Polynomial-time computing over quadratic maps I : sampling in real algebraic sets

Grigoriev, Dima.; Pasechnik, Dmitrii V.

2005

Grigoriev, D., & Pasechnik, D. V. (2005). Polynomial-time computing over quadratic maps I: sampling in real algebraic sets. *Computational Complexity*, 14, 20-52.

<https://hdl.handle.net/10356/92360>

<https://doi.org/10.1007/s00037-005-0189-7>

© 2005 Birkhäuser Verlag AG, Basel. This is the author created version of a work that has been peer reviewed and accepted for publication by *Computational Complexity*, Birkhäuser Verlag AG, Basel. It incorporates referee's comments but changes resulting from the publishing process, such as copyediting, structural formatting, may not be reflected in this document. The published version is available at the following DOI:
<http://dx.doi.org/10.1007/s00037-005-0189-7>.

Downloaded on 26 Apr 2023 21:07:12 SGT

POLYNOMIAL-TIME COMPUTING OVER QUADRATIC MAPS I: SAMPLING IN REAL ALGEBRAIC SETS.

DIMA GRIGORIEV AND DMITRII V. PASECHNIK

Abstract. Given a quadratic map $Q : \mathbb{K}^n \rightarrow \mathbb{K}^k$ defined over a computable subring D of a real closed field $\mathbb{K}$, and $p \in D[Y_1, \dots, Y_k]$ of degree d we consider the zero set $Z = Z(p(Q(X)), \mathbb{K}^n) \subseteq \mathbb{K}^n$ of $p(Q(X_1, \dots, X_n)) \in D[X_1, \dots, X_n]$. We present a procedure that computes, in $(dn)^{O(k)}$ arithmetic operations in D , a set $\mathcal{S}$ of (real univariate representations of) sampling points in $\mathbb{K}^n$ that intersects nontrivially each connected component of Z . As soon as $k = o(n)$, this is faster than the standard methods that all have exponential dependence on n in the complexity. In particular, our procedure is polynomial-time for constant k . In contrast, the best previously known procedure is only capable of deciding in $n^{O(k^2)}$ operations the nonemptiness (rather than constructing sampling points) of the set Z in the case of $p(Y) = \sum_i Y_i^2$ and homogeneous Q .

A by-product of our procedure is a bound $(dn)^{O(k)}$ on the number of connected components of Z .

The procedure consists of exact symbolic computations in D and outputs vectors of algebraic numbers. It involves extending $\mathbb{K}$ by infinitesimals and subsequent limit computation by a novel procedure that utilizes knowledge of an explicit isomorphism between real algebraic sets.

Keywords. symbolic computation, complexity, semialgebraic set, quadratic map, univariate representation, infinitesimal deformation.

Subject classification. 68W30, 13P10, 14Q20, 14Pxx

1. Introduction and the results

The algorithmic problem of finding points in real algebraic sets has received considerable attention, in particular as it forms a building block for a lot of procedures in real algebraic geometry [5]. Even if the algebraic sets one is interested in are subsets of $\mathbb{R}^n$, the algorithms with the best known complexity bounds use transcendental infinitesimals extending $\mathbb{R}$ to perform necessary geometric deformations of the sets. Hence it is natural to describe the procedures

as operating over an arbitrary real closed field $\mathbb{K}$, with the input data, i.e. the polynomials, lying in $D[X_1, \dots, X_n] = D[X]$, with $D \subset \mathbb{K}$ a computable (in sense discussed e.g. in [5, Sect. 8.1]) subring of $\mathbb{K}$. In the case $\mathbb{K} = \mathbb{R}$ one usually assumes $D = \mathbb{Z}$.

Let $\mathcal{S}_i$'s be the connected components of the real algebraic set $\mathcal{S} = Z(f, \mathbb{K}^n)$, $f \in D[X]$. In general the number of the $\mathcal{S}_i$'s is bounded by $d^{O(n)}$, where $d = \deg f$, and this bound is sharp, see [5, Theorem 7.23, Remark 7.22] and Remark 1.7 below. We are interested in point-finding (also called *sampling*) algorithms that produce a finite set of points that intersects each $\mathcal{S}_i$. Such algorithms with the best known complexity bounds need at most $d^{O(n)}$ arithmetic operations in D . Often such sets are exactly what is needed in applications. We describe here a procedure that finds a point in each connected component of our class of algebraic sets, namely the sets of the form $Z(p(Q(X)), \mathbb{K}^n)$, for $p \in D[Y_1, \dots, Y_k]$ of degree $\deg p = d$ and $Q = (Q_1(X), \dots, Q_k(X)) \in D[X]^k$ a quadratic map, i.e. $\deg Q_j \leq 2$ for $1 \leq j \leq k$ with the complexity $(dn)^{O(k)}$.

