) which enables us to evaluate these polynomials at any given point in $\mathbf{K}^n$. The number of instructions in the program is called the *length* of the slp (for a precise definition we refer to [4, Definition 4.2]; see also [12]). From the dense encoding of a family of m polynomials in $\mathbf{K}[x_1, \dots, x_n]$ of degrees bounded by d , we can obtain an slp of length $O(m \binom{d+n}{n})$ encoding them. Unless otherwise stated, throughout the paper, univariate polynomials will be encoded in dense form.

To estimate complexities we will use the following results. Operations between univariate polynomials with coefficients in a field $\mathbf{K}$ of degree bounded by d in dense form can be done using $O(d \log(d) \log \log(d))$ operations in $\mathbf{K}$ (see [26, Chapters 8 and 9]) and gcd or resultant computations by means of the Extended Euclidean Algorithm (see [26, Chapter 11]) can be performed within $O(d \log^2(d) \log \log(d))$ operations in $\mathbf{K}$. Given an slp of length L encoding a family of m univariate polynomials of degree at most d , we can obtain their dense form within $O(dL + md \log^2(d) \log \log(d))$ operations in $\mathbf{K}$ (see [26, Corollary 10.12]).

From an slp of length L encoding a polynomial $f \in \mathbf{K}[x_1, \dots, x_n]$, we can compute an slp of length $O(L)$ encoding f and all its first order partial derivatives (see [2]).

2.2 Geometric Resolutions

A way of representing zero-dimensional affine varieties which is widely used in computer algebra nowadays is a *geometric resolution* (see, for instance, [7]). The precise definition we are going to use is the following:

Let K be a field of characteristic 0 and $V = \{z_1, \dots, z_D\} \subset \mathbb{A}_{\bar{K}}^n$ be a zero-dimensional variety defined by polynomials in $K[x_1, \dots, x_n]$. Given a *separating* linear form $\ell = \alpha_1 x_1 + \dots + \alpha_n x_n \in K[x_1, \dots, x_n]$ for V (that is, a linear form ℓ such that $\ell(z_i) \neq \ell(z_j)$ if $i \neq j$), the following polynomials completely characterize the variety V :

- the *minimal polynomial* $p := \prod_{1 \leq i \leq D} (u - \ell(z_i)) \in K[u]$ of ℓ over the variety V (where u is a new variable),
- polynomials $v_1, \dots, v_n \in K[u]$ with $\deg(v_j) < D$ for every $1 \leq j \leq n$ satisfying $z_i = (v_1(\ell(z_i)), \dots, v_n(\ell(z_i)))$ for $1 \leq i \leq D$,

since they satisfy that

$$V = \{(v_1(\xi), \dots, v_n(\xi)) \in \bar{K}^n \mid \xi \in \bar{K}, p(\xi) = 0\}.$$

The family of univariate polynomials $(p, v_1, \dots, v_n)$ is called a *geometric resolution* of V (associated with the linear form ℓ). Note that if K is a subfield of $\mathbb{R}$, the real roots of p correspond to the real points of the variety V .

Given geometric resolutions $(p_1, v_{11}, \dots, v_{1n})$ and $(p_2, v_{21}, \dots, v_{2n})$ of two zero-dimensional varieties V_1 and V_2 in $\mathbb{A}^n$ consisting of D_1 and D_2 points respectively, associated with the same linear form ℓ which separates the points in $V_1 \cup V_2$, we can obtain a geometric resolution of $V_1 \cup V_2$ within complexity $O(nD \log^2(D) \log \log(D))$, where $D = \max\{D_1, D_2\}$, by means of the Chinese Remainder Theorem using the Extended Euclidean Algorithm.

2.3 Thom Encoding of Real Algebraic Numbers

The *Thom encoding* of real algebraic numbers provides an algebraic approach to distinguish the different real roots of a real univariate polynomial. We recall here its definition and main properties (see [1, Chapter 2]).

Given $p \in \mathbf{K}[u]$ and a real root ξ of p , the Thom encoding of ξ as a root of p is the sequence $(\text{sign}(p'(\xi)), \dots, \text{sign}(p^{(\deg p)}(\xi)))$, where we represent the sign with an element of the set $\{-1, 0, 1\}$. If the sign of the leading coefficient of p is known, the Thom encoding can be shortened to $(\text{sign}(p'(\xi)), \dots, \text{sign}(p^{(\deg p - 1)}(\xi)))$.

Two different real roots of p have different Thom encodings. In addition, given the Thom encodings $(\sigma_{1,1}, \dots, \sigma_{1,\deg p})$ and $(\sigma_{2,1}, \dots, \sigma_{2,\deg p})$ of two different real roots ξ_1 and ξ_2 of p , it is possible to decide which is the smallest between ξ_1 and ξ_2 as follows: Consider the largest value of k such that $\sigma_{1,k} \neq \sigma_{2,k}$; then $k < \deg p$, since $p^{(\deg p)}$ is a constant. Also $\sigma_{1,k+1} = \sigma_{2,k+1} \neq 0$, since otherwise ξ_1 and ξ_2 would have the same Thom encoding with respect to the polynomial $p^{(k+1)}$ and therefore $\xi_1 = \xi_2$. Then

- if $\sigma_{1,k+1} = \sigma_{2,k+1} = 1$, we have that $\xi_1 < \xi_2$ if and only if $\sigma_{1,k} < \sigma_{2,k}$,
- if $\sigma_{1,k+1} = \sigma_{2,k+1} = -1$, we have that $\xi_1 < \xi_2$ if and only if $\sigma_{1,k} > \sigma_{2,k}$.

3 The Deformation

3.1 Defining the Deformation

Here we introduce the deformation we use. We denote:

- $q_0 := n + 1$ and $q_1 < \dots < q_m$ the first m prime numbers greater than $n + 1$. Let $A \in \mathbb{Q}^{(m+1) \times (n+1)}$, $A = (a_{ij})_{0 \leq i \leq m, 0 \leq j \leq n}$ be the Cauchy matrix defined by $a_{ij} = \frac{1}{q_i - j}$ (note that each submatrix of A has maximal rank and $a_{ij} > 0$ for every i, j).
- For $e \in \mathbb{N}$, T_e the Tchebychev polynomial of degree e (see [17, Section 6.1]).
- $\tilde{g}(x) = \sum_{1 \leq j \leq n} a_{0j} T_d(x_j)$ and, for every $1 \leq i \leq m$, $\tilde{f}_i(x) = a_{i0} + \sum_{1 \leq j \leq n} a_{ij} (T_d(x_j) + 1)$.
- $G(t, x) = t g(x) + (1 - t) \tilde{g}(x)$ and, for every $1 \leq i \leq m$, $F_i^+(t, x) = t f_i(x) + (1 - t) \tilde{f}_i(x)$ and $F_i^-(t, x) = t f_i(x) - (1 - t) \tilde{f}_i(x)$.
- For every $S \subset \{1, \dots, m\}$ and $\sigma \in \{+, -\}^S$,

$$\hat{V}_{S,\sigma} = \left\{ (t, x, \lambda) \in \mathbb{A} \times \mathbb{A}^n \times \mathbb{P}^{|S|} \mid F_i^{\sigma_i}(t, x) = 0 \text{ for every } i \in S, \right. \\ \left. \lambda_0 \nabla_x G(t, x) = \sum_{i \in S} \lambda_i \nabla_x F_i^{\sigma_i}(t, x) \right\}.$$

We consider the decomposition of $\hat{V}_{S,\sigma}$ as $\hat{V}_{S,\sigma} = V_{S,\sigma}^{(t)} \cup V_{S,\sigma}$, where

- $V_{S,\sigma}^{(t)}$ is the union of the irreducible components of $\hat{V}_{S,\sigma}$ included in $t = t_0$ for some $t_0 \in \mathbb{C}$,
- $V_{S,\sigma}$ is the union of the remaining irreducible components of $\hat{V}_{S,\sigma}$.
- For a group of variables y , Π_y the projection to the coordinates y .

