Continued Fraction Expansion of Real Roots of Polynomial Systems

Angelos Mantzaflaris, Bernard Mourrain, Elias Tsigaridas
GALAAD, INRIA Méditerranée
BP 93, 06902 Sophia Antipolis, France
[FirstName.LastName]@sophia.inria.fr

Abstract

We present a new algorithm for isolating the real roots of a system of multivariate polynomials, given in the *monomial basis*. It is inspired by existing subdivision methods in the Bernstein basis; it can be seen as generalization of the univariate continued fraction algorithm or alternatively as a fully analog of Bernstein subdivision in the monomial basis. The representation of the subdivided domains is done through *homographies*, which allows us to use only integer arithmetic and to treat efficiently unbounded regions. We use univariate bounding functions, projection and preconditioning techniques to reduce the domain of search. The resulting boxes have optimized rational coordinates, corresponding to the first terms of the *continued fraction expansion* of the real roots. An extension of *Vincent's Theorem* to multivariate polynomials is established and used to prove termination of the algorithm. New complexity bounds are provided for a simplified version of the algorithm and for a family of subdivision algorithms in terms of the real *condition number* of the system. Examples computed with our C++ implementation illustrate the approach.

1. Introduction

The problem of computing roots of univariate polynomials has a long mathematical history [18]. Recently, some new investigations focused on subdivision methods, where root localization is based on simple tests such as *Descartes' Rule of Signs* and its variant in the Bernstein basis [17, 11, 8]. Complexity analysis was developed for univariate integer polynomial (real) solving taking into account the bitsize of the coefficients, and providing a good understanding of their behavior from a theoretical and practical point of view. Approximation and bounding techniques have been developed [2] to improve the local speed of convergence to the roots.

Even more recently, new attention has been given to continued fraction algorithms (CF), see e.g. [20, 23] and references therein. They differ from previous subdivision-based algorithms in that instead of bisecting a given initial interval and thus producing a binary expansion of the real roots, they compute continued fraction expansions of these roots. The algorithm relies heavily on computations of lower bounds of the positive real roots, and different ways of computing such bounds lead to different variants of the algorithm. The best known worst-case complexity of CF is $\widetilde{\mathcal{O}}_B(d^5\tau^2)$ [20], while its average complexity is $\widetilde{\mathcal{O}}_B(d^3\tau)$, thus being the only complexity result that matches, even in the average the complexity bounds of numerical algorithms [19]. Moreover, the algorithm seems to be the most efficient in practice [10, 23].

Subdivision methods for the approximation of isolated roots of multivariate systems are also investigated but their analysis is much less advanced. In [21], the authors used tensor product representation in Bernstein basis and domain reduction techniques based on the convex hull property to speed up the convergence and reduce the number of subdivisions. In [9], the emphasis is put on the subdivision process, and stopping criterion based on the normal cone attached to the surface patch. In [16], this approach has been improved by introducing pre-conditioning and univariate-solver steps. The complexity of the method is also analyzed in terms of intrinsic differential invariants.

This work is in the spirit of [16]. The novelty of our approach is the presentation of a tensor-monomial basis algorithm that generalizes the univariate continued fraction algorithm and does not assume generic

position. We apply a subdivision approach also exploiting sign properties of the Bernstein polynomial representation, by proving a correspondence between the latter and a specific sequence of homography transformations.

Our contributions are as follows. We propose a new adaptive algorithm for polynomial system real solving that acts in monomial basis, and exploits the continued fraction expansion of (the coordinates of) the real roots. This yields the best rational approximation of the real roots. All computations are performed with integers, thus this is a division-free algorithm. We propose a first step towards the generalization of Vincent's theorem to the multivariate case (Th. 4.2). We perform a (bit) complexity analysis of the algorithm, when oracles for lower bounds and counting the real roots are available (Proposition 5.2) and we propose non-trivial improvements for reducing the total complexity even more (Sec. 5.3). In all cases the bounds that we derive for the multivariate case, match the best known ones for the univariate case, if we restrict ourselves to $n = 1$. Finally, using an inclusion test based on α-theorems (Sec. 6), we provide an output-sensitive complexity bound which involves the real condition number a of the system.

1.1. Notation

For a polynomial $f \in \mathbb{R}[x_1, .., x_n] = \mathbb{R}[\underline{x}]$, $\deg(f)$ denotes its total degree, while $\deg_{x_i}(f)$ denotes its degree with respect to x_i. Let $f(\underline{x}) = f(x_1, .., x_n) \in \mathbb{R}[x_1, .., x_n]$ with $\deg_{x_k} f = d_k$, $k = 1, .., n$. If not specified, we denote $d = d(f) = \max\{d_1, .., d_n\}$.

We are interested in isolating the real roots of a system of polynomials $f_1(\underline{x}), .., f_s(\underline{x}) \in \mathbb{Z}[x_1, .., x_n]$, in a box $I_0 = [u_1, v_1] \times \cdots \times [u_n, v_n] \subset \mathbb{R}^n$, $u_k, v_k \in \mathbb{Q}$. We denote by $\mathcal{Z}_{\mathbb{K}^n}(f) = \{p \in \mathbb{K}^n; f(p) = 0\}$ the solution set in $\mathbb{K}^n$ of the equation $f(x) = 0$, where $\mathbb{K}$ is $\mathbb{R}$ or $\mathbb{C}$.

For a homogeneous polynomial $f(x_0, \ldots, x_n) = \sum_{|\alpha|=d} c_\alpha \underline{x}^\alpha \in \mathbb{R}[x]$ of degree d, we define

$$\|f\| = \left(\sum_{|\alpha|=d} |c_\alpha|^2 \binom{d}{\alpha}^{-1} \right)^{\frac{1}{2}}.$$

For an affine polynomial $f(x_1, \ldots, x_n)$ of degree d, we define $\|f\|$ as the norm of its homogenization in degree d. For a system $f = (f_1, \ldots, f_n)$ of polynomials f_i of degree d_i, we define $\|f\| = (\|f_1\|^2 + \cdots + \|f_n\|^2)^{\frac{1}{2}}$. An important property of this norm is that it stays invariant by a unitary change of coordinates.

For $\underline{v} = (v_1, \ldots, v_n) \in \mathbb{C}^n$, $\|v\| := (v_1^2 + \cdots + v_n^2)^{\frac{1}{2}}$, $\|v\|_\infty := \max\{v_i; i = 1, \ldots, n\}$.

For $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{C}$, $x \in \mathbb{K}^n$ and $\rho > 0$, we denote by

- $B_{\mathbb{K}}(x, \rho) = \{y \in \mathbb{K}^n; \|y - x\| < \rho\}$ the ball of center x and radius ρ;

- $I_{\mathbb{K}}(x, \rho) = \{y \in \mathbb{K}^n; \|y - x\|_\infty < \rho\}$ the box of center x and radius ρ;

If $I = I_1 \times \cdots \times I_n \subset \mathbb{R}^n$, we denote by $I_{\mathbb{C}}$ the product $D_{\mathbb{C}}(I_1) \times \cdots \times D_{\mathbb{C}}(I_n) \subset \mathbb{C}^n$ of discs $D_{\mathbb{C}}(I_j) \subset \mathbb{C}$ of diameter I_j.

In what follows $\mathcal{O}_B$, resp. $\mathcal{O}$, means bit, resp. arithmetic, complexity and the $\widetilde{\mathcal{O}}_B$, resp. $\widetilde{\mathcal{O}}$, notation means that we are ignoring logarithmic factors. For $a \in \mathbb{Q}$, $\mathcal{L}(a) \geq 1$ is the maximum bit size of the numerator and the denominator. For a polynomial $f \in \mathbb{Z}[x_1, .., x_n]$, we denote by $\mathcal{L}(f)$ the maximum of the bitsize of its coefficients (including a bit for the sign). In the following, we will consider classes of polynomials such that $\log(d(f)) = \mathcal{O}(\mathcal{L}(f))$.

Also, to simplify the notation we introduce multi-indices, for the variable vector $\underline{x} = (x_1, .., x_n)$, $\underline{x}^{\underline{i}} := x_1^{i_1} \cdots x_n^{i_n}$, the sum $\sum_{\underline{i}=\underline{0}}^{\underline{d}} := \sum_{i_1=0}^{d_1} \cdots \sum_{i_n=0}^{d_n}$, and $\binom{\underline{d}}{\underline{i}} := \binom{d_1}{i_1} \cdots \binom{d_n}{i_n}$. The tensor Bernstein basis polynomials of multidegree degree $\underline{d}$ of a box I are denoted $B(\underline{x}; \underline{i}, \underline{d}; I) := B_{d_1}^{i_1}(x_1; u_1, u_1) \cdots B_{d_n}^{i_n}(x_n; u_n, u_n)$ where $I = [\underline{u}, \underline{v}] := [u_1, v_1] \times \cdots \times [u_n, v_n]$.

1.2. The general scheme

In this section, we describe the family of algorithms that we consider. The main ingredients are

- a suitable representation of the equations in a given (usually rectangular) domain, in terms of a basis of $\mathbb{Z}[\underline{x}]$, for instance a representation in the Bernstein basis or in the monomial basis;

- an algorithm to split the representation into smaller sub-domains;

- a reduction procedure to shrink the domain.

Different choices for each of these ingredients lead to algorithms with different practical behavior. The general process is summarized in Algorithm 1.

Algorithm 1.1: Subdivision scheme

Input: A set of equations $f_1, f_2, .., f_s \in \mathbb{Z}[\underline{x}]$ represented over a domain I.
Output: A list of disjoint domains, each containing one and only one real root of $f_1 = \cdots = f_s = 0$.
Initialize a stack Q and add $(I, f_1, .., f_s)$ on top of it;
While Q is not empty do
    a) Pop a system $(I, f_1, .., f_s)$ and:
    b) Perform a precondition process and/or a reduction process to refine the domain.
    c) Apply an exclusion test to identify if the domain contains no root.
    d) Apply an inclusion test to identify if the domain contains a single root. In this case output $(I, f_1, .., f_s)$.
    e) If both tests fail split the representation into a number of sub-domains and push them to Q.

The instance of this general scheme that we obtain generalizes the continued fraction method for univariate polynomials; the realization of the main steps (b-e) can be summarized as follows:

    b) Perform a precondition process and compute a lower bound on the roots of the system, in order to reduce the domain.
    c) Apply interval analysis or sign inspection to identify if some f_i has constant sign in the domain, i.e. if the domain contains no root.
    d) Apply Miranda test to identify if the domain contains a single root. In this case output $(I, f_1, .., f_s)$.
    e) If both tests fail, split the representation at $(1, .., 1)$ and continue.

In the following sections, we are going to describe more precisely these specific steps and analyze their complexity. In Sec. 2, we describe the representation of domains via homographies and the connection with the Bernstein basis representation. Subdivision, based on shifts of univariate polynomials, reduction and preconditioning are analyzed in Sec. 3. Exclusion and inclusion tests as well as a generalization of Vincent's theorem to multivariate polynomials, are presented in Sec. 4. In Sec. 5, we recall the main properties of Continued Fraction expansion of real numbers and use them to analyze the complexity of a subdivision algorithm following this generic scheme. In Sec. 6, we bound the complexity of the subdivision method using the α-inclusion test in terms of the real condition number of the system. We conclude with examples produced by our C++ implementation in Sec. 7.

2. Representation: Homographies

A widely used representation of a polynomial f over a rectangular domain is the tensor-Bernstein representation. De Casteljau's algorithm provides an efficient way to split this representation to smaller domains. A disadvantage is that converting integer polynomials to Bernstein form results in rational or, if one uses machine numbers, approximate Bernstein coefficients. We follow an alternative approach that does not require basis conversion since it applies to monomial basis: We introduce a tensor-monomial representation,

i.e. a representation in the monomial basis over $\mathbb{P}^1 \times \cdots \times \mathbb{P}^1$ and provide an algorithm to subdivide this representation analogously to the Bernstein case.