The result of [3], that bounded, in particular, the sum of the Betti numbers of the set of real solutions of a system of quadratic equations $Q_1(X) = \dots = Q_k(X) = 0$ (that can obviously be written as $Z(p(Q(X)), \mathbb{R}^n)$ with $p(Y) = \sum_j Y_j^2$) by a polynomial in k and n of degree $O(k)$ was perhaps the earliest indication that in the case $Z(p(Q(X)), \mathbb{K}^n)$ the number of connected components has only polynomial dependence on n . However, until the present work, an algorithmic procedure with the similar complexity bound $(dn)^{O(k)}$ for finding points in $Z(p(Q(X)), \mathbb{K}^n)$ was unknown. Even procedures that decide non-emptiness of $Z(p(Q(X)), \mathbb{K}^n)$ in time polynomial in n and d for fixed k were, for general p , unknown; in [2] such a procedure was described for $p(Y) = \sum_j Y_j^2$, homogeneous Q_i 's and $\mathbb{K} = \mathbb{R}$.

The technique we use is that of symbolic computation. All the data is represented exactly, as (real) algebraic numbers, if necessary. More precisely, elements of $\mathbb{K}^n$ that we compute with are given by real univariate representations. The latter are defined as follows. A *sign condition* for a set of polynomials $\mathcal{P} = \{P_1, \dots, P_s\} \subset \mathbb{K}[Y]$ is specified by $\sigma \in \{-1, 0, 1\}^s$ so that $\sigma = (\text{sign } P_1(Y), \dots, \text{sign } P_s(Y))$. *Thom encoding* [5, Lemma 2.38, Sect. 10.4] of a root $\alpha \in \mathbb{K}$ of $f \in \mathbb{K}[T]$ is a sign condition σ_α on the derivatives of f , that is $\sigma_\alpha = (\text{sign } f'(\alpha), \dots, \text{sign } f^{(\deg f - 1)}(\alpha))$. Note that σ_α and f determine $\alpha \in \mathbb{K}$. Let $\overline{\mathbb{K}}$ denote the algebraic closure of $\mathbb{K}$. A *univariate representation* of $u \in \overline{\mathbb{K}}^m$ is an $(m + 2)$ -tuple

$$(1.1) \quad u(T) = (f, g_0, g_1, \dots, g_m)$$

of univariate polynomials in $D[T]$ satisfying $u = \frac{1}{g_0(\alpha)}(g_1(\alpha), \dots, g_m(\alpha))$ for a

root $\alpha \in \overline{\mathbb{K}}$ of f , and such that f and g_0 are coprime. Obviously, each $u(T)$ can represent as many as $\deg(f)$ distinct elements of $\overline{\mathbb{K}}^m$. A *real* univariate representation of $u \in \mathbb{K}^m$ is a pair $u(T), \sigma_\alpha$, where σ_α is the Thom encoding of a root $\alpha \in \mathbb{K}$ of f .

The main result of the paper is as follows.

THEOREM 1.2. *Let $Q = (Q_1, \dots, Q_k) \in D[X_1, \dots, X_n]^k$ be a quadratic map $X \mapsto Q(X)$, and let $p \in D[Y_1, \dots, Y_k]$ satisfy $\deg p \leq d$. A set of real univariate representations $u(T), \sigma_\alpha$ of a set of points in $Z = Z(p(Q(X)), \mathbb{K}^n)$ meeting each connected component of Z can be computed in $(dn)^{O(k)}$ arithmetic operations in D . The degrees of polynomials in $u(T)$ are bounded by $(dn)^{O(k)}$. When $D = \mathbb{Z}$, the coefficients of $u(T)$ and the intermediate polynomial data will be bounded by $(dn)^{O(k)}$ times the bitsize of the input data p, Q .*

From now on whenever we talk about finding points in $\mathbb{K}^n$, they are meant to be given as real univariate representations.