In [16], a similar deformation is defined using powers of the variables instead of Tchebychev polynomials; however, it is not suitable for the algorithmic purposes of the present work. In the algorithmic framework, deformations based on Tchebychev polynomials have already been applied in [15]. The main properties of this deformation on which our algorithms rely are stated in Lemma 6.

3.2 Geometric Properties

Let C be a compact connected component of $E^{\min}$. For $\delta > 0$, we write

- $C_{=\delta} = \{x \in \mathbb{R}^n \mid \text{dist}(x, C) = \delta\}$,
- $C_{\leq \delta} = \{x \in \mathbb{R}^n \mid \text{dist}(x, C) \leq \delta\}$,
- $C_{< \delta} = \{x \in \mathbb{R}^n \mid \text{dist}(x, C) < \delta\}$.

Let $\mu > 0$ such that $C_{\leq \mu}$ and $E_{\min} \setminus C$ do not intersect.

We consider

$$\begin{aligned}\tilde{E} = \{ & (t, x) \in \mathbb{R} \times \mathbb{R}^n \mid F_1^+(t, x) \geq 0, \dots, \\ & F_l^+(t, x) \geq 0, F_{l+1}^+(t, x) \geq 0, \dots, F_m^+(t, x) \geq 0, \\ & F_1^-(t, x) \leq 0, \dots, F_l^-(t, x) \leq 0\}\end{aligned}$$

and, for every $0 \leq t \leq 1$,

$$E_t = \Pi_x(\tilde{E} \cap (\{t\} \times \mathbb{R}^n)) \subset \mathbb{R}^n.$$

Note that $E_0 = \mathbb{R}^n$, $E_1 = E$ and for $0 \leq t_1 \leq t_2 \leq 1$, $E_{t_2} \subset E_{t_1}$.

Lemma 2 *There exists $0 < \varepsilon < 1$ such that for every $1 - \varepsilon \leq t \leq 1$, the minimum value that $G(t, \cdot)$ takes on $E_t \cap C_{\leq \mu}$ is not attained at any point in $E_t \cap C_{=\mu}$.*

Proof The minimum value of $g(\cdot) = G(1, \cdot)$ at $E_1 \cap C_{\leq \mu}$ is $g_{\min}$, which over this set is only attained at points in C and, therefore, not in $E_1 \cap C_{=\mu}$. Assume the claim does not hold. Then there exists a strictly increasing sequence of positive numbers $(t_k)_{k \in \mathbb{N}}$ converging to 1 such that the minimum value of $G(t_k, \cdot)$ at $E_{t_k} \cap C_{\leq \mu}$ is attained at a point $z_k \in E_{t_k} \cap C_{=\mu}$.

The sequence $(z_k)_{k \in \mathbb{N}}$ is contained in the compact set $C_{=\mu}$; therefore, without loss of generality, we may assume that it converges to a point $z \in C_{=\mu}$. On the other hand, the sequence $(t_k, z_k)_{k \in \mathbb{N}}$ is contained in $\tilde{E}$ and converges to $(1, z)$, then $(1, z) \in \tilde{E}$ and $z \in E_1$.

We will prove that $g(z) = g_{\min}$. Take $z' \in E_1 \cap C_{\leq \mu}$, then $z' \in E_{t_k} \cap C_{\leq \mu}$ for every $k \in \mathbb{N}$ and

$$g(z') = G(1, z') = \lim_{k \rightarrow \infty} G(t_k, z') \geq \lim_{k \rightarrow \infty} G(t_k, z_k) = G(1, z) = g(z).$$

Then g attains its minimum on $E_1 \cap C_{\leq \mu}$ at z , which is $g_{\min}$. This leads to a contradiction since this value is not attained at any point in $E_1 \cap C_{=\mu}$. $\square$

The following proposition shows that in order to obtain minimizers for the polynomial function g on the compact connected component C , it is enough to consider at most as many of the equations and inequations defining E as the number of variables. We define the set

$$\begin{aligned}S = \{ & (S, \sigma) \mid S \subset \{1, \dots, m\} \text{ with } 0 \leq |S| \leq n \text{ and} \\ & \sigma \in \{+, -\}^S \text{ with } \sigma_i = + \text{ for } l+1 \leq i \leq m\}.\end{aligned}$$

Proposition 3 *Let C be a compact connected component of $E^{\min}$. There exist $z \in C$ and $(S, \sigma) \in S$ such that $z \in \Pi_x(V_{S, \sigma} \cap \{t = 1\})$.*

We point out that if E (resp. $E^{\min}$) is compact, we can easily prove Proposition 3 adapting the arguments in the proof of [16, Proposition 7] (resp. [16, Theorem 14]). Now we prove it under our weaker assumptions.

Proof of Proposition 3 Consider $0 < \varepsilon < 1$ such that

- for every $1 - \varepsilon \leq t \leq 1$, the minimum value that $G(t, \cdot)$ takes on $E_t \cap C_{\leq \mu}$ is not attained at any point in $E_t \cap C_{=\mu}$,
- for every $S \subset \{1, \dots, m\}$ and $\sigma \in \{+, -\}^S$, $\Pi_t(V_{S,\sigma}^{(t)}) \cap (1 - \varepsilon, 1) = \emptyset$,
- for every $1 - \varepsilon \leq t \leq 1$, $S \subset \{1, \dots, m\}$ with $|S| > n$ and $\sigma \in \{+, -\}^S$, the set

$$\{x \in \mathbb{A}^n \mid F_i^{\sigma_i}(t, x) = 0 \text{ for every } i \in S\}$$

is empty.

The existence of such an ε follows from Lemma 2, the finiteness of $\Pi_t(V_{S,\sigma}^{(t)})$ for every S and σ , and an adaptation of the arguments in [15, Lemma 21] or [16, Lemma 4].

Let $(t_k)_{k \in \mathbb{N}}$ be an increasing sequence converging to 1 with $t_1 > 1 - \varepsilon$ and let $z_k \in E_{t_k} \cap C_{< \mu}$ be a point such that $G(t_k, \cdot)$ attains its minimum value on the set $E_{t_k} \cap C_{\leq \mu}$ at z_k . Without loss of generality, we may assume that the sequence $(z_k)_{k \in \mathbb{N}}$ is convergent to a point $z \in E_1 \cap C_{\leq \mu}$, and proceeding as in the proof of Lemma 2, we have that $z \in C$.

Now, for every $k \in \mathbb{N}$ and every $x \in \mathbb{R}^n$, at most one $F_i^+(t_k, x)$ and $F_i^-(t_k, x)$ may vanish. Let

$$S_k = \{i \in \{1, \dots, l\} \mid F_i^+(t_k, z_k) = 0 \text{ or } F_i^-(t_k, z_k) = 0\} \cup \\ \{i \in \{l+1, \dots, m\} \mid F_i^+(t_k, z_k) = 0\};$$

then $0 \leq |S_k| \leq n$. Without loss of generality, we may assume that S_k is the same set S for every $k \in \mathbb{N}$; moreover, we may assume that, for each $i \in S \cap \{1, \dots, l\}$, it is always the same, $F_i^+(t_k, z_k)$ or $F_i^-(t_k, z_k)$, the one which vanishes, thus defining a function $\sigma \in \{+, -\}^S$ with $\sigma_i = +$ for $l+1 \leq i \leq m$. Then, $(S, \sigma) \in \mathcal{S}$.