In a tensor-monomial representation a polynomial is represented as a tensor (higher dimensional matrix) of coefficients in the natural monomial basis, that is,

$$f(\underline{x}) = \sum_{i_1,..,i_n}^{d_1,..,d_n} c_{i_1..i_n} \underline{x}^{(i_1,...,i_n)} = \sum_{\underline{i}=\underline{0}}^{\underline{d}} c_{\underline{i}} \underline{x}^{\underline{i}},$$

for every equation f of the system. Splitting this representation is done using homographies. The main operation in this computation is the Taylor shift.

Definition 2.1. *A homography (or Mobius transformation) is a bijective projective transformation* $\mathcal{H} = (\mathcal{H}_1,..,\mathcal{H}_n)$, *defined over* $\mathbb{P}^1 \times \cdots \times \mathbb{P}^1$ *as*

$$x_k \mapsto \mathcal{H}_k(x_k) = \frac{\alpha_k x_k + \beta_k}{\gamma_k x_k + \delta_k}$$

with $\alpha_k, \beta_k, \gamma_k, \delta_k \in \mathbb{Z}$, $\gamma_k \delta_k \neq 0$, $k = 1,..,n$.

Using simple calculations, we can see that the inverse

$$\mathcal{H}_k^{-1}(x_k) = \frac{\delta_k x_k - \beta_k}{\gamma_k x_k - \alpha_k}$$

is also a homography, hence the set of homographies is a group under composition. Also, notice that if $\det \mathcal{H} > 0$ then, taking proper limits when needed, we can write

$$\mathbb{R}_+ \mapsto \mathcal{H}_k(\mathbb{R}_+) = \left[\frac{\beta_k}{\delta_k}, \frac{\alpha_k}{\gamma_k} \right] \tag{1}$$

hence $H(f) : \mathbb{R}_+^n \to \mathbb{R}$,

$$H(f) := \prod_{k=1}^{n} (\gamma_k x_k + \delta_k)^{d_k} \cdot (f \circ \mathcal{H})(x)$$

is a polynomial defined over $\mathbb{R}_+^n$ that corresponds to the (possibly unbounded) box

$$I_H = \mathcal{H}(\mathbb{R}_+^n) = \left[\frac{\beta_1}{\delta_1}, \frac{\alpha_1}{\gamma_1} \right] \times \cdots \times \left[\frac{\beta_n}{\delta_n}, \frac{\alpha_n}{\gamma_n} \right], \tag{2}$$

of the initial system, in the sense that the zeros of the initial system in I_H are in one-to-one correspondence with the positive zeros of $H(f)$.

We focus on the computation of $H(f)$. We use the basic homographies $T_k^c(f) = f|_{x_k = x_k + c}$ (translation by c) or simply $T_k(f)$ if $c = 1$, $C_k^c(f) = f|_{x_k = cx_k}$ (contraction by c) and $R_k(f) = x_k^{d_k} f|_{x_k = 1/x_k}$ (reciprocal polynomial). These notations are naturally extended to variable vectors; for instance $T^{\underline{c}} = (T_1^{c_1},..,T_n^{c_n})$, $c = (c_1,..,c_n) \in \mathbb{Z}^n$. Complexity results for these computations appear in the following sections. We can see that they suffice to compute any homography:

Lemma 2.2. *The group of homographies H with coefficients in $\mathbb{Z}$ is generated by* R_k, C_k^c, T_k^c, $k = 1,..,n$, $c \in \mathbb{Z}$.

PROOF. It can be verified that $H_k(f)$ is computed as

$$H_k(f) = C_k^{\gamma_k} R_k C_k^{\delta_k} T_k R_k C_k^{\beta_k/\delta_k - \alpha_k/\gamma_k} T_k^{\beta_k/\delta_k}(f)$$

where the product denotes composition. We abbreviate $C_k^{1/c} = R_k C_k^c R_k$ and $T_k^{1/c} = C_k^u T_k^1 R_k T_k^c R_k$.

Representation via homography is in an interesting correspondence to the Bernstein representation:

Lemma 2.3. *Let $f = \sum_{\underline{i}=0}^{d} b_i B_i^n(\underline{x}, I_H)$ the Bernstein expansion of f in the box I_H yielded by a homography H. If*

$$H(f) = C^\gamma RC^\delta T^1 RC^{\beta/\delta - \alpha/\gamma} T^{\beta/\delta}(f) = \sum_{\underline{i}=0}^{d} c_i \underline{x}^{\underline{i}}$$

then $c_i = \binom{d}{\underline{i}} \underline{\gamma}^{\underline{i}} \underline{\delta}^{\underline{d}-\underline{i}} b_i$.

PROOF. Let $[u_k, v_k] = \left[\frac{\alpha_k}{\gamma_k}, \frac{\beta_k}{\delta_k}\right]$. For a tensor-Bernstein polynomial $\binom{d}{\underline{i}} \frac{1}{(v-u)^d}(x-u)^i (v-x)^{d-i}$ we compute

$$C^\gamma RC^\delta T^1 RC^{v-u} T^u \left(\binom{d}{\underline{i}} \frac{1}{(v-u)^d}(x-u)^i (v-x)^{d-i}\right)$$

$$= C^\gamma RC^\delta RT^1 RC^{v-u} \left(\binom{d}{\underline{i}} \frac{1}{(v-u)^d} x^i (v-u-x)^{d-i}\right)$$

$$= C^\gamma RC^\delta RT^1 \left(\binom{d}{\underline{i}} (x-1)^{d-i}\right)$$

$$= C^\gamma RC^\delta \left(\binom{d}{\underline{i}} x^i\right) = \binom{d}{\underline{i}} \gamma^i \delta^{d-i} x^i$$

as needed.

Corollary 2.4. *The Bernstein expansion of f in I_H is*

$$\sum_{i=0}^{d} \frac{c_i}{\binom{d}{\underline{i}} \underline{\gamma}^{\underline{i}} \underline{\delta}^{\underline{d}-\underline{i}}} B(\underline{x}; \underline{i}, \underline{d}; I_H).$$

That is, the coefficients of $H(f)$ coincide with the Bernstein coefficients up to contraction and binomial factors.

Thus tensor-Bernstein coefficients and tensor-monomial coefficients in a sub-domain of $\mathbb{R}_+^n$ differ only by multiplication by positive constant. In particular they are of the same sign. Hence this corollary allows us to take advantage of sign properties (eg. the variation diminishing property) of the Bernstein basis that still hold in homography representation.

The resulting representation of the system consists of the transformed polynomials $H(f_1), .., H(f_n)$, represented as tensors of coefficients as well as $4n$ integers, $\alpha_k, \beta_k, \gamma_k, \delta_k$ for $k = 1, .., n$ from which we can recover the endpoints of the domain, using (2).

3. Subdivision and reduction

3.1. The subdivision step

We describe the subdivision step using the homography representation. This is done at a point $\underline{u} = (u_1, .., u_n) \in \mathbb{Z}_{\geq 0}^n$. It consists in computing up to 2^n new sub-domains (depending on the number of nonzero u_k's), each one having $\underline{u}$ as a vertex.

Given $H(f_1), .., H(f_s)$ that represent the initial system over some domain, we consider the partition of $\mathbb{R}_+^n$ defined by the hyperplanes $x_k = u_k$, $k = 1, .., n$. These intersect at $\underline{u}$ hence we call this *partition at $\underline{u}$*. Subdividing at $\underline{u}$ is equivalent to subdividing the initial domain into boxes that share the common vertex $\mathcal{H}(\underline{u})$ and have faces either parallel or perpendicular to those of the initial domain.

We need to compute a homography representation for every domain in this partition. The computation is done coordinate wise; observe that for any domain in this partition we have, for all k, either $x_k \in [0, u_k]$ or $x_k \in [u_k, \infty]$. It suffices to apply a transformation that takes these domains to $\mathbb{R}_+$. In the former case, we apply $T_k^1 R_k C_k^{u_k}$ to the current polynomials and in the latter case we shift them by u_k, i.e. we apply $T_k^{u_k}$. The integers $\alpha_k, \beta_k, \gamma_k, \delta_k$ that keep track of the current domain can be easily updated to correspond to the new subdomain.

We can make this process explicit in general dimension: every computed subdomain corresponds to a binary number of length n, where the $k-$th bit is 1 if $T_k^1 R_k C^{u_k}$ is applied or 0 if $T_k^{u_k}$ is applied.

In our continued fraction algorithm the subdivision is performed at $\underline{u} = \underline{1}$.

Illustration. Let us illustrate this process in dimension two. The system f_1, f_2 is defined over $\mathbb{R}^2_{>0}$. We subdivide this domain into $[0,1]^2$, $[0,1] \times \mathbb{R}_{>1}$, $\mathbb{R}_{>1} \times [0,1]$ and $\mathbb{R}_{>1} \times \mathbb{R}_{>1}$. Equivalently, we compute four new pairs of polynomials, as illustrated in Fig. 1 (we abbreviate $S_k = T_k^1 R_k$).


Figure 1: Subdividing the domain of $\underline{f}$.

Complexity of subdivision step. The transformation of a polynomial into two sub-domains, i.e. splitting with respect to one direction, consists in performing d^{n-1} univariate shifts, one for every coefficient $\in \mathbb{Z}[x_k]$ of $f \in \mathbb{Z}[x_k][x_1, .., \widehat{x_k}, .., x_n]$.

If the subdivision is performed in every direction, each transformation consists of d^{n-1} univariate shifts for every variable, i.e. nd^{n-1} shifts. There are 2^n sub-domains to compute, hence a total of $n2^n d^{n-1}$ shifts have to be performed in a single subdivision step. We must also take into account that every time a univariate shift is performed, the coefficient bitsize increases.

The operations

$$T_k(f) = f|_{x_k = x_k + 1} \quad \text{and} \quad T_k R_k(f) = (x_k + 1)^{d_k} f|_{x_k = \frac{1}{x_k + 1}}$$

are essentially of the same complexity, except that the second requires one to exchange the coefficient of $c_{i_1, .., i_k, .., i_n}$ with $c_{i_1, .., d_k - i_k, .., i_n}$ before translation, i.e. an additional $\mathcal{O}(d^n)$ cost. Hence we only need to consider the case of shifts for the complexity.

The continued fraction algorithm subdivides a domain using unit shifts and inversion. Successive operations of this kind increase the bitsize equivalently to a big shift by the sum of these units. Thus it suffices to consider the general computation of $f(\underline{x} + \underline{u})$ to estimate the complexity of the subdivision step.

Lemma 3.1 (Shift complexity). *The computation of $f(\underline{x} + \underline{u})$ with $\mathcal{L}(f) = \tau$ and $\mathcal{L}(u_k) \leq \sigma$, $k = 1, .., n$ can be performed in $\widetilde{\mathcal{O}}_B(n^2 d^n \tau + d^{n+1} n^3 \sigma)$.*

PROOF. We use known facts for the computation of $T_k^{u_k}(f)$ for univariate polynomials. If $\deg_k f = d_k$ and f is univariate, this operation is performed in $\widetilde{\mathcal{O}}_B(d_k^2 \sigma + d_k \tau)$; the resulting coefficients are of bitsize $\tau + d_k \sigma$ [25]. Hence $f(x_1, .., x_k + u_k, .., x_n)$ is computed in $\widetilde{\mathcal{O}}_B(d^{n-1}(d_k^2 \sigma + d_k \tau))$.

Suppose we have computed $f(x_1 + u_1, x_{k-1} + u_{k-1}, x_k, .., x_n)$ for some k. The coefficients are of bitsize $\tau + \sum_{i=1}^{k-1} \sigma_i$. The computation of shift with respect to $k-$th variable $f(x_1 + u_1, .., x_k + u_k, x_{k+1}, .., x_n)$ results in a polynomial of bitsize $\tau + \sum_{i=1}^{k} \sigma_i$ and consists of $d^{n-1} \widetilde{\mathcal{O}}_B(d^2 \sum_{i=1}^{k} \sigma_i + d\tau))$ operations. That is, we perform d^{n-1} univariate polynomial shifts, one for every coefficient of f in $\mathbb{Z}[x_k][x_1, .., \widehat{x_k}, .., x_n]$.