REMARK 1.3. *When $\mathbb{K}$ is archimedean, e.g. $\mathbb{K} = \mathbb{R}$, the approximations, in the ring of fractions of D , of the point in $\mathbb{K}^n$ given by a real univariate representation can be found efficiently as long as approximations $\tilde{\alpha}$ of α can be computed efficiently (indeed, then one can just compute $u(\tilde{\alpha})$). For instance when $D = \mathbb{Z}$ one can find an interval $\mathcal{I} = [\alpha_-, \alpha_+] \ni \alpha$ with $\alpha_\pm \in \mathbb{Q}$ so that α is the only root of f in $\mathcal{I}$, see e.g. [5, Sect. 10.2]. Once $\mathcal{I}$ is known, one can compute its repeated (rational) bisections to obtain approximations of α of needed precision; the complexity of the latter is analyzed e.g. in [11] (see also [17]).*

Note that by *connected* (component of) semialgebraic set, we mean *semialgebraically connected*, that is, connected in the semialgebraic topology, (component of) semialgebraic set, see e.g. [8]. It is well-known that for the semialgebraic sets over $\mathbb{R}$ semialgebraic connectedness implies connectedness (in the usual Euclidean topology), see e.g. [5, Thm. 5.21].

Theorem 1.2 is proved in Section 5 by exhibiting a procedure that does the claimed task. It immediately implies the following.

COROLLARY 1.4. *The number of connected components of the set Z is at most $(dn)^{O(k)}$. $\square$*

An extra argument, to be published elsewhere, allowed us to show that the latter bound holds for the sum of Betti numbers of Z , and not only for the 0-th one, i.e. the number of components. As well, one can modify the procedure of Theorem 1.2 to prove

THEOREM 1.5. *Under the assumptions of Theorem 1.2, computing the exact minimum and a minimizer (i.e. a point where the minimum is attained) of $r(Q(X))$, for $r \in D[Y]$, $\deg r \leq d$, on $Z(p(Q(X)), \mathbb{K}^n)$, or checking that the minimum is not attained and computing the infimum, can be done within the same number of operations, and for $D = \mathbb{Z}$ within the same bitwise complexity, as the computation of Theorem 1.2.*

A proof of the latter, and a number of applications in mathematical programming, will appear in the continuation of the present paper.

An easier than optimization problem is the problem of checking whether the set $Z(p(Q(X)), \mathbb{K}^n)$ is empty, i.e. the feasibility problem. Our immediate predecessor here is [2], where it was shown that for homogeneous Q the emptiness of $Z(Q_1(X)^2 + \dots + Q_k(X)^2, \mathbb{R}^n - \{0\})$ can be checked in $n^{O(k^2)}$ operations in D .

For the sake of completeness, we state the following straightforward implication of Theorem 1.2.

COROLLARY 1.6. *The emptiness of $Z(p(Q(X)), \mathbb{K}^n)$ can be checked within the same complexity bound as in Theorem 1.2. $\square$*

REMARK 1.7. *It is easy to see that the bounds of Theorem 1.2 are close to best possible. Indeed, any real solution of degree 4 equation*

$$(X_1^2 - 1)^2 + \dots + (X_n^2 - 1)^2 = 0,$$

or the system of n quadratic equations

$$X_1^2 = X_2^2 = \dots = X_n^2 = 1$$

has coordinates 1 or -1 , and there are in total 2^n of them. In the continuation of the present paper we will further sharpen this by showing a similar result for a system of one cubic and two quadratic equations.

In a nutshell, the procedure at the core of Theorem 1.2 that we are going to describe works as follows. First, we write down the equations for the critical points of the projection map $X \mapsto X_1$ on $Z = Z(p(Q(X)), \mathbb{K}^n)$ by equating the gradient of $p(Q(X))$ with the vector proportional to the gradient of $X \mapsto X_1$, that is with a vector of the form $(\lambda, 0, \dots, 0)$. These equations have a rather special structure: the variables X occur either within $Q(X)$, or linearly. By introducing k new variables $(Y_1, \dots, Y_k) = Y = Q(X)$, we thus obtain a system of *linear* equations $A(Y)X = b(Y)$ in X . The next step is to solve this system;

we simply loop through all the maximal (by inclusion) candidates for invertible submatrices A_{UW} of $A(Y)$ and the corresponding partition $X_W \cup X_{\overline{W}}$ of X into X_W and the remaining variables $X_{\overline{W}}$. For each of them we rewrite the system to express X_W as rational functions $X_W = A'(Y, X_{\overline{W}})$ of Y and $X_{\overline{W}}$. This certainly only makes sense, from the complexity point of view, when $\text{rk}(A(Y))$ never drops below certain threshold. We make sure by means of an infinitesimal deformation that $\text{rk}(A(Y)) \geq n - k$. Then $|\overline{W}| \leq k$ and the coordinates of X are expressed as rational functions in at most $2k$ variables Y and $X_{\overline{W}}$. We are able to describe an isomorphism of a semialgebraic subset, that we call, following [2], *piece*, of the critical points of $X \mapsto X_1$ on Z , that corresponds to a particular A_{UW} being maximal and invertible, to a semialgebraic subset of $\mathbb{F}^{k+|\overline{W}|}$, for $\mathbb{F}$ being a real closed extension of $\mathbb{K}$, that is defined by polynomials of degree $O(nd)$. These pieces cover the whole set of the critical points just mentioned.