For $k \in \mathbb{N}$, if the set $\{\nabla_x F_i^{\sigma_i}(t_k, z_k), i \in S\}$ is linearly independent, since the function $G(t_k, \cdot)$ attains a local minimum at the point z_k when restricted to the set $E_{t_k} \cap C_{< \mu}$, by the Lagrange Multiplier Theorem, there exists $(\lambda_{i,k})_{i \in S}$ such that

$$\nabla_x G(t_k, z_k) = \sum_{i \in S} \lambda_{i,k} \nabla_x F_i^{\sigma_i}(t_k, z_k).$$

We take $\lambda_{0,k} = 1$ and we have that $(t_k, z_k, (\lambda_{0,k}, (\lambda_{i,k})_{i \in S})) \in \hat{V}_{S,\sigma}$; but since $t_k \notin \Pi_t(V_{S,\sigma}^{(t)})$, we have that $(t_k, z_k, (\lambda_{0,k}, (\lambda_{i,k})_{i \in S})) \in V_{S,\sigma}$. On the other hand, if $\sum_{i \in S} \lambda_{i,k} \nabla_x F_i^{\sigma_i}(t_k, z_k) = 0$ with $(\lambda_{i,k})_{i \in S} \neq 0$, we take $\lambda_{0,k} = 0$ and, as in the previous case, $(t_k, z_k, (\lambda_{0,k}, (\lambda_{i,k})_{i \in S})) \in V_{S,\sigma}$.

Without loss of generality, we may assume that $(\lambda_{0,k}, (\lambda_{i,k})_{i \in S})_{k \in \mathbb{N}}$ converges to a point $(\lambda_0, (\lambda_{i,0})_{i \in S}) \in \mathbb{P}^{|S|}$; then $(1, z, (\lambda_0, (\lambda_{i,0})_{i \in S})) \in V_{S,\sigma}$ and, therefore, $z \in \Pi_x(V_{S,\sigma} \cap \{t = 1\})$ as we wanted to prove. $\square$

Proposition 4 *For every $(S, \sigma) \in \mathcal{S}$, we have that $\Pi_x(V_{S,\sigma} \cap \{t = 1\})$ is a finite set.*

The proof of Proposition 4 will follow from arguments in the next section. From Propositions 3 and 4 we deduce the following:

Corollary 5 *Under the assumptions of Theorem 1, the set*

$$\bigcup_{(S,\sigma) \in \mathcal{S}} \Pi_x(V_{S,\sigma} \cap \{t = 1\})$$

is finite and contains a point in every compact connected component of $E^{\min}$.

We point out that in this paper we focus on semialgebraic sets over the field of real numbers; in particular, the proofs of Lemma 2 and Proposition 3 use compactness arguments that do not hold over arbitrary real closed fields.

4 A Geometric Resolution

Due to Corollary 5, to solve the problem we are considering, we will now focus on describing the set $\Pi_x(V_{S,\sigma} \cap \{t = 1\})$ for every $(S, \sigma) \in \mathcal{S}$.

For simplicity, throughout this section, we consider fixed $(S, \sigma) \in \mathcal{S}$ and denote $\hat{V} = \hat{V}_{S,\sigma}$ and $V = V_{S,\sigma}$; moreover, we suppose $S = \{1, \dots, s\}$ and $\sigma = \{+\}^S$. For $1 \leq j \leq n$, let

$$g_j(x, \lambda) = \lambda_0 \frac{\partial g}{\partial x_j} - \sum_{1 \leq i \leq s} \lambda_i \frac{\partial f_i}{\partial x_j} \in \mathbf{K}[x_1, \dots, x_n, \lambda_0, \dots, \lambda_s],$$

$$\tilde{g}_j(x, \lambda) = \lambda_0 \frac{\partial \tilde{g}}{\partial x_j} - \sum_{1 \leq i \leq s} \lambda_i \frac{\partial \tilde{f}_i}{\partial x_j} \in \mathbf{K}[x_1, \dots, x_n, \lambda_0, \dots, \lambda_s],$$

$$G_j(t, x, \lambda) = t g_j(x, \lambda) + (1 - t) \tilde{g}_j(x, \lambda) \in \mathbf{K}[t, x_1, \dots, x_n, \lambda_0, \dots, \lambda_s].$$

These polynomials are homogeneous of degree 1 in the variables λ ; therefore, by the multihomogeneous Bézout theorem (see, for instance, [24, Chapter 4, Section 2.1]), the degree of the varieties $\hat{V} \cap \{t = t_0\}$ for $t_0 \in \mathbb{C}$ is bounded by

$$D_s := \binom{n}{s} d^s (d - 1)^{n-s}.$$

The next lemma shows the key properties of the initial system in the deformation based on Tchebychev polynomials introduced in Sect. 3.1 (cf. [15, Lemma 20]). In addition to what is stated, its proof shows that the solution set of the initial system can be partitioned into a finite union of solution sets of square systems in separated variables, which will play a key role in our algorithms (see the proof of Proposition 8).

Lemma 6 *The polynomials $\tilde{f}_1, \dots, \tilde{f}_s, \tilde{g}_1, \dots, \tilde{g}_n$ define a 0-dimensional variety in $\mathbb{A}^n \times \{\lambda_0 \neq 0\} \subset \mathbb{A}^n \times \mathbb{P}^s$ with D_s distinct points $w_1, \dots, w_{D_s}$ satisfying $\Pi_x(w_i) \neq \Pi_x(w_j)$ for $i \neq j$, and the Jacobian determinant of $\tilde{f}_1, \dots, \tilde{f}_s$ and the polynomials obtained from $\tilde{g}_1, \dots, \tilde{g}_n$ dehomogenizing with $\lambda_0 = 1$ does not vanish at any of these points.*

Proof Recalling that $\tilde{g}(x) = \sum_{1 \leq j \leq n} a_{0j} T_d(x_j)$ and $\tilde{f}_i = a_{i0} + \sum_{1 \leq j \leq n} a_{ij} (T_d(x_j) + 1)$, we have that, for every $1 \leq j \leq n$,

$$\tilde{g}_j(x, \lambda) = T'_d(x_j) \left(a_{0j} \lambda_0 - \sum_{1 \leq i \leq s} a_{ij} \lambda_i \right).$$

Therefore, the solution set of the system $\tilde{f}_1, \dots, \tilde{f}_s, \tilde{g}_1, \dots, \tilde{g}_n$ can be decomposed as

$$\bigcup_{B \subset \{1, \dots, n\}} \left\{ T'_d(x_j) = 0 \text{ for all } j \in B, \tilde{f}_1(x) = 0, \dots, \tilde{f}_s(x) = 0 \right\} \\ \times \left\{ a_{0j} \lambda_0 - \sum_{1 \leq i \leq s} a_{ij} \lambda_i = 0 \text{ for all } j \notin B \right\}.$$

By our assumption on the matrix A , if $|B| = n - s$, the linear system $a_{0j} \lambda_0 - \sum_{1 \leq i \leq s} a_{ij} \lambda_i = 0$ for all $j \notin B$ has a unique solution $\Lambda_B \in \mathbb{P}^s$; moreover, this solution lies in $\{\lambda_0 \neq 0\}$.