This gives a total cost for computing $f(\underline{x} + \underline{u})$ of

$$d^{n-1} \sum_{k=1}^{n} \left(d^2 \sum_{i=1}^{k} \sigma_i + d\tau \right) = nd^n\tau + d^{n+1} \sum_{k=1}^{n} (n+1-k)\sigma_k.$$

The latter sum implies that it is faster to apply the shifts with increasing order, starting with the smallest number u_k. Since $\sigma_k = \mathcal{O}(\sigma)$ for all k, and we must shift a system of $\mathcal{O}(n)$ polynomials we obtain the stated result.

Let us present an alternative way to compute a sub-domain using contraction, preferable when the bitsize of $\underline{u}$ is big. The idea behind this is the fact that T_k^c and $T_k^1 C^c$ compute the same sub-domain, in two different ways.

Lemma 3.2. *If $f = \sum_{\underline{i}=\underline{0}}^{d} c_{\underline{i}} \underline{x}^{\underline{i}}$, $\mathcal{L}(f) = \tau$, then the coefficients of $C^u(f)$, $\mathcal{L}(u_k) \le \sigma$, $k = 1, .., n$, can be computed in $\widetilde{\mathcal{O}}_B(d^n\tau + nd^{n+1}\sigma)$.*

PROOF. The operation, i.e. computing the new coefficients $c_{\underline{i}} \underline{u}^{\underline{i}}$ can be done with $\widetilde{\mathcal{O}}(d^n)$ multiplications: Since $\underline{u}^{(i_1,...,i_k,..,i_n)} = u_k \underline{u}^{(i_1,...,i_k-1,...,i_n)}$, if these powers are computed successively then every coefficient is computed using two multiplications. Moreover, it suffices to keep in memory the n powers $u^{(i_1,...,i_{k-1},i_k-1,i_{k+1},...,i_n)}$, $k = 1, .., n$ in order to compute any $\underline{u}^{\underline{i}} c_{\underline{i}}$. Geometrically this can be understood as a stencil of n points that sweeps the coefficient tensor and updates every element using one neighbor at each time. The bitsize of the multiplied numbers is $\mathcal{O}(\tau + d\sigma)$ hence the result follows.

Now if we consider a contraction followed by a shift by 1 with respect to x_k for $\mathcal{O}(n)$ polynomials we obtain $\widetilde{\mathcal{O}}_B(n^2 d^n\tau + n^3 d^{n+1} + nd^{n+1}\sigma)$ operations for the computation of the domain. The disadvantage is that the resulting coefficients are of bitsize $\mathcal{O}(\tau + d\sigma)$ instead of $\mathcal{O}(\tau + n\sigma)$ with the use of shifts. Also note that this operation would compute a expansion of the real root which differs from continued fraction expansion.

3.2. Reduction: Bounds on the range of f

In this section we define univariate polynomials whose graph in $\mathbb{R}^{n+1}$ bounds the graph of f. For every direction k, we provide two polynomials bounding the values of f in $\mathbb{R}^n$ from below and above respectively.

Define

$$m_k(f; x_k) = \sum_{i_k=0}^{d_k} \min_{i_1,..,\widehat{i_k},..,i_n} c_{i_1..i_n} x_k^{i_k} \tag{3}$$

$$M_k(f; x_k) = \sum_{i_k=0}^{d_k} \max_{i_1,..,\widehat{i_k},..,i_n} c_{i_1..i_n} x_k^{i_k} \tag{4}$$

Lemma 3.3. *For any $x \in \mathbb{R}_+^n$, $n > 1$ and any $k = 1, .., n$, we have*

$$m_k(f; x_k) \le \frac{f(x)}{\prod_{s \neq k} \sum_{i_s=0}^{d_s} x_s^{i_s}} \le M_k(f; x_k). \tag{5}$$

PROOF. For $x \in \mathbb{R}_+^n$, we can directly write

$$f(x) \le \left(\sum_{i_k=0}^{d_k} \max_{i_1,..,\widehat{i_k},..,i_n} c_{i_1..i_n} x_k^{i_k} \right) \prod_{s \neq k} \sum_{i_s=0}^{d_s} x_s^{i_s}$$

The product of power sums is greater than 1; divide both sides by it. Analogously for $M_k(f; x_k)$.

7

Corollary 3.4. *Given $k \in \{1, .., n\}$, if $u \in \mathbb{R}_+^n$ with $u_k \in]0, \mu_k]$, where*

$$
\mu_k = \begin{cases}
\text{min. pos. root of } M_k(f, x_k) & \text{if } M_k(f; 0) < 0 \\
\text{min. pos. root of } m_k(f, x_k) & \text{if } m_k(f; 0) > 0 \\
0 & \text{otherwise}
\end{cases},
$$

then $f(u) \neq 0$. Consequently, all positive roots of f lie in $\mathbb{R}_{>\mu_1} \times \cdots \times \mathbb{R}_{>\mu_n}$. Also, for $u \in \mathbb{R}_+^n$ with $u_k \in [\mathcal{M}_k, \infty]$,

$$
\mathcal{M}_k = \begin{cases}
\text{max. pos. root of } M_k(f, x_k) & \text{if } M_k(f; \infty) < 0 \\
\text{max. pos. root of } m_k(f, x_k) & \text{if } m_k(f; \infty) > 0 \\
\infty & \text{otherwise}
\end{cases},
$$

it is $f(u) \neq 0$, i.e. all pos. roots are in $\mathbb{R}_{<\mathcal{M}_1} \times \cdots \times \mathbb{R}_{<\mathcal{M}_n}$.

Combining both bounds we deduce that $[\mu_1, \mathcal{M}_1] \times \cdots \times [\mu_n, \mathcal{M}_n]$ is a bounding box for $f^{-1}(\{0\}) \cap \mathbb{R}_+^n$.

PROOF. The denominator in (5) is always positive in $\mathbb{R}_+^n$. Let $\underline{u} \in \mathbb{R}^n$ with $u_k \in [0, \mu_k]$. If $M_k(f, 0) < 0$ then also $M_k(f, u) < 0$ and it follows $f(\underline{u}) < 0$. Similarly $m_k(f, 0) > 0 \Rightarrow m_k(f, u) > 0 \Rightarrow f(\underline{u}) < 0$. The same arguments hold for $[\mathcal{M}_k, \infty]$, $M_k(f; \infty) = R(M_k(f; x_k))(0)$, $m_k(f; \infty) = R(m_k(f; x_k))(0)$, and $R(f)$, since lower bounds on the zeros of $R(f)$ yield upper bounds on the zeros of f.

Thus lower and upper bounds on the $k-$th coordinates of the roots of $(f_1, .., f_s)$ are given by

$$
\max_{i=1,..,s} \{\mu_k(f_i)\} \qquad \text{and} \qquad \min_{i=1,..,s} \{\mathcal{M}_k(f_i)\} \tag{6}
$$

respectively, i.e. the intersection of these bounding boxes.

We would like to remain in the ring of integers all along the process, thus integer lower or upper bounds will be used. These can be the floor or ceil of the above roots of univariate polynomials, or even known bounds for them, e.g. Cauchy's bound.

If the minimum and maximum are taken with the ordering of coefficients defined as $c_i \prec c_j \iff c_i \binom{d}{j} \gamma^j \delta^{d-j} < c_j \binom{d}{i} \gamma^i \delta^{d-i}$ then different $m_k(f, x_k)$, $M_k(f, x_k)$ polynomials are obtained. By Corollary 2.4 their control polygon is the lower and upper hull respectively of the projections of the tensor-Bernstein coefficients to the $k-$th direction and are known to converge quadratically to simple roots when preconditioning (described in the following paragraph) is utilized [16, Corollary 5.3].

Complexity analysis. The analysis of the subdivision step in Sec. 3.2 applies as well to the reduction step, since reducing the domain means to compute a new subdomain and ignore the remaining volume of the initial box.

If a lower bound $\underline{l} = (l_1, \ldots, l_n)$ is known, with $\mathcal{L}(l_k) = \widetilde{\mathcal{O}}(\sigma)$, then the reduction step is performed in $\widetilde{\mathcal{O}}_B(n^2 d^n \tau + d^{n+1} n^3 \sigma)$. This is an instance of Lemma 3.1.

The projections of Lemma 3.3 are computed using $\mathcal{O}(d^n)$ comparisons. The computation of $\underline{l}$ costs $\widetilde{\mathcal{O}}_B(d^3 \tau)$ in average, for solving these projections using univariate CF algorithm. Another option would be to compute well known lower bounds on their roots, for instance Cauchy's bound in $\mathcal{O}(d)$.

Illustration. Consider a bi-quadratic $f_0 \in \mathbb{R}[x, y]$, namely, $\deg_{x_1} f_0 = \deg_{x_2} f_0 = 2$ with coefficients c_{ij}. Suppose that $f_0 = H(f)$ for $I_0 = I_H$. We compute

$$
m_1(f, x_1) = \sum_{i=0}^{2} \min_{j=0,\ldots,2} c_{ij} x_1^i \quad \text{and} \quad M_1(f, x_1) = \sum_{i=0}^{2} \max_{j=0,\ldots,2} c_{ij} x_1^i.
$$

thus $m(f, x_1) \leq \frac{f(x_1, x_2)}{1 + x_2 + x_2^2} \leq M_1(f, x_1)$. Fig. 2 shows how these univariate quadratics bound the graph of f in I_0.

Figure 2: The enveloping polynomials $M_1(x)$, $m_1(x)$ in domain I_0 for a bi-quadratic polynomial $f(x,y)$.

3.3. Preconditioning

To improve the reduction step, we use preconditioning. The aim of a preconditioner is to tune the system so that it can be tackled more efficiently; in our case we aim at improving the bounds of Corollary 3.4.

A preconditioning matrix P is an invertible $s \times s$ matrix that transforms a system $(f_1, .., f_s)^t$ into the equivalent one $P \cdot (f_1, .., f_s)^t$. This transformation does not alter the roots of the system, since the computed equations generate the same ideal. The bounds obtained on the resulting system can be used directly to reduce the domain of the equations before preconditioning. Preconditioning can be performed to a subset of the equations.

Since we use a reduction process using Corollary 3.4 we want to have among our equations n of them whose zero locus $f^{-1}(\{0\})$ is orthogonal to the $k-$th direction, for all k.

Assuming a square system, we precondition $H(f_1), .., H(f_n)$ to obtain a locally orthogonal to the axis system; an ideal preconditioner would be the Jacobian of the system evaluated at a common root; instead, we evaluate $J_{H(f)}$ at the image of the center $\underline{u}$ of the initial domain I_H, $u_k = \frac{\alpha_k \delta_k + \beta_k \gamma_k}{2 \gamma_k \delta_k}$. Thus we must compute the inverse of the Jacobian matrix $J_{H(\underline{f})}(x) = [\partial_{x_i} H(f_j)(x)]_{1 \leq i,j, \leq n}$ evaluated at $\underline{u}' := \mathcal{H}(\underline{u}) = (\delta_1/\gamma_1, .., \delta_n/\gamma_n)$.

Precondition step complexity. Computing $J_{H(f)}(\underline{u}) \cdot (H(f_1), .., H(f_n))^t$ is done with cost $\widetilde{\mathcal{O}}_B(n^2 d^n)$ and evaluating at u' has cost $\widetilde{\mathcal{O}}_B(n^2 d^{n-1})$. We also need $\widetilde{\mathcal{O}}_B(n^2)$ for inversion and $\mathcal{O}(n^2 d^n)$ for multiplying polynomials times scalar as well as summing polynomials. This gives a precondition cost of order $\mathcal{O}(n^2 d^n)$.