Finally, we find representatives of connected components of the pieces over $\mathbb{F}$, obtaining Y and $X_{\overline{W}}$ with values in $\mathbb{F}$, and recover X_W and $X_{\overline{W}}$ in the original field $\mathbb{K}$ by computing the limit.

The actual implementation of this procedure is more involved. Section 4 describes in detail the candidates for invertible submatrices A_{UW} of A and presents the explicit isomorphisms of pieces to semialgebraic subsets in $\mathbb{F}^{k+|\overline{W}|}$ mentioned above. In order to apply the result of Section 4 to Z , one needs to deform p in such a way that 0 becomes a regular value of $p(Q(X))$ and of $p(Y)$. Further, one needs to deform Q so that the number of pieces of the set of critical points of $X \mapsto X_1$ on Z does not exceed $(dn)^{O(k)}$. In fact, our deformation will give us a better bound, $n^{O(k)}$, on the latter. Lastly, one has to ensure (again, using a deformation) that Z is bounded, otherwise we miss connected components of Z whose projection on X_1 is open.

Our deformations are done by extending $\mathbb{K}$ with a number of infinitesimals. Subsequent limit computations are needed to recover elements in the original set Z by using the following Theorem 1.10. To state it, let us recall some notation. For a field $\mathbb{F}$ and a transcendental ζ , we denote by $\mathbb{F}\langle\zeta\rangle \subset \mathbb{F}((\zeta^{\frac{1}{\infty}}))$ the subfield of Puiseux series algebraic over $\mathbb{F}(\zeta)$. For

$$(1.8) \quad a = \sum_{i \geq \nu} a_i \zeta^{i/q} \in \mathbb{F}((\zeta^{\frac{1}{\infty}})), \quad 0 < q \in \mathbb{Z}$$

with the order $\nu/q \geq 0$, $a_\nu \neq 0$, define the *standard part* (cf. e.g. [10]) of a to be a_0 ; in [5] it is called the *limit* $a_0 = \lim_{\zeta} a$. Note that if $\nu < 0$ then $\lim_{\zeta} a$ is not defined. When ζ is a vector of infinitesimals $\zeta_1 \gg \zeta_2 \gg \dots \gg \zeta_\ell$, the notation $\lim_{\zeta} a$ is a shorthand for $\lim_{\zeta_1}(\lim_{\zeta_2}(\dots(\lim_{\zeta_\ell} a)\dots))$. It is often helpful to

view ζ as a *parameter* and computing $\lim_{\zeta} a$ as computing $\lim_{\zeta \rightarrow 0} a$, where $\lim$ is understood in the usual sense. Note that $\lim_{\zeta}$ is a ring homomorphism of the ring $\mathbb{F}\langle\zeta\rangle_b = \{a \in \mathbb{F}\langle\zeta\rangle \mid \nu(a) \geq 0\}$, of all the elements of $\mathbb{F}\langle\zeta\rangle$ *bounded* over $\mathbb{F}$, to $\mathbb{F}$.

Let $\mathbb{F}\langle\varepsilon\rangle$ be a real closed extension of a real closed $\mathbb{F}$ with infinitesimals $\varepsilon = (\varepsilon_1, \dots, \varepsilon_\ell)$ such that $\varepsilon_1 \gg \varepsilon_2 \gg \dots \gg \varepsilon_\ell$, and let $D \subset \mathbb{F}$ be a computable subring of $\mathbb{F}$. For $F \in D[\varepsilon][Y_1, \dots, Y_{q-1}]$, let $Z_F = Z(F(Y), \mathbb{F}\langle\varepsilon\rangle^{q-1})$, and let

$$(1.9) \quad \begin{aligned} &\Psi : \mathbb{F}\langle\varepsilon\rangle^{q-1} \rightarrow \mathbb{F}\langle\varepsilon\rangle^m \quad \text{be a rational mapping} \\ &Y \mapsto \left(\frac{\Omega_1(Y)}{\Lambda(Y)}, \dots, \frac{\Omega_m(Y)}{\Lambda(Y)} \right), \quad \Omega_i, \Lambda \in D[\varepsilon][Y], \quad 1 \leq i \leq m. \end{aligned}$$