For a fixed $B \subset \{1, \dots, n\}$ with $|B| = n - s$, taking into account that T'_d has $d - 1$ real roots and T_d takes the value 1 or -1 at each of these roots, we have that

$$S_B := \{T'_d(x_j) = 0 \text{ for all } j \in B, \tilde{f}_1(x) = 0, \dots, \tilde{f}_s(x) = 0\}$$

decomposes as the union of the sets

$$S_{B,e} := \{T'_d(x_j) = 0, T_d(x_j) = e(j) \text{ for all } j \in B, \tilde{f}_1^{B,e} = 0, \dots, \tilde{f}_s^{B,e} = 0\}$$

for all $e : B \rightarrow \{1, -1\}$, where $\tilde{f}_i^{B,e} \in \mathbf{K}[x_j; j \notin B]$ denotes the polynomial obtained from $\tilde{f}_i$ by replacing $T_d(x_j) = e(j)$ for every $j \in B$.

Without loss of generality, in order to simplify notation, assume $B = \{s+1, \dots, n\}$. Then, for $e : B \rightarrow \{1, -1\}$, the system $\tilde{f}_1^{B,e} = 0, \dots, \tilde{f}_s^{B,e} = 0$ can be written in the form

$$A_B \begin{pmatrix} T_d(x_1) + 1 \\ \vdots \\ T_d(x_s) + 1 \end{pmatrix} = \begin{pmatrix} \alpha_1^{B,e} \\ \vdots \\ \alpha_s^{B,e} \end{pmatrix}$$

where $A_B := (a_{ij})_{1 \leq i, j \leq s}$ and $\alpha_i^{B,e} = -a_{i0} - \sum_{s+1 \leq j \leq n} a_{ij} (e(j) + 1)$ for $1 \leq i \leq s$. Since A_B is invertible, we can solve the underlying linear system for $T_d(x_1) + 1, \dots, T_d(x_s) + 1$. By applying Cramer's rule, it can be seen that the coordinates of the solution to this linear system are rational numbers where the denominators are a multiple of the prime number q_s , whereas the numerators are relatively prime with q_s ; therefore, no coordinate of a solution is an integer number. We deduce that the above system is equivalent to a system of the form

$$T_d(x_1) = c_1^{B,e}, \dots, T_d(x_s) = c_s^{B,e}$$

where $c_i^{B,e} \neq \pm 1$ for every $1 \leq i \leq s$. It follows that each of the equations has d distinct roots, none of which equals a root of T_d' (thus, the sets $S_{B,e}$ are mutually disjoint).

Moreover, the Jacobian matrix of $\tilde{f}_1, \dots, \tilde{f}_s$ and the polynomials obtained from $\tilde{g}_1, \dots, \tilde{g}_n$ dehomogenizing with $\lambda_0 = 1$ evaluated at any of its solutions is of the form

$$\begin{array}{l} s \\ s \\ n-s \end{array} \left\{ \underbrace{\left(\begin{array}{c|c|c} C_1 & 0 & 0 \\ \hline * & 0 & C_2 \\ \hline * & C_3 & * \end{array} \right)}_{\substack{s \quad n-s \quad s}} \right\}.$$

It is easy to see that C_1 , C_2 and C_3 are invertible matrices and so, the Jacobian determinant does not vanish.

We conclude that S_B consists of $(d-1)^{n-s}d^s$ distinct points in $\mathbb{A}^n$ for every B with $|B| = n-s$. Hence, the system $\tilde{f}_1, \dots, \tilde{f}_s, \tilde{g}_1, \dots, \tilde{g}_n$ has D_s isolated solutions in $\mathbb{A}^n \times \mathbb{P}^s$ whose projections to $\mathbb{A}^n$ are all distinct. Since D_s is an upper bound for the degree of the variety the system defines, it follows that these are all its solutions. $\square$

As a consequence of Lemma 6, it follows that all the irreducible components of V intersect the set $\{t=0\}$ and have dimension 1. Proposition 4 is immediate from this fact. Moreover, the following further properties of the induced deformation hold.

Lemma 7 *The variety defined in $\mathbb{A}_{\mathbf{K}(t)}^n \times \mathbb{P}_{\mathbf{K}(t)}^s$ by $F_1, \dots, F_s, G_1, \dots, G_n$ is 0-dimensional and has D_s distinct points $W_1, \dots, W_{D_s}$ in $\{\lambda_0 \neq 0\}$ such that $\Pi_x(W_i) \neq \Pi_x(W_j)$ for $i \neq j$. Moreover, by means of the map $\phi: (\mathbb{A}_{\mathbf{K}(t)}^n \times \mathbb{P}_{\mathbf{K}(t)}^s) \cap \{\lambda_0 \neq 0\} \rightarrow \mathbb{A}_{\mathbf{K}(t)}^n \times \mathbb{A}_{\mathbf{K}(t)}^s$, $\phi(x, (\lambda_0, \lambda_1, \dots, \lambda_s)) = (x, (\frac{\lambda_1}{\lambda_0}, \dots, \frac{\lambda_s}{\lambda_0}))$, these points can be considered as elements in $\overline{\mathbf{K}}[[t]]^{n+s}$.*

Proof If $w_1, \dots, w_{D_s}$ are the common zeros of $\tilde{f}_1, \dots, \tilde{f}_s, \tilde{g}_1, \dots, \tilde{g}_n$, the Jacobian with respect to $x_1, \dots, x_n, \lambda_1, \dots, \lambda_s$ of $F_1, \dots, F_s$, and the polynomials obtained from $G_1, \dots, G_n$ dehomogenizing with $\lambda_0 = 1$ at $t = 0$ and $(x, \lambda_1, \dots, \lambda_s) = \phi(w_i)$ is nonzero, where by abuse of notation ϕ is the map from the statement of the lemma considered over the base field $\overline{\mathbf{K}}$. Applying the Newton–Hensel lifting to the dehomogenized system and the initial points $\phi(w_1), \dots, \phi(w_{D_s})$ (for example as in [11, Lemma 3]) we obtain D_s distinct points in $\overline{\mathbf{K}}[[t]]^{n+s}$ which are solutions of this system. Finally, $W_1, \dots, W_{D_s}$ are obtained by applying ϕ^{-1} to these points. They are all the common solutions to $F_1, \dots, F_s, G_1, \dots, G_n$ since the multihomogeneous Bézout Theorem states that the degree of the variety the system defines is bounded by D_s . $\square$

Consider now new variables $y_1, \dots, y_n$ and define $\ell(x, \lambda, y) = \ell(x, y) = \sum_{1 \leq j \leq n} y_j x_j$. For $\alpha_1, \dots, \alpha_n \in \mathbb{C}$, let $\ell_\alpha(x, \lambda) = \ell_\alpha(x) = \sum_{1 \leq j \leq n} \alpha_j x_j$. Let

$$\begin{aligned}
 P(t, u, y) &= \prod_{1 \leq i \leq D_s} (u - \ell(W_i, y)) \\
 &= \frac{\sum_{0 \leq h \leq D_s} p_h(t, y) u^h}{q(t)} = \frac{\hat{P}(t, u, y)}{q(t)} \in \mathbf{K}(t)[u, y],
 \end{aligned}$$

with $\hat{P}(t, u, y) \in \mathbf{K}[t, u, y]$ with no factors in $\mathbf{K}[t]$. Let $Q(u, y) = \gcd(\hat{P}(1, u, y), \partial \hat{P} / \partial u(1, u, y))$. Then, for generic $\alpha \in \mathbb{C}^n$, if

$$p(u) := \frac{\hat{P}(1, u, \alpha)}{Q(u, \alpha)}$$

and, for every $1 \leq j \leq n$,

$$v_j(u) := -\frac{\frac{\partial \hat{P}}{\partial y_j}(1, u, \alpha)}{Q(u, \alpha)} \left(\frac{\frac{\partial \hat{P}}{\partial u}(1, u, \alpha)}{Q(u, \alpha)} \right)^{-1} \pmod{p(u)},$$

we have that $(p(u), v_1(u), \dots, v_n(u))$ is the geometric resolution associated to ℓ_α of a finite set $\mathcal{P}$ containing $\Pi_x(V \cap \{t = 1\})$ (see, for instance, [7, Algorithm 9] or [15, Proposition 8]).