4. Regularity tests

A subdivision scheme 1 is able to work when two tests are available: one that identifies empty domains (exclusion test) and one that identifies domains with exactly one zero (inclusion test). If these two tests are negative, a domain cannot be neither included nor excluded so we need to apply further reduction/subdivision steps to it. Nevertheless, if the result of the test is affirmative, then this is certified to be correct.

4.1. Exclusion test

The bounding functions defined in the previous section provide a fast filter to exclude empty domains. Define $\min \varnothing = \infty$ and $\max \varnothing = 0$.

Corollary 4.1. *If for some $k \in \{1, .., n\}$ and for some $i \in \{1, .., s\}$ it is $\mu_k(f_i) = \infty$ or $\mathcal{M}_k(f_i) = 0$ then the system has no solutions. Also, if $\max_{i=1,..,s}\{\mu_k(f_i)\} > \min_{i=1,..,s}\{\mathcal{M}_k(f_i)\}$ then there can be no solution to the system.*

PROOF. For the former statement observe that f_i has no real positive roots, thus the system has no roots. The latter statement means that the reduced domains of each f_i, $i = 1, .., s$ do not intersect, thus there are no solutions.

We can use interval arithmetic to identify additional empty domains; if the sign of some initial f_i is constant in $I_H = \mathcal{H}(\mathbb{R}^n_{>0})$ then this domain is discarded. We can also simply inspect the coefficients of each $H(f_i)$; if there are no sign changes then the corresponding box contains no solution.

The accuracy of these criteria greatly affects the performance of the algorithm. In particular, the sooner an empty domain is rejected the less subdivisions will take place and the process will terminate faster. We justify that the exclusion criteria will eventually succeed on an empty domain by proving a generalization of Vincent's theorem to the tensor multivariate case.

Theorem 4.2. *Let $f(\underline{x}) = \sum_{\underline{i}=0}^{d} c_{\underline{i}} \, \underline{x}^{\underline{i}}$ be a polynomial with real coefficients, such that it has no (complex) solution with $\Re(z_k) \geq 0$ for $k = 1, .., n$. Then all its coefficients $c_{i_1,...,i_n}$ are of the same sign.*

PROOF. We prove the result by induction on n, the number of variables. For $n = 1$, this is the classical Vincent's theorem [1].

Consider now a polynomial

$$f(x_1, x_2) = \sum_{0 \leq i_1 \leq d_1, 0 \leq i_2 \leq d_2} c_{i_1, i_2} \, x_1^{i_1} \, x_2^{i_2}$$

in two variables with no (complex) solution such that $\Re(x_i) \geq 0$ for $i = 1, 2$. We prove the result for $n = 2$, by induction on the degree $d = d_1 + d_2$. The property is obvious for polynomials of degree $d = 0$. Let us assume it for polynomials of degree less than d.

By hypothesis, for any $z_1 \in \mathbb{C}$ with $\Re(z_1) \geq 0$, the univariate polynomial $f(z_1, x_2)$ has no root with $\Re(x_2) \geq 0$. According to Lucas theorem [15], the complex roots of $\partial_{x_2} f(z_1, x_2)$ are in the convex hull of the complex roots of $f(z_1, x_2)$. Thus, there is no root of $\partial_{x_2} f(x_1, x_2)$ with $\Re(x_1) \geq 0$ and $\Re(x_2) \geq 0$. By induction hypothesis, the coefficients of $\partial_{x_2} f(x_1, x_2)$ are of the same sign. We decompose P as

$$f(x_1, x_2) = f(x_1, 0) + f_1(x_1, x_2)$$

where $f_1(x_1, x_2) = \sum_{0 \leq i_1 \leq d_1, 1 \leq i_2 \leq d_2} c_{i_1, i_2} \, x_1^{i_1} \, x_2^{i_2}$ with c_{i_1, i_2} of the same sign, say positive. By Vincent theorem in one variable, as $f(x_1, 0)$ has no root with $\Re(x_1) \geq 0$, the coefficients $c_{i_1, 0}$ of $f(x_1, 0)$ are also of the same sign. If this sign is different from the sign of c_{i_1, i_2} for $i_2 \geqslant 1$ (ie. negative here), then $f(0, x_2)$ has one sign variation in its coefficients list. By Descartes rule, it has one real positive root, which contradicts the hypothesis on f. Thus, all the coefficients have the same sign.

Assume that the property has been proved for polynomials in $n - 1$ variables and let us consider a polynomial $f(\underline{x}) = \sum_{\underline{i}=0}^{d} c_{\underline{i}} \, \underline{x}^{\underline{i}}$ in n variables with no (complex) solution such that $\Re(x_k) \geq 0$ for $k = 1, .., n$. For any $z_1, .., z_{n-1} \in \mathbb{C}$ with $\Re(z_k) \geq 0$, for $k = 1, .., n-1$, the polynomial $f(z_1, .., z_{n-1}, x_n)$ and $\partial_{x_n} f(z_1, .., z_{n-1}, x_n)$ has no root with $\Re(x_n) \geq 0$. By Lucas theorem and induction hypothesis on the degree, $\partial_{x_n} f(\underline{x})$ has coefficients of the same sign. We also have $f(x_1, .., x_{n-1}, 0)$ with coefficients of the same sign, by induction hypothesis on the number of variables. If the two signs are different, then $f(0, .., 0, x_n)$ has one sign variation in its coefficients and thus one real positive root, say ζ_n, which cannot be the case, since $(0, .., 0, \zeta_n)$ would yield a real root of f. We deduce that all the coefficients of f are of the same sign.

This completes the induction proof of the theorem.

We can reformulate this result for bounded domains, using homography transformations, as follows.

Corollary 4.3. *Let $H(f) = \sum_{\underline{i}=0}^{d} c_{\underline{i}} \, \underline{x}^{\underline{i}}$ be the representation of f through $\mathcal{H}$ in a box $I_H = [\underline{u}, \underline{v}]$. If there is no root $\underline{z} \in \mathbb{C}^n$ of f such that*

$$\left| z_k - \frac{u_k + v_k}{2} \right| \leq \frac{v_k - u_k}{2}, \ for \ k = 1, .., n,$$

then all the coefficients $c_{i_1,...,i_n}$ are of the same sign.

That is, if $dist_\infty(\mathcal{Z}_{\mathbb{C}^n}(f), m) > \dfrac{\delta}{2}$, where m is the center of I_H of size δ, then I_H is excluded by sign conditions.

PROOF. The interval $[u_k, v_k]$ is transformed by $\mathcal{H}^{-1}$ into $[0, +\infty]$ and the disk $\left|z_k - \frac{u_k+v_k}{2}\right| \leq \frac{v_k-u_k}{2}$ is transformed into the half complex plane $\Re(z_k) \geq 0$. We deduce that $H(f)$ has no root with $\Re(z_k) \geq 0$, $k = 1, \ldots, n$. By Theorem 4.2, the coefficients of $H(f)$ are of the same sign.

We deduce that if a domain is far enough from the zero locus of some f_i then it will be excluded, hence redundant empty domains concentrate only in a neighborhood of $\underline{f} = \underline{0}$.

The regions which will be excluded during the subdivision algorithm can be quantitatively related to the regions where $\|f(x)\|_\infty$ is large, using the Lipschitz constant of f:

Definition 4.4. *For a system $f = (f_1, \ldots, f_s)$ of polynomials $f_i \in \mathbb{R}[\underline{x}]$ $(i = 1, \ldots, s)$ and a box $I = I_1 \times \cdots \times I_n \subset \mathbb{R}^n$, let*

$$\lambda_I(f) = \max\{\frac{\|f(x) - f(y)\|}{\|x - y\|}; x \neq y \in I_\mathbb{C}\}$$

be the Lipschitz constant of f on $I_\mathbb{C}$.

By definition, we have $\|f(x) - f(y)\| \leq \lambda_I(f)\|x - y\|$ for all $x, y \in I_\mathbb{C}$.

Lemma 4.5. *Suppose that $I_\mathbb{C} \subset S_{\mathbb{C}^n}(0, \rho)$ with $\rho \geq 1$ and $f \in \mathbb{R}[\underline{x}]$ is of degree d, then*

$$\lambda_I(f) \leq \sqrt{2}\, d\, \|f\|\rho^{d-1}.$$

PROOF. Let $x, y \in I_\mathbb{C} \subset S_{\mathbb{C}^n}(0, \rho)$. We consider $g(t) := \Re(f(x + t\,(y - x)) - f(x))$. By the intermediate value theorem, there exists $\tau \in [0, 1]$ such that

$$|\Re(f(y) - f(x))| = |g'(\tau)| \leq \|D_{y-x}(f)(x + \tau\,(y - x))|$$

The same results applies if we take the imaginery part. By [22, III, Prop. 1 p. 484] we have

$$|f(y) - f(x)| \leq \sqrt{2}\|D_{y-x}(f)(x + \tau\,(y - x))| \leq \sqrt{2}\|D_{y-x}(f)\|\rho^{d-1}$$

since $x + \tau(y - x) \in I_\mathbb{C} \subset S_{\mathbb{C}^n}(0, \rho)$. By [22, III, Lem. 2, p. 485] we deduce that

$$|f(y) - f(x)| \leq \sqrt{2}\, d\, \|f\|\rho^{d-1}\|y - x\|,$$

which concludes the proof of the Lemma.

Proposition 4.6. *Let H be a homography of $\mathbb{R}^n$ and $\varepsilon > 0$ is such that $|I| < 2\varepsilon$, where $I = I_H$. If $\|f(x)\|_\infty > \sqrt{n}\,\lambda_I(f)\,\varepsilon$ then the coefficients of at least one of the functions $H(f_i)$ are of constant sign.*

PROOF. Suppose that $\|f(x)\|_\infty = |f_{i_0}(x)|$ and let $z \in \mathbb{C}^n$ be the closest point to x such that $f_{i_0}(z) = 0$. We have

$$|f_{i_0}(x)| = |f_{i_0}(x) - f_{i_0}(z)| \leq \lambda_I(f)\,\|x - z\| \leq \lambda_I(f)\,\sqrt{n}\,\|x - z\|_\infty.$$

and thus there is no root of f_{i_0} in $I_\mathbb{C}(x, \varepsilon)$ for $\varepsilon < \frac{|f_{i_0}(x)|}{\sqrt{n}\,\lambda_I(f)}$. By Corollary 4.3, we deduce that the coefficients of $H(f_{i_0})$ are of constant sign.

To analyze more precisely the subdivision process, we are going to relate the number of boxes which are not removed with some integral geometry invariants [28]:

Definition 4.7. *The tubular neighborhood of size ε of f_i is the set*

$$\tau_\varepsilon(f_i) = \{x \in \mathbb{R}^n \ : \ \exists z \in \mathbb{C}^n,\, f_i(z) = 0,\, s.t.\ \|z - x\|_\infty < \varepsilon\}.$$

We bound the number of boxes that are not excluded at each level of the subdivision tree.

11

Lemma 4.8. *Assume that $I = I_1 \times \cdots \times I_n$ is bounded. There exists $N_f^*(I) \in \mathbb{N}^+$ such that the number of boxes of size $\varepsilon > 0$ kept by the algorithm is less than $N_f^*(I)$ and such that*

$$V(f, \varepsilon) := \text{volume}\left(\cap_{i=1}^s \tau_\varepsilon(f_i) \cap I\right) \leq N_f^*(I)\,\epsilon^n.$$

PROOF. Consider a subdivision of the domain I_0 into boxes of size $\varepsilon > 0$. We will bound the number $N_f^\varepsilon(I)$ of boxes in this subdivision that are not rejected by the algorithm. By Corollary 4.3 if a box is not rejected, then we have for all $i = 1, \ldots, s$ $\text{dist}_\infty(\mathcal{Z}_{\mathbb{C}^n}(f_i), m) < \frac{\varepsilon}{2}$, where m is the center of the box. Thus all the points of this box (at distance $< \frac{\varepsilon}{2}$ to m) are at distance $< \varepsilon$ to $\mathcal{Z}_{\mathbb{C}^n}(f_i)$ that is in $\cap_{i=1}^s \tau_\varepsilon(f_i) \cap I$.