THEOREM 1.10. *Let F and Ψ be as above, with the Y -degree of F at most d and the Y -degrees of Ω_i and Λ less than $d-1$, and their ε -degrees at most d . A set of univariate representations $u(T)$ of a set of points meeting each connected component of $\lim_{\varepsilon} \Psi(Z_F) \subseteq \mathbb{F}^m$ can be computed in $(m+d)^{O(q\ell)}$ arithmetic operations in D .*

The degrees of the polynomials in $u(T)$ are at most $d^{O(q)}$. When $D = \mathbb{Z}$, the bitsizes of the coefficients of $u(T)$ and of the intermediate data are bounded by a polynomial in d, m and $d^{O(q\ell)}$ times the bitsize of the input data.

Theorem 1.10 generalizes [5, Alg. 11.61] to non-identity mappings Ψ .

The remainder of the paper begins with presenting the procedures behind Theorem 1.10, along with its proof, in Sections 2 and 3. Then Section 4 presents the aforementioned decomposition of the zero set of $p(Q(X))$ under the regularity conditions. Finally, Section 5 describes the deformations of $p(Q(X))$ that are needed and completes the proof of the main Theorem 1.2.

2. Limits of solution images: dimension 0

As the first part of the proof of Theorem 1.10, in this section we address the problem of finding limits of the images $P(x)$ of the real roots $x \in Z(\mathcal{B}, \mathbb{F}\langle\varepsilon\rangle^q)$ of a 0-dimensional polynomial system $\mathcal{B} \subset \mathbb{F}\langle\varepsilon\rangle^q$ with respect to $\varepsilon = (\varepsilon_1, \dots, \varepsilon_\ell) \rightarrow 0$ under a polynomial mapping $P : \mathbb{F}\langle\varepsilon\rangle^q \rightarrow \mathbb{F}\langle\varepsilon\rangle^m$.

The separating element based methods for finding limits of the real roots of $\mathcal{B}$, such as [16], [5, Chapter 11], cannot deal directly with this situation, even in the simplest case of P being the orthogonal projection onto a subset of coordinates.

In this section we generalize these methods to accommodate our needs. We introduce *P -separating elements* $a \in \overline{\mathbb{F}}\langle\varepsilon\rangle[S_1, \dots, S_m]$ such that the map

$P(x) \mapsto a(P(x))$ is injective on $P(\overline{Z})$, where $\overline{Z} = Z(\mathcal{B}, \overline{\mathbb{F}}\langle\varepsilon\rangle^q)$, that is to say that $a(P(y)) \neq a(P(x))$ whenever $P(y) \neq P(x)$ for $x, y \in \overline{Z}$. We introduce below, in Section 2.2, the corresponding notion for the limit setting, *well- P -separating elements*.

It turns out that the machinery of [16], see also [1], generalizes here relatively smoothly. Let an ideal $(\mathcal{W}) \subseteq \overline{\mathbb{F}}\langle\varepsilon\rangle[S_1, \dots, S_n] = \overline{\mathbb{F}}\langle\varepsilon\rangle[S]$ be generated by its Gröbner basis $\mathcal{W} \subseteq D[\varepsilon][S]$. That is, we fixed a particular monomial ordering on $\overline{\mathbb{F}}\langle\varepsilon\rangle[S]$, and the leading (with respect to this ordering) terms of $\mathcal{W}$ generate the ideal of leading terms of $(\mathcal{W})$. Following [5, Sect. 11.3], we call $\mathcal{W}$ *parametrized special* if it is of the form $\mathcal{W} = \{b_1 S_1^{d_1} + U_1, \dots, b_n S_n^{d_n} + U_n\}$, where the leading terms are $b_1 S_1^{d_1}, \dots, b_n S_n^{d_n}$, and $\deg(U_i) < d_i$, $\deg_{S_j}(U_i) < d_j$ for $1 \leq i, j \leq n$. Note that the quotient algebra $\overline{\mathbb{F}}\langle\varepsilon\rangle[S]/(\mathcal{W})$ has the natural basis $U(\mathcal{W})$ of *monomials under the staircase*, that is, of monomials $S^\alpha = S_1^{\alpha_1} \dots S_n^{\alpha_n}$ with $\alpha_i < d_i$ for $1 \leq i \leq n$. In particular, the dimension of the quotient algebra is $d_1 \dots d_n$. In order to keep doing the arithmetic in $D[\varepsilon]$ when reducing with respect to $\mathcal{W}$ (and this is one of the purposes of a parametrized special Gröbner basis), one works in the basis $\overline{U}(\mathcal{W}) = \{b^{|\alpha|} S^\alpha \mid S^\alpha \in U(\mathcal{W})\}$, where $|\alpha| = \alpha_1 + \dots + \alpha_n$ and b a common multiple of $b_1, \dots, b_n \in D[\varepsilon]$.