The computation of $\hat{P}$ will be done in Sect. 5 by means of a Newton–Hensel based approximation (see Algorithm `GeometricResolution`). The required precision is obtained from the degree bound $\deg_t \hat{P}(t, u, y) \leq nD_s$, which can be proved as in [15, Lemma 9].

5 The Algorithm

In this section we describe our algorithms and prove the main result of the paper. For the sake of readability, for each algorithm we first present a schematic description containing its main steps. A precise description of how computations are performed at each step, the complexity analysis and proof of correctness are given right after the presentation of the algorithm.

From Corollary 5, we know that the finite set

$$\bigcup_{(S, \sigma) \in \mathcal{S}} \Pi_x(V_{S, \sigma} \cap \{t = 1\})$$

contains a point in every compact connected component of $E^{\min}$; nevertheless, for a fixed $(S, \sigma) \in \mathcal{S}$, the set $\Pi_x(V_{S, \sigma} \cap \{t = 1\})$ is not necessarily contained in $E^{\min}$, or may even have an empty intersection with E . The idea of our main algorithm is to compute first finite sets $\mathcal{P}_{S, \sigma}$ containing $\Pi_x(V_{S, \sigma} \cap \{t = 1\})$; then, look for the points of each $\mathcal{P}_{S, \sigma}$ that lie in E and finally, compare the values that the function g takes at these points.

First, we introduce three auxiliary subroutines we use to construct our main procedure.

Our first subroutine is an algorithm to compute the geometric resolution of the finite set $\mathcal{P}_{S,\sigma}$ containing $\Pi_x(V_{S,\sigma} \cap \{t = 1\})$ introduced in the previous section. This algorithm relies on the global Newton lifting from [7] and it is essentially the procedure underlying [15, Proposition 13]; we include it here for completeness. In order to simplify notation, we assume that $S = \{1, \dots, s\}$ and $\sigma = \{+\}^S$.

Algorithm GeometricResolution

INPUT: Polynomials $f_1, \dots, f_s, g \in \mathbf{K}[x_1, \dots, x_n]$ encoded by an slp of length L , an even integer $d \geq \deg(f_i), \deg(g)$, and a linear form $\ell_\alpha \in \mathbf{K}[x_1, \dots, x_n]$.

OUTPUT: The geometric resolution $(p, v_1, \dots, v_n)$ associated to ℓ_α of a finite set $\mathcal{P}$ containing $\Pi_x(V \cap \{t = 1\})$.

1. Compute the geometric resolution associated to $\ell_\alpha(x) = \alpha_1 x_1 + \dots + \alpha_n x_n$ of the variety defined in $\mathbb{A}^{n+s}$ by the (dehomogenized) system $\tilde{f}_1, \dots, \tilde{f}_s, \tilde{g}_1, \dots, \tilde{g}_n$ as follows:
 - (a) For every $B \subset \{1, \dots, n\}$ and $e : B \rightarrow \{-1, 1\}$, compute the geometric resolution associated to $\ell_\alpha(x)$ of the variety $S_{B,e}$.
 - (b) Compute the geometric resolution associated to $\ell_\alpha(x)$ of the variety $\bigcup_{B,e} S_{B,e}$.
2. Compute the geometric resolution associated to $\ell(x, y) = y_1 x_1 + \dots + y_n x_n$ of the variety defined by the (dehomogenized) system $\tilde{f}_1, \dots, \tilde{f}_s, \tilde{g}_1, \dots, \tilde{g}_n$ over $\mathbf{K}(y)$, modulo the ideal $(y_1 - \alpha_1, \dots, y_n - \alpha_n)^2$.
3. Compute $P(t, u, y) \bmod ((t)^{2nD_s+1} + (y_1 - \alpha_1, \dots, y_n - \alpha_n)^2) \mathbf{K}[[t]][u, y]$.
4. Compute $\hat{P}(1, u, \alpha) = \sum_{0 \leq h \leq D_s} p_h(1, \alpha) u^h$ and $\frac{\partial \hat{P}}{\partial y_j}(1, u, \alpha) = \sum_{0 \leq h \leq D_s} \frac{\partial p_h}{\partial y_j}(1, \alpha) u^h$ as follows:
 - (a) Compute $p_h(t, \alpha)$ and $\frac{\partial p_h}{\partial y_j}(t, \alpha)$ ($1 \leq j \leq n$, $0 \leq h \leq D_s$).
 - (b) Evaluate $t = 1$.
5. Compute $Q(u, \alpha) = \gcd(\hat{P}(1, u, \alpha), \frac{\partial \hat{P}}{\partial u}(1, u, \alpha))$, $p(u) = \frac{\hat{P}(1, u, \alpha)}{Q(u, \alpha)}$ and, for every $1 \leq j \leq n$, $v_j(u) := -\frac{\frac{\partial \hat{P}}{\partial y_j}(1, u, \alpha)}{Q(u, \alpha)} \left(\frac{\partial \hat{P}}{\partial u}(1, u, \alpha) \right)^{-1} \bmod p(u)$.

Proposition 8 Given a generic $\alpha \in \mathbf{K}^n$ and polynomials $f_1, \dots, f_m \in \mathbf{K}[x_1, \dots, x_n]$ of degree bounded by an even integer d and encoded by an slp of length L , Algorithm GeometricResolution computes the geometric resolution associated to the linear form $\ell_\alpha(x) = \sum_{1 \leq j \leq n} \alpha_j x_j$ of a finite set $\mathcal{P}_{S,\sigma}$ with at most D_s points, containing $\Pi_x(V_{S,\sigma} \cap \{t = 1\})$, within complexity $O(n^3(L + dn + n^{\Omega-1})D_s^2 \log^2(D_s) \log \log^2(D_s))$.

Proof STEP 1(a). The variety $S_{B,e}$ is defined by a square polynomial system in separated variables; then, the required computation can be achieved as in [14, Section 5.2.1] within complexity $O(D_{B,e}^2 \log^2(D_{B,e}) \log \log(D_{B,e}))$, where $D_{B,e}$ is the cardinality of $S_{B,e}$.

STEP 1(b). This step can be done within complexity $O(nD_s \log^3(D_s) \log \log(D_s))$ following the procedure in Sect. 2.2 and the strategy described in [26, Algorithm 10.3].

STEP 2. This step can be done applying [7, Algorithm 1] within complexity $O((dn^3 + n^{\Omega+1}) D_s \log(D_s) \log \log(D_s))$.

STEP 3. Since $F_1, \dots, F_s, G_1, \dots, G_n$ can be encoded by an slp of length $O(L + (d + s)n)$, a geometric resolution of the variety they define associated with the linear form $\ell(x, y)$ modulo the ideal $(t)^{2nD_s+1} + (y_1 - \alpha_1, \dots, y_n - \alpha_n)^2$ can be obtained from the previously computed geometric resolution by applying [7, Algorithm 1] within complexity $O(n^3(L + dn + n^{\Omega-1})D_s^2 \log^2(D_s) \log \log^2(D_s))$.