To bound $N_f^\varepsilon(I)$, it suffices to estimate the n−dimensional volume $V(f, \varepsilon)$, since we have:

$$N_f^\varepsilon(I)\varepsilon^n \leq \text{volume}\left(\cap_{i=1}^s \tau_\varepsilon(f_i) \cap I\right) = V(f, \varepsilon).$$

When ε tends to 0, this volume becomes equivalent to a constant times ε^n. For a square system with simple roots in I, it becomes equivalent to the sum for all real roots ζ in I of the volumes of parallelotopes in n dimensions of height 2ε and edges proportional to the gradients of the polynomials at ζ; More precisely, it is bounded by $\varepsilon^n 2^n \sum_{\zeta \in I, f(\zeta) = 0} \frac{\prod_i \|\nabla f_i(\zeta)\|}{|J_f(\zeta)|}$. We deduce that there exists an integer $N_f^*(I) \geq 2^n \sum_{\zeta \in I, f(\zeta) = 0} \frac{\prod_i \|\nabla f_i(\zeta)\|}{|J_f(\zeta)|}$ such that $V(f, \varepsilon) \leq N_f^*(I)\,\varepsilon^n < \infty$. For overdetermined systems, the volume is bounded by a similar expression.

Since $V(f, \varepsilon)\,\varepsilon^{-n}$ has a limit when ε tends to 0, we deduce the existence of the finite constant $N_f^*(I) = \max_{\varepsilon > 0} N_f^\varepsilon(I)$, which concludes the proof of the lemma.

Notice that preconditioning operations can be used here to improve this bound.

4.2. Inclusion test

We consider two types of inclusion tests (which can easily be combined) and analyze the complexity of the corresponding subdivision process in the following sections.

4.3. Miranda's test

We present a first test that discovers common solutions, in a box, or equivalently in $\mathbb{R}_+^n$, through homography. To simplify the statements we assume that the system is square, i.e. $s = n$.

Definition 4.9. *The* lower face *polynomial of f with respect to direction k is $\text{low}(f, k) = f|_{x_k = 0}$. The* upper face *polynomial of f with respect to k is $\text{upp}(f, k) = f|_{x_k = \infty} := R_k(f)|_{x_k = 0}$.*

Lemma 4.10 (Miranda Theorem [26]). *If for some permutation $\pi : \{1, \ldots, n\} \to \{1, \ldots, n\}$, we have $\text{sign}(\text{low}(H(f_k), \pi(k)))$ and $\text{sign}(\text{upp}(H(f_k), \pi(k)))$ are constant and opposite for all $k = 1, \ldots, n$, then the equations $(f_1, \ldots, f_n)$ have at least one root in I_H.*

The implementation of Miranda's test can be done efficiently if we compute a $0 - 1$ matrix with (i, j)−th entry 1 iff $\text{sign}(\text{low}(H(f_i), j))$ and $\text{sign}(\text{upp}(H(f_i), j))$ are opposite. Then, Miranda test is satisfied iff there is no zero row and no zero column. To see this observe that the matrix is the sum of a permutation matrix and a $0 - 1$ matrix iff this permutation satisfies Miranda's test.

Combined with the following simple fact, we have a test that identifies boxes with a single root.

Lemma 4.11. *If $\det J_f(x)$ has constant sign in a box I, then there is at most one root of $f = (f_1, \ldots, f_n)$ in I.*

PROOF. Suppose $u, v \in I$ are two distinct roots; by the mean value theorem there is a point w on the line segment $\overline{uv}$, and thus in I, such that $J_f(w) \cdot (u - v) = f(u) - f(v) = \mathbf{0}$ hence $\det J_f(w) = 0$.

4.3.1. Complexity of the test

Miranda test can be decided with $\mathcal{O}(n^2)$ evaluations on interval (cf. [13]) as well as one evaluation of $J_{\underline{f}}$, overall $\mathcal{O}(n^2 d^n)$ operations. The cost of the inclusion test is dominated by the cost of evaluating $\mathcal{O}(n)$ polynomials of size $\mathcal{O}(d^n)$ on an interval, i.e. $\mathcal{O}(nd^n)$ operations suffice.

Proposition 4.12. *If the real roots of the square system in the initial domain I are simple, then Algorithm 1 stops with boxes isolating the real roots in I_0.*

PROOF. If the real roots of $f = (f_1, \ldots, f_n)$ in I_0 are simple, in a small neighborhood of them the Jacobian of f has a constant sign. By the inclusion test, any box included in this neighborhood will be output if and only if it contains a single root and has no real roots of the jacobian. Otherwise, it will be further subdivided or rejected. Suppose that the subdivision algorithm does not terminate. Then the size of the boxes kept at each step tends to zero. By Corollary 4.3, these boxes are in the intersection of the tubular neighborhoods $(\cap_{i=1}^s \tau_\varepsilon(f_i)) \cap \mathbb{R}^n$ for $\varepsilon > 0$ the maximal size of the kept boxes. If ε is small enough, these boxes are in a neighborhood of a root in which the Jacobian has a constant size, hence the inclusion test will succeed. By the exclusion criteria, a box domain is not subdivided indefinitely, but is eventually rejected when the coefficients become positive. Thus the algorithm either outputs isolating boxes that contains a real root of the system or rejects empty boxes. This shows, by contradiction, the termination of the subdivision algorithm.

4.4. α-inclusion test

In this section, we consider another inclusion test, based on α-theory and properties of convergence of Newton's method. This test of existence and unicity of a root in a neighborhood of a point involve the following constants:

Definition 4.13.

- $\beta_f(x) = \|Df(x)^{-1}f(x)\|$

- $\gamma_f(x) = \sup_{k \geq 2} \|\frac{1}{k!}Df(x)^{-1}D^k f(x)\|^{\frac{1}{k-1}}$

- $\alpha_f(x) = \beta_f(x)\,\gamma_f(x).$

Let $\delta(u) := \frac{1}{4}(1 + u - \sqrt{(1+u)^2 - 8\,u})$. We recall here a well-known theorem [22], [3] which is the foundation of the theory:

Theorem 4.14. *If $\alpha_f(x) < \alpha_0 \leq \frac{1}{4}(13 - 3\sqrt{17}) \sim 0.1577$ then f has a unique zero ζ in the ball $B(x, \frac{\delta_0}{\gamma_f(x)})$, with $\delta_0 = \delta(\alpha_0) \leq \frac{2-\sqrt{2}}{2} \sim 0.2929$. Moreover, for each point $z \in I(x, \frac{\delta_0}{\gamma_f(x)})$, Newton's method starting from z converges quadratically to ζ.*

This yields the following definition:

Definition 4.15. *A box $I(x, \varepsilon)$ is an α_0-inclusion box if*

- $\alpha_f(x) < \alpha_0$ *and*

- $\varepsilon < \sqrt{n}\frac{\delta(\alpha_0)}{\gamma_f(x)}.$

By theorem 4.14, if $I(x, \varepsilon)$ is an α_0-inclusion box then there is a unique root in the ball $B(x, \frac{\delta_0}{\gamma_f(x)})$ where $\delta_0 = \delta(\alpha_0)$. Moreover, we have $I(x, \varepsilon) \subset B(x, \frac{\delta_0}{\gamma_f(x)})$.

As we will see, the α_0-inclusion boxes that derive from the subdivision process (1) determine regions which isolate the roots of the system $f(\underline{x}) = 0$. In fact, if we are able to decompose a domain I_0 into a union of boxes which either do contain no roots or are α_0-inclusion boxes, then

- each root is in a connected component of the union of the α_0-inclusion boxes and

- each connected component of the union of the α_0-inclusion boxes contains a unique root of $f(x) = 0$.

More precisely, if such a connected component is $\cup_{i=1}^{s} I(x_i, \varepsilon_i)$ with $x_i \in D_0$ and $\varepsilon_i > 0$, then by theorem 4.14, there is a unique root ζ_i in $B(x_i, \frac{\delta_0}{\gamma_f(x_i)})$. As the balls $B(x_i, \frac{\delta_0}{\gamma_f(x_i)})$, $B(x_j, \frac{\delta_0}{\gamma_f(x_j)})$ of two adjacent α_0-inclusion boxes intersect, we must have

$$\zeta_i = \zeta_j \in B(x_i, \frac{\delta_0}{\gamma_f(x_i)}) \cap B(x_j, \frac{\delta_0}{\gamma_f(x_j)}).$$

Since the connected component $\cup_{i=1}^{s} I(x_i, \varepsilon_i)$ is a union of boxes in which every box has at least one adjacent, by recursive application of the above argument we derive that it contains a unique root ζ, which moreover is in $\cap_{i=1}^{s} B(x_i, \varepsilon_i)$.

5. The complexity of MCF

In this section we compute an upper bound on the complexity of the algorithm that exploits the continued fraction expansion of the real roots of the system. Hereafter, we call this algorithm MCF (Multivariate Continued Fractions). Since the analysis of the reduction steps of Sec. 3 and the Exclusion-Inclusion test of Sec. 4 would require much more developments, we simplify the situation and analyze a variant of this algorithm. We assume that two oracles are available. The first one computes, exactly, the partial quotients of the positive real roots of the system, that is the integer part of the coordinates. The second counts exactly the number of real roots of the system inside a hypercube in the open positive orthant, namely $\mathbb{R}_+^n$. Actually the latter suffices for the realization/implementation of the former. In what follows, we will assume the cost of the first oracle is bounded by $\mathcal{C}_1$, while the cost of the second is bounded by $\mathcal{C}_2$, and we shall derive the total complexity of the algorithm with respect to these parameters. In any case the number of reduction or subdivision steps that we derive is a lower bound on the number of steps that every variant of the algorithm will perform. The next section presents some preliminaries on continued fractions, and then we detail the complexity analysis.

5.1. About continued fractions

Our presentation follows closely [23], and we refer the reader to, e.g., [27, 4, 24] for additional details. A *simple (regular) continued fraction* is a (possibly infinite) expression of the form

$$c_0 + \cfrac{1}{c_1 + \cfrac{1}{c_2 + \cdots}} = [c_0, c_1, c_2, \ldots],$$

where the numbers c_i are called *partial quotients*, $c_i \in \mathbb{Z}$ and $c_i \geq 1$ for $i > 0$. Notice that c_0 may have any sign, however, in our real root isolation algorithm $c_0 \geq 0$, without loss of generality. By considering the recurrent relations

$$P_{-1} = 1, \quad P_0 = c_0, \quad P_{n+1} = c_{n+1} P_n + P_{n-1},$$
$$Q_{-1} = 0, \quad Q_0 = 1, \quad Q_{n+1} = c_{n+1} Q_n + Q_{n-1},$$

it can be shown by induction that $R_n = \frac{P_n}{Q_n} = [c_0, c_1, \ldots, c_n]$, for $n = 0, 1, 2, \ldots$.

If $\gamma = [c_0, c_1, \ldots]$ then $\gamma = c_0 + \frac{1}{Q_0 Q_1} - \frac{1}{Q_1 Q_2} + \cdots = c_0 + \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{Q_{n-1} Q_n}$ and since this is a series of decreasing alternating terms, it converges to some real number γ. A finite section $R_n = \frac{P_n}{Q_n} = [c_0, c_1, \ldots, c_n]$ is called the n−th *convergent* (or *approximant*) of γ and the tails $\gamma_{n+1} = [c_{n+1}, c_{n+2}, \ldots]$ are known as its *complete quotients*. That is $\gamma = [c_0, c_1, \ldots, c_n, \gamma_{n+1}]$ for $n = 0, 1, 2, \ldots$. There is an one to one correspondence between the real numbers and the continued fractions, where evidently the finite continued fractions correspond to rational numbers.