STEP 4(a). By expanding $P(t, u, y) = \sum_{0 \leq h \leq D_s} \frac{p_h(t, y)}{q(t)} u^h \in \mathbf{K}[[t]][u, y]$ into powers of u , $(y_1 - \alpha_1), \dots, (y_n - \alpha_n)$, we have that the coefficients corresponding to u^h and $u^h(y_j - \alpha_j)$ ($1 \leq j \leq n$, $0 \leq h \leq D_s$) are $p_h(t, \alpha)/q(t)$ and $\frac{\partial p_h}{\partial y_j}(t, \alpha)/q(t)$ respectively. As the degrees of the polynomials involved in these fractions are bounded by nD_s , they are uniquely determined by their power series expansions modulo $(t)^{2nD_s+1} \mathbf{K}[[t]]$ (see [26, Corollary 5.21]) that were computed at Step 3. By using [26, Corollary 5.24 and Algorithm 11.4] and converting all rational fractions to a common denominator, the computation is done within complexity $O(n^2 D_s^2 \log^2(D_s) \log \log(D_s))$.

STEP 5. This step is achieved by means of the Extended Euclidean algorithm and polynomial divisions with remainder within complexity $O(nD_s \log^2(D_s) \log \log(D_s))$. $\square$

The second subroutine presented here tells us, for a finite set $\mathcal{P}_{S,\sigma}$ given by a geometric resolution, if the set $\mathcal{P}_{S,\sigma} \cap E$ is empty, and, if not, it computes the Thom encoding of all the real roots of the minimal polynomial $p_{S,\sigma}$ corresponding to the points where the minimum value of g on $\mathcal{P}_{S,\sigma} \cap E$ is attained.

Algorithm MinimumInGeometricResolution

INPUT: A geometric resolution $(p_{S,\sigma}, v_{S,\sigma,1}, \dots, v_{S,\sigma,n})$ in $\mathbf{K}[u]$ of a finite set $\mathcal{P}_{S,\sigma}$, polynomials $f_1, \dots, f_m, g \in \mathbf{K}[x_1, \dots, x_n]$ encoded by an slp of length L , an integer $0 \leq l \leq m$ and an integer $d \geq \deg(f_i), \deg(g)$.

OUTPUT: A boolean variable “Empty” with the truth value of the statement “The set $\mathcal{P}_{S,\sigma} \cap E$ is empty” and a list of elements $\tau_1, \dots, \tau_k \in \{-1, 0, 1\}^{\deg p_{S,\sigma}-1}$ with $k = 0$ if Empty = True, representing the Thom encodings of all the real roots of $p_{S,\sigma}$ corresponding to the points where the minimum value of g on $\mathcal{P}_{S,\sigma} \cap E$ is attained.

1. Compute the list of realizable sign conditions for $f_1(v_{S,\sigma}(u)), \dots, f_m(v_{S,\sigma}(u))$ over the real zeros of $p_{S,\sigma}(u)$.
2. Determine Empty going through the obtained list of realizable sign conditions: Empty = False if and only the list computed in step 1 contains an element $\eta = (\eta_1, \dots, \eta_m)$ such that $\eta_i = 0$ for every $1 \leq i \leq l$ and $\eta_i \in \{0, 1\}$ for every $l+1 \leq i \leq m$.
3. If Empty = False:
 - (a) Compute $h(u) := \text{Res}_{\tilde{u}}(p_{S,\sigma}(\tilde{u}), u - g(v_{S,\sigma}(\tilde{u})))$.
 - (b) Compute the list of realizable sign conditions for $f_1(v_{S,\sigma}(u)), \dots, f_m(v_{S,\sigma}(u)), p'_{S,\sigma}(u), \dots, p_{S,\sigma}^{(\deg p_{S,\sigma}-1)}(u), h'(g(v_{S,\sigma}(u))), \dots, h^{(\deg p_{S,\sigma}-1)}(g(v_{S,\sigma}(u)))$ over the real zeros of $p_{S,\sigma}(u)$.
 - (c) Determine $\tau_1, \dots, \tau_k$ going through the obtained list of realizable sign conditions.

Proposition 9 *Given a geometric resolution $(p_{S,\sigma}, v_{S,\sigma,1}, \dots, v_{S,\sigma,n})$ of a set $\mathcal{P}_{S,\sigma}$ with at most D_s points and the polynomials $f_1, \dots, f_m, g \in \mathbf{K}[x_1, \dots, x_n]$ encoded by an slp of length L , Algorithm MinimumInGeometricResolution decides whether $\mathcal{P}_{S,\sigma} \cap E$ is empty or not and computes the Thom encodings of the real roots of $p_{S,\sigma}$ corresponding to the points where the minimum value of g on $\mathcal{P}_{S,\sigma} \cap E$ is attained within complexity $O(nd_s^2 + LdD_s + (m + D_s)D_s^2 \log^3(D_s))$.*

Proof STEP 1. First we compute the dense encoding of the polynomials $f_1(v_{S,\sigma}(u)), \dots, f_m(v_{S,\sigma}(u))$ and their remainders in the division by $p_{S,\sigma}$, within complexity $O((nD_s + L)dD_s + mdD_s \log^2(D_s) \log \log(D_s))$. Note that the value that a univariate polynomial takes at a root of $p_{S,\sigma}$ coincides with the value that its remainder in the division by $p_{S,\sigma}$ takes at the same root. Then we apply to the remainders computed the sign determination algorithm from [1, Section 10.3] and [5], following [22, Corollary 2], within complexity $O(mD_s^2 \log^3(D_s))$.

STEP 2. This step can be done within complexity $O(mD_s)$.

STEP 3(a). At this step the algorithm computes the monic polynomial h whose roots are the values of g at the points in $\mathcal{P}_{S,\sigma}$. First we compute the dense encoding of the polynomial $g(v_{S,\sigma}(u))$, its remainder $\tilde{g}_{S,\sigma}$ in the division by $p_{S,\sigma}$ and then, the resultant polynomial (using $\tilde{g}_{S,\sigma}$ instead of $g(v_{S,\sigma}(u))$) by multi-point evaluation and interpolation within complexity $O((nD_s + L)dD_s + D_s^2 \log^2(D_s) \log \log(D_s))$.

STEP 3(b). We continue the sign determination algorithm adding the polynomials $p'_{S,\sigma}(u), \dots, p_{S,\sigma}^{(\deg p_{S,\sigma} - 1)}(u)$, and the remainders of $h'(g(v_{S,\sigma}(u))), \dots, h^{(\deg p_{S,\sigma} - 1)}(g(v_{S,\sigma}(u)))$ in the division by $p_{S,\sigma}$ to what we have already computed at Step 1. Note that, for a fixed $1 \leq j \leq \deg p_{S,\sigma} - 1$, by evaluating $h^{(j)}$ at $\tilde{g}_{S,\sigma}(u)$ modulo $p_{S,\sigma}$, we can obtain the dense encoding of the remainder of $h^{(j)}(g(v_{S,\sigma}(u)))$ in the division by $p_{S,\sigma}$ within complexity $O(D_s^2 \log(D_s) \log \log(D_s))$. Then, the whole step can be done within complexity $O((m + D_s)D_s^2 \log^3(D_s))$.

STEP 3(c). The list of sign conditions computed at Step 3(b) enables us to know the Thom encoding of every real root ξ of $p_{S,\sigma}$, and to relate ξ with the Thom encoding of $g(v_{S,\sigma}(\xi))$ as a root of h . This information is enough to compare the different values of $g(v_{S,\sigma}(\xi))$ (see Sect. 2.3) and, so, we can give the Thom encoding as roots of $p_{S,\sigma}$ of the roots giving the points where the minimum value of g is attained. This step is done within complexity $O((m + D_s)D_s)$. $\square$

Since we know that the minimum value of g on E is also the minimum value of g on

$$\bigcup_{(S,\sigma) \in \mathcal{S}} \Pi_x(V_{S,\sigma} \cap \{t = 1\}) \cap E,$$

the set of sample minimizing points will be obtained by comparing the minimum values that g takes on $\mathcal{P}_{S,\sigma} \cap E$ for all $(S, \sigma) \in \mathcal{S}$. This task is achieved by the following subroutine.

Algorithm ComparingMinima

INPUT: Geometric resolutions $(p_{S_1,\sigma_1}, v_{S_1,\sigma_1,1}, \dots, v_{S_1,\sigma_1,n})$ and $(p_{S_2,\sigma_2}, v_{S_2,\sigma_2,1}, \dots, v_{S_2,\sigma_2,n})$ of finite sets $\mathcal{P}_{S_1,\sigma_1}$ and $\mathcal{P}_{S_2,\sigma_2}$ associated with a linear form ℓ_α , $g \in$

$\mathbf{K}[x_1, \dots, x_n]$ encoded by an slp of length L , an integer $d \geq \deg(g)$, and Thom encodings $\tau_1 \in \{-1, 0, 1\}^{\deg p_{S_1, \sigma_1} - 1}$ and $\tau_2 \in \{-1, 0, 1\}^{\deg p_{S_2, \sigma_2} - 1}$ of real roots of p_{S_1, σ_1} and p_{S_2, σ_2} corresponding to points where the minimum of g on $\mathcal{P}_{S_1, \sigma_1} \cap E$ and $\mathcal{P}_{S_2, \sigma_2} \cap E$ is attained.

OUTPUT: An integer “Sign” from the set $\{-1, 0, 1\}$ representing the sign of the minimum value that g takes on $\mathcal{P}_{S_1, \sigma_1} \cap E$ minus the minimum value that g takes on $\mathcal{P}_{S_2, \sigma_2} \cap E$.

1. Compute a geometric resolution $(p, v_1, \dots, v_n)$ of the union of the sets described by $(p_{S_1, \sigma_1}, v_{S_1, \sigma_1, 1}, \dots, v_{S_1, \sigma_1, n})$ and $(p_{S_2, \sigma_2}, v_{S_2, \sigma_2, 1}, \dots, v_{S_2, \sigma_2, n})$.
2. Compute $h(u) := \text{Res}_{\tilde{u}}(p(\tilde{u}), u - g(v(\tilde{u})))$.
3. Compute the list of realizable sign conditions for $p_{S_1, \sigma_1}(u), p'_{S_1, \sigma_1}(u), \dots, p_{S_1, \sigma_1}^{(\deg p_{S_1, \sigma_1} - 1)}(u), p_{S_2, \sigma_2}(u), p'_{S_2, \sigma_2}(u), \dots, p_{S_2, \sigma_2}^{(\deg p_{S_2, \sigma_2} - 1)}(u), h'(g(v(u))), \dots, h^{(\deg p)}(g(v(u)))$ over the real zeros of $p(u)$.
4. Determine Sign going through the obtained list of realizable sign conditions.

Proposition 10 *Given geometric resolutions $(p_{S_1, \sigma_1}, v_{S_1, \sigma_1, 1}, \dots, v_{S_1, \sigma_1, n})$ of a finite set $\mathcal{P}_{S_1, \sigma_1}$ with at most D_{s_1} points and $(p_{S_2, \sigma_2}, v_{S_2, \sigma_2, 1}, \dots, v_{S_2, \sigma_2, n})$ of a finite set $\mathcal{P}_{S_2, \sigma_2}$ with at most D_{s_2} , and the Thom encodings of a real root of p_{S_1, σ_1} and p_{S_2, σ_2} corresponding to points where the minimum of g on $\mathcal{P}_{S_1, \sigma_1} \cap E$ and $\mathcal{P}_{S_2, \sigma_2} \cap E$ is attained, Algorithm ComparingMinima compares these minima within complexity $O(ndD_{s_1, s_2}^2 + LdD_{s_1, s_2} + D_{s_1, s_2}^3 \log^3(D_{s_1, s_2}))$ where $D_{s_1, s_2} := \max\{D_{s_1}, D_{s_2}\}$ with $s_1 = |S_1|$ and $s_2 = |S_2|$.*

Proof STEP 1. This step is achieved within complexity $O(nD_{s_1, s_2} \log^2(D_{s_1, s_2}) \log \log(D_{s_1, s_2}))$ as explained in Sect. 2. STEPS 2, 3 AND 4. Similar to Algorithm MinimumInGeometricResolution Step 3(a), (b) and (c). The overall complexity of these steps is $O(ndD_{s_1, s_2}^2 + LdD_{s_1, s_2} + D_{s_1, s_2}^3 \log^3(D_{s_1, s_2}))$. $\square$

We give now the main algorithm of the paper.

Algorithm FindingMinimum

INPUT: Polynomials $f_1, \dots, f_m, g \in \mathbf{K}[x_1, \dots, x_n]$ encoded by an slp of length L , an integer $0 \leq l \leq m$ and an even integer $d \geq \deg(f_i), \deg(g)$.

OUTPUT: A family $\{(p_i, v_{i, 1}, \dots, v_{i, n}), \tau_i\}_{i \in \mathcal{I}}$ where $\mathcal{I}$ is a finite set and for every $i \in \mathcal{I}$, $(p_i, v_{i, 1}, \dots, v_{i, n})$ is a geometric resolution in $\mathbf{K}[u]$ and $\tau_i \in \{-1, 0, 1\}^{\deg p_i}$ is the Thom encoding of a real root ξ of p_i .

1. Take $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbf{K}^n$ at random and set $\ell_\alpha := \alpha_1 x_1 + \dots + \alpha_n x_n$.
2. $\mathcal{S} = \{(S, \sigma) \mid S \subset \{1, \dots, m\} \text{ with } 0 \leq |S| \leq n \text{ and } \sigma \in \{+, -\}^S \text{ with } \sigma_i = + \text{ for } l+1 \leq i \leq m\}$.
3. Take $(S_1, \sigma_1) \in \mathcal{S}$ and remove it from $\mathcal{S}$.
4. $(p_{S_1, \sigma_1}, v_{S_1, \sigma_1, 1}, \dots, v_{S_1, \sigma_1, n}) = \text{GeometricResolution}(f_i (i \in S_1), g, \sigma_1, d, \ell_\alpha)$.
5. (Empty, $\tau_{S_1, \sigma_1, 1}, \dots, \tau_{S_1, \sigma_1, k}$) = MinimumInGeometricResolution($p_{S_1, \sigma_1}, v_{S_1, \sigma_1, 1}, \dots, v_{S_1, \sigma_1, n}, f_1, \dots, f_m, g, l, d$).