It is known that $Q_n \geq F_{n+1}$ and that $F_{n+1} < \phi^n < F_{n+2}$, where F_n is the $n-$th Fibonacci number and $\phi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. Continued fractions are the best rational approximation(for a given denominator size). This is as follows:

$$\frac{1}{Q_n(Q_{n+1} + Q_n)} \leq \left| \gamma - \frac{P_n}{Q_n} \right| \leq \frac{1}{Q_n Q_{n+1}} < \phi^{-2n+1}. \tag{7}$$

Let $\gamma = [c_0, c_1, \ldots]$ be the continued fraction expansion of a real number. The Gauss-Kuzmin distribution [4] states that for almost all real numbers γ, that is the set of exceptions has Lebesgue measure zero, the probability for a positive integer δ to appear as an element c_i in the continued fraction expansion of γ is

$$Prob[c_i = \delta] \simeq \lg \frac{(\delta + 1)^2}{\delta(\delta + 2)}, \quad \text{for any fixed } i > 0. \tag{8}$$

The Gauss-Kuzmin law induces that we can not bound the mean value of the partial quotients or in other words that the expected value (arithmetic mean) of the partial quotients is diverging, i.e.

$$E[c_i] = \sum_{\delta=1}^{\infty} \delta \, Prob[c_i = \delta] = \infty, \text{ for } i > 0.$$

However, the geometric, as well as the harmonic, mean is not only asymptotically bounded, but is bounded by a constant, for almost all $\gamma \in \mathbb{R}$. For the geometric mean this is the famous Khintchine's constant [14], i.e.

$$\lim_{n \to \infty} \sqrt[n]{\prod_{i=1}^{n} c_i} = \mathcal{K} = 2.685452001...$$

It is not known if $\mathcal{K}$ is a transcendental number. The expected value of the bit size of the partial quotients is a constant for almost all real numbers, when $n \to \infty$ or n sufficiently big [14]. Notice that in (8), $i > 0$, thus $\gamma \in \mathbb{R}$ is uniformly distributed in $(0, 1)$. Let $\mathcal{L}(c_i) \triangleq b_i$, then

$$E[b_i] = \mathcal{O}(\lg \mathcal{K}) = \mathcal{O}(1). \tag{9}$$

5.2. Complexity results

We denote by σ the upper bound on the bitsize of the partial quotients that appear during the execution of the algorithm.

Lemma 5.1. *The number of reduction and subdivision steps that the algorithm performs is $\widetilde{\mathcal{O}}(2^n \, n(d + n + \tau)d^{2n-1})$.*

PROOF. Let $\zeta = (\zeta_1, \ldots, \zeta_n)$ be a real root of the system. It suffices to consider the number of steps needed to isolate the i coordinate of ζ. We remind the reader that we are working in the positive orthant and we can compute exactly the next partial quotient in each coordinate; in other words a vector $\underline{l} = (l_1, \ldots, l_n)$, where each l_i, $1 \leq i \leq n$, is the partial quotient of a coordinate of a positive real[1] solution of the system.

Let $k_i(\zeta)$ be the number of steps needed to isolate the i^{th} coordinate of the real root ζ. The analysis is similar to the univariate case. We may consider the whole process of the subdivision algorithm as a 2^n-ary tree, where at each node we associate a, possible, open hypercube, and to the root of the tree we associate the positive orthant. The leaves of the tree form a partition of the positive orthant, and they contain at most one real root of the system. The number of nodes of the tree correspond to the number of steps if the algorithms. We prune some leaves of the tree to make the counting easier. We prune all the leaves that

---

[1]Actually the analysis holds even in the case where each l_i is the partial quotient of the positive imaginary part of a coordinate of a solution of the system.

have siblings that are not leaves. We prune all the leaves that do not contain a real root. Notice that these leaves have at least one sibling that contains a real root, since otherwise the subdivision process would have stopped one step before. All the remaining leaves contain a real root. If there are siblings that are all leaves then we keep arbitrarily one of them. The number of nodes in the original tree is at most 2^n times the number of nodes of the pruned tree.

Now all the leaves of the pruned tree correspond to a hypercube that contains exactly one real root, say ζ, of the system. The edges of the hypercube correspond to successive approximations of ζ_i by consecutive approximants. The $k_i(\zeta)$-th approximant is $\frac{P_{k_i(\zeta)}}{Q_{k_i(\zeta)}}$ and following (7) should satisfy

$$\left| \frac{P_{k_i(\zeta)}}{Q_{k_i(\zeta)}} - \zeta_i \right| \leq \frac{1}{Q_{k_i(\zeta)} Q_{k_i(\zeta)+1}} < \phi^{-2k_i(\zeta)+1}.$$

In order to isolate ζ_i, it suffices to have

$$\left| \frac{P_{k_i(\zeta)}}{Q_{k_i(\zeta)}} - \zeta_i \right| \leq \Delta_i(\zeta),$$

where $\Delta_i(\zeta)$ is the local separation bound of ζ_i, that is the smallest distance between ζ_i and all the other i-coordinates of the positive real solutions of the system. The number of nodes from the hypercube that isolates ζ to the root of the tree is, in the worst case, $k(\zeta) = \max_i k_i(\zeta)$.

Combining the last two equations, we deduce that to achieve the desired approximation, we should have $\phi^{-2k_i(\zeta)+1} \leq \Delta_i(\zeta)$, or $k_i(\zeta) \geq \frac{1}{2} - \frac{1}{2}\lg \Delta_i(\zeta)$. That is to achieve the desired approximation it suffices to compute $\mathcal{O}(-\frac{1}{2}\lg \Delta_i(\zeta))$ approximants. In other words, from the leaf that corresponds to a hypercube that isolates ζ to the root of the tree there are $\mathcal{O}(-\frac{1}{2}\lg \Delta(\zeta))$ nodes, where $\Delta = \min \Delta_i$.

To compute the total number of steps, i.e. the total number of nodes of the pruned tree, we need to sum over all the real roots that appear at the leaves of the tree; hence

$$\sum_{\zeta \in V} k(\zeta) \leq \frac{1}{2}R - \frac{1}{2}\sum_{\zeta \in V} \lg \Delta(\zeta) = \frac{1}{2}R - \frac{1}{2}\lg \prod_{\zeta \in V} \Delta(\zeta),$$

where $|V| \leq R$, V is the set of positive real roots at the leaves of the pruned tree and R and the total number of positive real roots.

To bound the logarithm of the product, we use $\mathtt{DMM}_n$ [12], i.e. aggregate separation bounds for multivariate, zero-dimensional polynomial systems. It holds

$$
\begin{array}{rcl}
\prod_{\zeta \in V} \Delta(\zeta) & \geq & 2^{-(3+4\lg n+4n\lg d)d^{2n}} \, 2^{-2n(1+n\lg d+\tau)d^{2n-1}} \\
-\log \prod_{\zeta \in V} \Delta(\zeta) & \leq & (3 + 4\lg n + 4n\lg d)d^{2n} + 2n(1 + n\lg d + \tau)d^{2n-1}, \\
-\log \prod_{\zeta \in V} \Delta(\zeta) & = & \widetilde{\mathcal{O}}(nd^{2n} + (n^2 + n\tau)d^{2n-1}).
\end{array}
$$

Taking into account that $R \leq d^n$ we conclude that the total number of nodes of the pruned tree is $\widetilde{\mathcal{O}}(nd^{2n} + n(n+\tau)d^{2n-1})$, and hence the number of steps of the algorithms is $\widetilde{\mathcal{O}}(2^n nd^{2n} + 2^n n(n+\tau)d^{2n-1})$.

Proposition 5.2. *The total complexity of the algorithm is $\widetilde{\mathcal{O}}_B(2^{3n}n^5(n^2+d^2+\tau^2)d^{5n-1}\sigma + (\mathcal{C}_1+\mathcal{C}_2)2^n n(d+n+\tau)d^{2n-1})$.*

PROOF. At each h-th step of algorithm, if there are more than one roots of the corresponding system in the positive orthant, (let the cost of estimating this is be $\mathcal{C}_2$), we compute the corresponding partial quotients $l_h = (l_{h,1}, \ldots, l_{h,n})$, where $\mathcal{L}(h_{h,i}) \leq \sigma_h$ (let the cost of this computation be $\mathcal{C}_1$). Then, for each polynomial of the system, f, we perform the shift operation $f(x_1+l_1, \ldots, x_n+l_n)$, and then we split the positive orthant to 2^n subdomains. Let us estimate the cost of the last two operations.

A shift operation on a polynomial of degree $\leq d$, by a number of bitsize σ, increases the bitsize of the polynomial by an additive factor $nd\sigma$. At the h step of the algorithm, the polynomials of the corresponding

system are of bitsize $\mathcal{O}(\tau + nd \sum_{i=1}^{h} \sigma_h)$, and we need to perform a shift operation to all the variables, with number of bitsize σ_{h+1}. The cost of this operation is $\widetilde{\mathcal{O}}_B(nd^n\tau + n^2d^{n+1}\sum_{k=1}^{h+1} \sigma_k)$, and since we have n polynomials the costs becomes $\widetilde{\mathcal{O}}_B(n^2d^n\tau + n^3d^{n+1}\sum_{k=1}^{h+1} \sigma_k)$, The resulting polynomial has bitsize $\mathcal{O}(\tau + nd\sum_{k=1}^{h+1} \sigma_k)$.

To compute the cost of splitting the domain, we proceed as follows. The cost is bounded by the cost of performing $n2^n$ operations $f(x_1 + 1, \ldots, x_n + 1)$, which in turn is $\widetilde{\mathcal{O}}_B(nd^n\tau + n^2d^{n+1}\sum_{k=1}^{h+1} \sigma_k + n^2d^{n+1})$. So the total cost becomes $\widetilde{\mathcal{O}}_B(2^nn^2d^n\tau + 2^nn^3d^{n+1}\sum_{k=1}^{h+1} \sigma_k)$. It remains to bound $\sum_{k=1}^{h+1} \sigma_k$. If σ is a bound on the bitsize of all the partial quotients that appear during and execution of the algorithm, then $\sum_{k=1}^{h+1} \sigma_k = \mathcal{O}(h\sigma)$.

Moreover, $h \leq \#(T) = \widetilde{\mathcal{O}}(2^n nd^{2n} + 2^n n(n + \tau)d^{2n-1})$ (lem. 5.1), and so the cost of each step is $\widetilde{\mathcal{O}}_B(2^{2n}n^4(n + d + \tau)d^{3n}\sigma)$.

Finally, multiplying by the number of steps (lem. 5.1) we get a bound of $\widetilde{\mathcal{O}}_B(2^{3n}n^5(n^2 + d^2 + \tau^2)d^{5n-1}\sigma)$.

To derive the total complexity we have to take into account that at each step we compute some partial quotients and and we count the number of real root of the system in the positive orthant. Hence the total complexity of the algorithm is $\widetilde{\mathcal{O}}_B(2^{3n}n^5(n^2 + d^2 + \tau^2)d^{5n-1}\sigma + (\mathcal{C}_1 + \mathcal{C}_2)2^n n(d + n + \tau)d^{2n-1})$.

In the univariate case ($n = 1$), if we assume that (9) holds for real algebraic numbers, then the cost of $\mathcal{C}_1$ and $\mathcal{C}_2$ is dominated by that of the other steps, that is the splitting operations, and the (average) complexity becomes $\widetilde{\mathcal{O}}_B(d^3\tau)$ and matches the one derived in [23] (without scaling).

5.3. Further improvements

We can reduce the number of steps that the algorithm performs, and thus improve the total complexity bound of the algorithm, using the same trick as in [23]. The main idea is that the continued fraction expansion of a real root of a polynomial does not depend on the initial computed interval that contains all the roots. Thus, we spread away the roots by scaling the variables of the polynomials of the system by a carefully chosen value.

If we apply the map $(x_1, \ldots, x_n) \mapsto (x_1/2^\ell, \ldots, x_n/2^\ell)$, to the initial polynomials of the system, then the real roots are multiplied by 2^ℓ, and thus their distance increases. The key observation is that the continued fraction expansion of the real roots does not depend on their integer part. Let ζ be the roots of the system, and γ, be the roots after the scaling. It holds $\gamma = 2^\ell \zeta$. From [12] it holds that

$$-\log \prod_{\zeta \in V} \Delta_i(\zeta) \leq (3 + 4\lg n + 4n\lg d)d^{2n} + 2n(1 + n\lg d + \tau)d^{2n-1},$$

and thus

$$-\log \prod_{\zeta \in V} \Delta_i(\gamma) = -\log 2^{R\ell} \prod_{\zeta \in V} \Delta_i(\zeta)$$
$$\leq (2n\tau d^{2n-1} + 2nd^n)\lg(nd^{2n}) - R\ell.$$

If we choose $\ell = \widetilde{\mathcal{O}}(nd^{n-1}(n+d+\tau))$ and assume that $R \leq d^n$ which is the worst case, then $-\log \prod_{\zeta \in V} \Delta_i(\gamma) = \widetilde{\mathcal{O}}(1)$. Thus, following the proof of Lemma 5.1, the number of steps that the algorithm performs is $\mathcal{O}(2^n d^n)$.

The bitsize of the scaled polynomials becomes $\widetilde{\mathcal{O}}(n^2d^{n+1} + n^2d^n\tau)$. The total complexity of algorithm is now

$$\widetilde{\mathcal{O}}_B(2^{2n}n^3d^{3n}(n + 2^nd\sigma + n\tau) + +2^nd^n(\mathcal{C}_1 + \mathcal{C}_2)),$$

where σ the maximum bitsize of the partial quotient appear during the execution of the algorithm. If we assume that (9) holds for real algebraic numbers, then $\sigma = \mathcal{O}(1)$. Notice that in this case, when $n = 1$, the bound becomes $\widetilde{\mathcal{O}}_B(d^3\tau)$, which agrees with the one proved in [23].

The discussion above combined with Proposition 5.2 lead us to:

Theorem 5.3. *The total complexity of the algorithm is* $\widetilde{\mathcal{O}}_B(2^{2n}n^3d^{3n}(n + 2^nd\sigma + n\tau) + 2^nd^n(\mathcal{C}_1 + \mathcal{C}_2))$.

17

6. Complexity and condition number

We are going to analyze the behavior of the subdivision algorithm in terms of condition numbers of the input polynomial system. This analysis uses the same tools as in [6]. However, we provide a bound which is not exponential but linear in the logarithm of this condition number. This bounds are closely connected to the results in [16] or [5]. They are formulated in terms of this condition number defined in [7], which takes into account the distance to systems with degenerate real roots.

In the following, we consider a system $f = (f_1, \ldots, f_n)$ of polynomials $f_i \in \mathbb{R}[\underline{x}]$ of degree $d_i := \deg(f_i)$. We denote by $d = \max\{d_1, \ldots, d_n\}$. We assume that the roots of the system in the box $I_0 \subset \mathbb{R}^n$ are simple. By a homographic transformation, we will assume that the initial domain is $I_0 := [-1, 1]^n$, so that $I_{0,\mathbb{C}} \subset B_{\mathbb{C}}(0, \rho)$ with $\rho = \sqrt{n}$. We assume that the system f has no multiple root in I_0. We consider a subdivision algorithm based on the exclusion test of Sec. 4.1 and the inclusion test of Sec. 4.4. We assume moreover that there is a constant $0 < \Phi < 1$ such that at each subdivision the size of a box which is kept is at most Φ times the size of its parent box. Consequently, if we apply k subdivision steps, the boxes which are kept are of size Φ^k times the size of the initial box.

We recall here the definitions that will be used in this complexity analysis.

Definition 6.1. *For a system $f = (f_1, \ldots, f_n)$ of polynomials $f_i \in \mathbb{R}[\underline{x}]$ with $\deg(f_i) = d_i$,*

$$\mu_f(x) := \|f\| \|Df(x)^{-1} \Delta(\sqrt{d_1} \|x\|_1, \ldots, \sqrt{d_n} \|x\|_1)\|,$$

where $\Delta(z_1, \ldots, z_n)$ is the diagonal matrix with entry z_i for the index (i, i) and 0 for the indices (i, j) with $1 \leq i \neq j \leq n$.

For a root $\zeta \in \mathbb{C}^n$, $\mu_f(\zeta)$ is the condition number of the system f at ζ. It measures the distance to the set of systems which are degenerated at ζ. See [3, p. 233]. This distance bounds the size of complex perturbation we can apply on our system, while staying in the safe region of systems with simple roots. However in practice, we usually consider real perturbations. To take into account the distance to real systems which are degenerate, we use the following *real condition number* [7]:

Definition 6.2. *The local condition number at $x \in \mathbb{R}^n$ for the system f is*

$$\kappa_f(x) := \frac{\|f\|}{(\|f\|^2 \mu_f(x)^{-2} + \|f(x)\|_\infty^2)^{\frac{1}{2}}} = \frac{1}{(\mu_f(x)^{-2} + \|f(x)\|_\infty^2 \|f\|^{-2})^{\frac{1}{2}}}.$$

This condition number $\kappa_f(x)$ is related to the distance to the set $\Sigma_\mathbb{R}(x, \underline{d})$ of systems of real polynomials $(f_1, \ldots, f_n)$ with $\underline{d} = (d_1, \ldots, d_n)$, $\deg(f_i) = d_i$ which are singular at x [7]:

$$\kappa_f(x) := \frac{1}{\text{dist}(f, \Sigma_\mathbb{R}(x, d))}.$$

Definition 6.3. $\kappa_I(f) := \max\{k(x, f); x \in I\}.$

In the univariate case $f \in \mathbb{R}[x]$, the value of κ_f will be large at the real roots $\zeta \in \mathbb{R}$ of f where $f'(\zeta)$ is small and at local extrema ξ where $|f(\xi)|$ is small.

Proposition 6.4. *For all $\sigma \geq \|f\|$ and $\varepsilon < \frac{4 \alpha_0 \|f\|^2}{d^{\frac{3}{2}} \kappa_f(x)^2 \sigma^2}$, we have*

- $\|f(x)\|_\infty > \sigma \varepsilon$, or

- $\alpha_f(x) < \alpha_0$.

18

PROOF. Let us suppose that $\alpha_f(x) \geq \alpha_0$ and prove that $\|f(x)\|_\infty > \sigma\varepsilon$. We consider two cases.

In the case where $\mu_f(x)^{-1} < \|f(x)\| \|f\|^{-1}$, we have $\kappa_f(x) > 2^{\frac{1}{2}}\mu_f(x)$. By [22, Proposition 2, p. 467], we have

$$\beta_f(x) \leq \|x\|_1 \mu_f(x) \frac{\|f(x)\|_\infty}{\|f\|},$$

where $\|x\|_1 = (1 + |x_1|^2 + \cdots + |x_n|^2)^{\frac{1}{2}}$. By [22, Proposition 3, p. 468],

$$\gamma_f(x) \leq \frac{1}{2\|x\|_1}\mu_f(x)d^{3/2}.$$

We deduce that

$$\frac{1}{4}\varepsilon\, d^{\frac{3}{2}}\kappa_f(x)^2\frac{\sigma^2}{\|f\|^2} < \alpha_0 \leq \alpha_f(x) = \beta_f(x)\,\gamma_f(x) \leq \frac{1}{2}\mu_f(x)^2 d^{\frac{3}{2}}\frac{\|f(x)\|_\infty}{\|f\|}$$

which implies that

$$\varepsilon\,\frac{\sigma^2}{\|f\|} < \|f(x)\|_\infty,$$

since $\kappa_f(x) > 2^{\frac{1}{2}}\mu_f(x)$. As $\sigma \geq \|f\|$, we deduce that

$$\varepsilon\,\sigma < \|f(x)\|_\infty.$$

In the case where $\mu_f(x)^{-1} \geq \frac{\|f(x)\|_\infty}{\|f\|}$, we have $\kappa_f(x) \geq 2^{\frac{1}{2}}\frac{\|f\|}{\|f(x)\|_\infty}$ and

$$2\,\varepsilon\frac{\|f\|^2}{\|f(x)\|_\infty^2}\frac{\sigma^2}{\|f\|^2} \leq 2\,\varepsilon\,\kappa_f(x)^2\frac{\sigma^2}{\|f\|^2} < 4\frac{\alpha_0}{d^{\frac{3}{2}}},$$

which implies that

$$\|f(x)\|_\infty > \sqrt{\frac{d^{\frac{3}{2}}}{2\,\alpha_0}}\sigma\varepsilon^{\frac{1}{2}} > \sigma\,\varepsilon,$$

since $\varepsilon < 1$ and $\frac{d^{\frac{3}{2}}}{2\,\alpha_0} > 1$.

Let $\alpha_0 = 0.1$ so that $\delta_0 = \delta(\alpha_0) \sim 0.1145$. We bound the complexity of the subdivision algorithm for the exclusion test of section 4.1 and the α_0-inclusion test of section 4.4.

Theorem 6.5. *The number of arithmetic operations needed to isolate the roots of a polynomial system f with simple roots in $I = I_1 \times \cdots \times I_n \subset \mathbb{R}^n$ with $I_\mathbb{C} \subset S_{\mathbb{C}^n}(0, \rho)$ and $\rho > 1$ is in*

$$\mathcal{O}(N_f^*(I)\, d^{n+1}\,(\log(\kappa_I(f)) + d\log(\rho) + \log(n)).$$

PROOF. By Lemma 4.8, the number of boxes of size ε kept during the subdivision is bounded by $N_f^*(I)$. The number of arithmetic operations is bounded by $N_f^*(I)$ times the cost of a subdivision step times the depth of the subdivision tree. The cost of a subdivision step is in $\mathcal{O}(d^{n+1})$.

We are going to bound the depth of the subdivision tree as follows. We will show that a box of size $\varepsilon < \frac{\alpha_0}{8\, n\, d^{\frac{7}{2}}\kappa_I(f)^2\rho^{2d-2}}$ either satisfies the exclusion test or is an α_0-inclusion box. By Proposition 6.4, for a box $I(x, \varepsilon)$ with $\varepsilon < \frac{\alpha_0}{8\, n\, d^{\frac{7}{2}}\kappa_I(f)^2\rho^{2d-2}} < \frac{4\,\alpha_0\|f\|^2}{d^{\frac{3}{2}}\kappa_I(f)^2(2\, n\, d^2\,\rho^{2d-2}\|f\|^2)}$, we have

- either $\|f(x)\| > \sqrt{2\, n}d\|f\|\rho^{d-1}\varepsilon$, which implies by Lemma 4.5 and Proposition 4.6 that the box $I(x, \varepsilon)$ satisfies the exclusion test;

- or $\alpha_f(x) < \alpha_0$ and by Theorem 4.14 there is unique root ζ of $f(\underline{x}) = 0$ in $B(x, \frac{\delta_0}{\gamma_f(x)})$.

19

Figure 3: Isolating boxes of the real roots of Left: (Σ_1), Right:(Σ_2).