6. While Empty = True:
 - (a) Discard (S_1, σ_1) , take a new (S_1, σ_1) and remove it from $\mathcal{S}$.
 - (b) $(p_{S_1, \sigma_1}, v_{S_1, \sigma_1, 1}, \dots, v_{S_1, \sigma_1, n}) = \text{GeometricResolution}(f_i(i \in S_1), g, \sigma_1, d, \ell_\alpha)$.
 - (c) $(\text{Empty}, \tau_{S_1, \sigma_1, 1}, \dots, \tau_{S_1, \sigma_1, k}) = \text{MinimumInGeometricResolution}(p_{S_1, \sigma_1}, v_{S_1, \sigma_1, 1}, \dots, v_{S_1, \sigma_1, n}, f_1, \dots, f_m, g, l, d)$.
7. $\mathcal{I} = \{((p_{S_1, \sigma_1}, v_{S_1, \sigma_1, 1}, \dots, v_{S_1, \sigma_1, n}), \tau_{S_1, \sigma_1, 1}), \dots, ((p_{S_1, \sigma_1}, v_{S_1, \sigma_1, 1}, \dots, v_{S_1, \sigma_1, n}), \tau_{S_1, \sigma_1, k})\}$.
8. While $\mathcal{S} \neq \emptyset$:
 - (a) Take $(S_2, \sigma_2) \in \mathcal{S}$ and remove it from $\mathcal{S}$.
 - (b) $(p_{S_2, \sigma_2}, v_{S_2, \sigma_2, 1}, \dots, v_{S_2, \sigma_2, n}) = \text{GeometricResolution}(f_i(i \in S_2), g, \sigma_2, d, \ell_\alpha)$.
 - (c) $(\text{Empty}, \tau_{S_2, \sigma_2, 1}, \dots, \tau_{S_2, \sigma_2, k}) = \text{MinimumInGeometricResolution}(p_{S_2, \sigma_2}, v_{S_2, \sigma_2, 1}, \dots, v_{S_2, \sigma_2, n}, f_1, \dots, f_m, g, l, d)$.
 - (d) If Empty = False :
 - i. Sign = ComparingMinima($p_{S_1, \sigma_1}, v_{S_1, \sigma_1, 1}, \dots, v_{S_1, \sigma_1, n}, p_{S_2, \sigma_2}, v_{S_2, \sigma_2, 1}, \dots, v_{S_2, \sigma_2, n}, g, d, \tau_{S_1, \sigma_1, 1}, \tau_{S_2, \sigma_2, 1}$).
 - ii. If Sign = 0 then $\mathcal{I} = \mathcal{I} \cup \{((p_{S_2, \sigma_2}, v_{S_2, \sigma_2, 1}, \dots, v_{S_2, \sigma_2, n}), \tau_{S_2, \sigma_2, 1}), \dots, ((p_{S_2, \sigma_2}, v_{S_2, \sigma_2, 1}, \dots, v_{S_2, \sigma_2, n}), \tau_{S_2, \sigma_2, k'})\}$.
 - iii. If Sign = 1
 - A. $\mathcal{I} = \{((p_{S_2, \sigma_2}, v_{S_2, \sigma_2, 1}, \dots, v_{S_2, \sigma_2, n}), \tau_{S_2, \sigma_2, 1}), \dots, ((p_{S_2, \sigma_2}, v_{S_2, \sigma_2, 1}, \dots, v_{S_2, \sigma_2, n}), \tau_{S_2, \sigma_2, k'})\}$.
 - B. $(S_1, \sigma_1) = (S_2, \sigma_2)$.

Now we prove Theorem 1.

Proof of Theorem 1. The correctness of the algorithm follows from the results in Sects. 3 and 4. The complexity upper bound is obtained by adding the complexity bounds for the subroutines `GeometricResolution` and `MinimumInGeometricResolution` applied to every element in $\mathcal{S}$, and `ComparingMinima` applied successively to pairs of elements from $\mathcal{S}$. $\square$

Acknowledgments We thank the reviewers for their helpful comments and suggestions. The authors were partially supported by the following grants: PIP 099/11 CONICET and UBACYT 20020090100069 (2010–2013) and UBACYT 20020120100133BA (2013–2016).

References

1. Basu, S., Pollack, R., Roy, M.-F.: Algorithms in real algebraic geometry. In: Algorithms and Computation in Mathematics, vol. 10, 2nd edn. Springer, Berlin (2006)
2. Baur, W., Strassen, V.: The complexity of partial derivatives. *Theor. Comput. Sci.* **22**(3), 317–330 (1983)
3. Berkowitz, S.: On computing the determinant in small parallel time using a small number of processors. *Inf. Process. Lett.* **18**(3), 147–150 (1984)
4. Bürgisser, P., Clausen, M., Shokrollahi, M.A.: Algebraic Complexity Theory. Grundlehren der Mathematischen Wissenschaften, vol. 315. Springer, Berlin (1997)
5. Canny, J.: Improved algorithms for sign determination and existential quantifier elimination. *Comput. J.* **36**(5), 409–418 (1993)

6. Emiris, I., Mourrain, B., Tsigaridas, E.: The DMM bound: multivariate (aggregate) separation bounds. In: ISSAC 2010—Proceedings of the 2010 International Symposium on Symbolic and Algebraic Computation, pp. 243–250. ACM, New York (2010)
7. Giusti, M., Lecerf, G., Salvy, B.: A Gröbner free alternative for polynomial system solving. *J. Complexity* **17**(1), 154–211 (2001)
8. Greuet, A., Guo, F., Safey El Din, M., Zhi, L.: Global optimization of polynomials restricted to a smooth variety using sums of squares. *J. Symb. Comput.* **47**(5), 503–518 (2012)
9. Greuet, A., Safey El Din, M.: Deciding reachability of the infimum of a multivariate polynomial. In: ISSAC 2011—Proceedings of the 36th International Symposium on Symbolic and Algebraic Computation, pp. 131–138. ACM, New York (2011)
10. Guo, F., Safey El Din, M., Zhi, L.: Global optimization of polynomials using generalized critical values and sums of squares. In: ISSAC 2010—Proceedings of the 2010 International Symposium on Symbolic and Algebraic Computation, pp. 107–114. ACM, New York (2010)
11. Heintz, J., Krick, T., Puddu, S., Sabia, J., Waissbein, A.: Deformation techniques for efficient polynomial equation solving. *J. Complexity* **16**(1), 70–109 (2000)
12. Heintz, J., Schnorr, C.-P.: Testing polynomials which are easy to compute. In: Logic and Algorithmic (Zurich, 1980). Monographs of L'Enseignement Mathématique, vol. 30, pp. 237–254. University of Genève, Geneva (1982)
13. Jeronimo, G., Perrucci, D.: On the minimum of a positive polynomial over the standard simplex. *J. Symb. Comput.* **45**(4), 434–442 (2010)
14. Jeronimo, G., Matera, G., Solernó, P., Waissbein, A.: Deformation techniques for sparse systems. *Found. Comput. Math.* **9**(1), 1–50 (2009)
15. Jeronimo, G., Perrucci, D., Sabia, J.: On sign conditions over real multivariate polynomials. *Discrete Comput. Geom.* **44**(1), 195–222 (2010)
16. Jeronimo, G., Perrucci, D., Tsigaridas, E.: On the minimum of a polynomial function on a basic closed semialgebraic set and applications. *SIAM J. Optim.* **23**(1), 241–255 (2013)
17. Kincaid, D., Cheney, W.: Numerical Analysis. Mathematics of Scientific Computing. Brooks/Cole Publishing Co., Pacific Grove, CA (1991)
18. Lasserre, J.B.: Global optimization with polynomials and the problem of moments. *SIAM J. Optim.* **11**(3), 796–817 (2001)
19. Nie, J., Demmel, J., Sturmfels, B.: Minimizing polynomials via sum of squares over the gradient ideal. *Math. Program.* **106**(3), 587–606 (2006)
20. Parrilo, P.A.: Semi-definite relaxations for semi-algebraic problems. *Math. Program.* **92**(2), 293–320 (2003)
21. Parrilo, P.A., Sturmfels, B.: Minimizing polynomial functions. In: Algorithmic and Quantitative Real Algebraic Geometry, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 60, pp. 83–99. American Mathematical Society, Providence, RI (2003)
22. Perrucci, D.: Linear solving for sign determination. *Theor. Comput. Sci.* **412**(35), 4715–4720 (2011)
23. Safey El Din, M.: Computing the global optimum of a multivariate polynomial over the reals. In: ISSAC, pp. 71–78. ACM, New York (2008)
24. Shafarevich, I.: Basic Algebraic Geometry, Study Edition. Springer, Berlin (1977)
25. Schweighofer, M.: Global optimization of polynomials using gradient tentacles and sums of squares. *SIAM J. Optim.* **17**(3), 920–942 (2006)
26. von zur Gathen, J., Gerhard, J.: Modern Computer Algebra. Cambridge University Press, New York (1999)