To prove that in the latter case, the box $I(x, \varepsilon)$ is an α_0-inclusion box, we need to check that

$$\frac{\delta_0}{\gamma_f(x)} \geq n^{\frac{1}{2}}\varepsilon.$$

By [22, Proposition 2, p. 476], it is

$$\frac{1}{\gamma_f(x)} \geq \frac{\nu_0}{\gamma_f(\zeta)},$$

where $\nu_0 = (2\delta_0^2 - 4\delta_0 + 1)(1 - \delta_0) \sim 0.5016$. By [22, Proposition 3, p. 468], we have

$$\frac{1}{\gamma_f(\zeta)} \geq \frac{2}{d^{\frac{3}{2}}\mu_f(\zeta)} \geq \frac{2}{d^{\frac{3}{2}}\kappa_I(f)},$$

since $\mu_f(\zeta) = \kappa(\zeta, f) \leq \kappa_I(f)$. This implies that

$$\frac{\delta_0}{\gamma_f(x)} \geq \frac{2\delta_0\nu_0}{d^{\frac{3}{2}}\kappa_I(f)}.$$

As $\varepsilon < \frac{\alpha_0}{8\,nd^{\frac{7}{2}}\kappa_I(f)^2\rho^{2d}}$ and $\varepsilon < 1$, we have $\frac{1}{\kappa_I(f)} > \sqrt{\frac{8\,\rho^{2d}\,d^{\frac{7}{2}}\,n\,\varepsilon}{\alpha_0}} > \frac{2\sqrt{2}\rho^d\,d^{\frac{7}{4}}\,n^{\frac{1}{2}}}{\sqrt{\alpha_0}}\varepsilon$. We deduce that

$$\frac{\delta_0}{\gamma_f(x)} > \frac{4\sqrt{2}\,\rho^d\,\delta_0\nu_0}{\sqrt{\alpha_0}}d^{\frac{1}{4}}n^{\frac{1}{2}}\varepsilon > n^{\frac{1}{2}}\varepsilon,$$

since $\rho \geq 1$, $d \geq 1$ and $\frac{4\sqrt{2}\delta_0\nu_0}{\sqrt{\alpha_0}} > 1$. This proves that the box $I(x, \varepsilon)$ is an α_0-inclusion box. Therefore the subdivision step must stop before this precision, which gives a bound of order $\mathcal{O}(\log(\kappa_I(f) + d\log(\rho))$ for the depth of the subdivision tree.

7. Implementation and Examples

We have implemented the algorithm in the C++ library `realroot` of MATHEMAGIX[2], which is an open source effort that provides fundamental algebraic operations such as algebraic number manipulation tools, different types of univariate and multivariate polynomial root solvers, resultant and GCD computations, etc.

---

[2] http://www.mathemagix.org/

Figure 4: Isolating boxes of the real roots of the system Left: (Σ_3), Right: (Σ_4).

The polynomials are internally represented as a vector of coefficients along with some additional data, such as a variable dictionary and the degree of the polynomial in every variable. This allows us to map the tensor of coefficients to the one-dimensional memory. The univariate solver that is used is the continued fraction solver; this is essentially the same algorithm with a different inclusion criterion, namely Descartes rule. The same data structures is used to store the univariate polynomials, and the same shift/contraction routines. The univariate solver outputs a lower bound on the smallest positive root, as a result of a breadth-first traverse of the subdivision tree. Our code is templated and support different types of coefficients. It can use the integer arithmetic of GMP, since long integers appear as the box size decreases.

The following four examples demonstrate the output of our implementation, which we visualize using AXEL[3].

First, we consider the system $f_1 = f_2 = 0$ (Σ_1), where $f_1 = x^2 + y^2 - xy - 1$, and $f_2 = 10xy - 4$. We are looking for the real solutions in the domain $I = [-2,3] \times [-2,2]$, which is mapped to $\mathbb{R}^2_+$, by an initial transformation. The isolating boxes of the real roots can be seen in Fig. 3.

In systems (Σ_2), (Σ_3), We multiply f_1 and f_2 by quadratic components, hence we obtain

$$(\Sigma_2) \begin{cases} f_1 = x^4 + 2x^2y^2 - 2x^2 + y^4 - 2y^2 - x^3y - xy^3 + xy + 1 \\ f_2 = 20x^3y - 10x^2y^2 - 10xy^3 - 8x^2 + 4xy + 4y^2 \end{cases}$$

and

$$(\Sigma_3) \begin{cases} f_1 = 10x^2y - 10xy^3 - 4x + 4y^2 \\ f_2 = x^4 - 2x^2y - 2x^2 + y^2x^2 - 2y^3 - y^2 - x^3y + \\ \quad + 2xy^2 + xy + 2y + 1 \end{cases}$$

The isolating boxes of this system could be seen in Fig. 4. Notice, that size of the isolation boxes that are returned in this case is considerably smaller.

Consider the system (Σ_4), which consists of $f_1 = x^4 - 2x^2 - y^4 + 1$ and f_2, which is a polynomial of bidegree $(8,8)$. The output of the algorithm, that is the isolating boxes of the real roots can be seen in Fig. 4. One important observation is that the isolating boxes *are not* squares, which verifies the adaptive nature of the proposed algorithm.

We provide execution details on these experiments in Table 1. Several optimizations can be applied to our code, but the results already indicate that our approach competes well with the Bernstein case.

---

[3]http://axel.inria.fr

System	Domain	Iters.	Subdivs.	Sols.	Excluded
Σ_1	$[0,10]^2$	53	26	4	25
Σ_2	$[-2,3]^2$	263	131	12	126
Σ_3	$[-2,3]^2$	335	167	8	160
Σ_4	$[-3,3]^2$	1097	548	16	533

Table 1: Execution data for Σ_1, Σ_2, Σ_3, Σ_4.

Degrees in (x,y)	Domain	MCF(integer)	MCF(GMP)	Bernstein
$(2,1),(3,1)$	$[0,2] \times [0,2]$	20	110	2
$(4,4),(2,1)$	$[0,1] \times [0,1]$	70	280	30
$(6,6),(3,2)$	$[-2,2] \times [-2,2]$	10	200	10
$(4,3),(7,6)$	$[-5,5] \times [-5,5]$	110	600	20
$(8,8),(6,7)$	$[0,10] \times [0,10]$	110	540	20
$(8,8),(6,7)$	$[-2,2] \times [-2,2]$	960	8820	490
$(16,16),(12,15)$	$[0,10] \times [0,10]$	460	6550	320

Table 2: Running times in ms for our implementation (MCF) and the Bernstein solver [16].

We compared our implementation with the Bernstein solver of [16] on a number of bi-variate systems, and we present the times in milliseconds in Table 2. The tests were run on the same machine and the timings are rounded averages over 10 executions. When using machine integers for representing the polynomials, the Bernstein solver is proved faster from MCF, but the timings for both solvers are of the same order. If we use GMP integers then our algorithm is 10-20 times slower than the Bernstein solver; this difference is expected since GMP integers ought to be slower than machine numbers. Also, when using machine numbers, large coefficient values may occur if the degree as well as the predefined precision are high. For this, we declare the coefficients as doubles, in order to take advantage of their big range of available numbers. But then not all integers in this range are feasible, thus we work locally with a nearby system. However, setting the rounding mode to $-\infty$ guarantees that the lower univariate bounding functions we compute are indeed lower envelopes of the real system.

Acknowledgments

[1] A. Alesina and M. Galuzzi. New proof of Vincent's Theorem. *Enseignement Mathématique*, 44:219–256, 1998.

[2] M. Bartoň and B. Jüttler. Computing roots of polynomials by quadratic clipping. *Comp. Aided Geom. Design*, 24:125–141, 2007.

[3] L. Blum, F. Cucker, M. Shub, and S. Smale. Complexity and Real Computation, Springer, 1998.

[4] E. Bombieri and A. van der Poorten. Continued fractions of algebraic numbers. In *Computational algebra and number theory (Sydney, 1992)*, pp. 137–152. Kluwer Acad. Publ., Dordrecht, 1995.

[5] G. Chéze, J.-C. Yakoubsohn, A. Galligo, and B. Mourrain. Computing nearest gcd with certification. In *Symbolic-Numeric Computation (SNC'09)*, pages 29–34, Japon Kyoto, 2009-08-27. ACM New York, NY, USA.

[6] F. Cucker, T. Krick, G. Malajovich, and M. Wschebor. A numerical algorithm for zero counting, I: Complexity and accuracy. *J. Complex.*, 24(5-6):582–605, 2008.

[7] F. Cucker, T. Krick, G. Malajovich, and M. Wschebor. A numerical algorithm for zero counting. II: Distance to ill-posedness and smoothed analysis. *Journal of Fixed Point Theory and Applications*, 10.1007/s11784-009-0127-4, 2009.

[8] A. Eigenwillig, V. Sharma, and C. K. Yap. Almost tight recursion tree bounds for the Descartes method. In *ISSAC 2006*, pp. 71–78. ACM, New York, 2006.

[9] G. Elber and M.-S. Kim. Geometric constraint solver using multivariate rational spline functions. In *Proc. of 6th ACM Symposium on Solid Modelling and Applications*, pp. 1–10. ACM Press, 2001.

[10] I. Emiris, M. Hemmer, M. Karavelas, S. Limbach, B. Mourrain, E. P. Tsigaridas, and Z. Zafeirakopoulos. Cross-benchmarks of univariate algebraic kernels. ACS-TR-363602-02, INRIA, MPI and NUA, 2008.

[11] I. Z. Emiris, B. Mourrain, and E. P. Tsigaridas. Real Algebraic Numbers: Complexity Analysis and Experimentation. In P. Hertling, C. Hoffmann, W. Luther, and N. Revol, ed., *Reliable Implementations of Real Number Algorithms: Theory and Practice*, *LNCS* vol. 5045, pp. 57–82. Springer Verlag, 2008.

[12] I. Z. Emiris, B. Mourrain, and E. P. Tsigaridas. The DMM bound: multivariate (aggregate) separation bounds. Research Report, INRIA, March 2009, `http://hal.inria.fr/inria-00393833`.

[13] J. Garloff and A. P. Smith. Investigation of a subdivision based algorithm for solving systems of polynomial equations. *Journal of Nonlinear Analysis*, 47(1):167–178, 2001.

[14] A. Khintchine. *Continued Fractions*. University of Chicago Press, Chicago, 1964.

[15] M. Marden. *Geometry of Polynomials*. American Mathematical Society, Providence, RI, 1966.

[16] B. Mourrain and J. Pavone. Subdivision methods for solving polynomial equations. *J. Symbolic Comp.* 44(3):292 – 306, 2009.

[17] B. Mourrain, F. Rouillier, and M.-F. Roy. *Bernstein's basis and real root isolation*, pp. 459–478. MSRI Publications. Cambridge University Press, 2005.

[18] V. Pan. Solving a polynomial equation: Some history and recent progress. *SIAM Rev.*, 39(2):187–220, 1997.

[19] V. Pan. Univariate polynomials: Nearly optimal algorithms for numerical factorization and rootfinding. *JSC*, 33(5):701–733, 2002.

[20] V. Sharma. Complexity of real root isolation using continued fractions. *Theor. Comput. Sci.*, 409(2):292–310, 2008.

[21] E. C. Sherbrooke and N. M. Patrikalakis. Computation of the solutions of nonlinear polynomial systems. *Comput. Aided Geom. Design*, 10(5):379–405, 1993.

[22] M. Shub and S. Smale. Complexity of bezout's theorem I: Geometric aspects. *Journal of the American Mathematical Society*, 6(2):459–501, 1993.

[23] E. P. Tsigaridas and I. Z. Emiris. On the complexity of real root isolation using Continued Fractions. *Theoretical Computer Science*, 392:158–173, 2008.

[24] A. van der Poorten. An introduction to continued fractions. In *Diophantine analysis*, pp. 99–138. Cambridge University Press, 1986.

[25] J. von zur Gathen and J. Gerhard. Fast Algorithms for Taylor Shifts and Certain Difference Equations. In *Proc. Annual ACM ISSAC*, pp. 40–47, 1997.

[26] M. N. Vrahatis. A short proof and a generalization of Miranda's existence theorem. *Proceedings of the American Mathematical Society*, 107(3):701–703, 1989.

[27] C. Yap. *Fundamental Problems of Algorithmic Algebra*. Oxford University Press, New York, 2000.

[28] Y. Yomdin and G. Comte. *Tame geometry with applications in smooth analysis*. LNM 1834. Springer-Verlag, 2004.