The basis $\mathcal{W}$ naturally appears when critical points of a coordinate projection of a certain special type of hypersurface are computed, as in [5, Sect. 11.6], in contrast to a common situation when calculation of a Gröbner basis of an ideal is computationally very costly (the latter can generally require doubly exponential, in the number of variables, running time, cf. [14]). Given $\mathcal{W}$, one can efficiently compute the multiplication table of the quotient algebra, see [5, Alg. 11.22]. Namely, when the degrees of the elements of $\mathcal{W}$ in S (resp. in ε) are bounded by d (resp. by λ) it takes $(d\lambda)^{O(n\ell)}$ operations in D , the ε -degrees never exceed $\lambda(nd)^{O(1)}$; when $D = \mathbb{Z}$, the bitsize of the data involved is bounded by a polynomial in n and $(\lambda d)^{n\ell}$ times the bitsizes of the elements of $\mathcal{W}$, cf. [5, pp. 381–382].

THEOREM 2.1. *For D a computable subring of $\mathbb{F}$, let $\mathcal{B} \subset D[\varepsilon][X_1, \dots, X_q]$ define a 0-dimensional polynomial system, that is its own special Gröbner basis with the LCM of the leading terms equal to $b_{\mathcal{B}} \in D[\varepsilon]$. Let*

$$N = \dim \overline{\mathbb{F}}\langle\varepsilon\rangle[X]/(\mathcal{B}) \quad \text{and} \quad \overline{Z} = Z(\mathcal{B}, \overline{\mathbb{F}}\langle\varepsilon\rangle^q).$$

Assume the degrees of the elements of $\mathcal{B}$ in X as well as in ε , and the degree of $b_{\mathcal{B}}$ in ε , bounded by d . Let P be a polynomial mapping $X \mapsto (P_1(X), \dots, P_m(X))$, with $P_i \in D[\varepsilon][X]$ of degree $< d$ in X and at most d in ε .

Then a set of at most $(m-1)N^3$ candidates for univariate representations

$u(T) \in D^{m+2}$ for the elements of $\lim_{\varepsilon} P(\overline{\mathbb{Z}})$ can be computed in $(m + N)^{O(\ell)}$ arithmetic operations in D . The degrees of the polynomials in $u(T)$ are at most N .

When $D = \mathbb{Z}$, the bitsizes of the coefficients of $u(T)$ and of the intermediate data are bounded by a polynomial in d , m and N^ℓ times the bitsize of the input data $\mathcal{B}$.

Note that $\mathcal{B}$ will be constructed in Theorem 3.1.

REMARK 2.2. The number $(m - 1)N^3$ of candidates can be reduced to N , see Remark 2.29 later in this section.

The remainder of the section is devoted to the proof of Theorem 2.1.

2.1. On P -separating elements. Here we prepare the ground for the limit computations. Denote by $\chi(a, T)$ the characteristic polynomial of a linear transformation $a \in A$. The Stickelberger's lemma [5, Thm. 4.69] states in particular that

$$(2.3) \quad \chi(a, T) = \prod_{x \in \overline{\mathbb{Z}}} (T - a(x))^{\mu(x)}.$$

where $\mu(x)$ is the multiplicity of x as a root of $\mathcal{B}$. For a given $a \in A$, denote by $[x] \subseteq \overline{\mathbb{Z}}$ the equivalence class of x with respect to the equivalence relation defined by a , so that x is equivalent to y when $a(x) = a(y)$. Then

$$(2.4) \quad \chi(a, T) = \prod_{[x] \subseteq \overline{\mathbb{Z}}} (T - a(x))^{\mu_{[x]}}, \quad \text{where} \quad \mu_{[x]} = \sum_{y \in [x]} \mu(y).$$

Let $0 \neq b \in A$ be an a -class function on the relation $[*]$ induced by a , that is, $b(y) = b(x)$ for any $y \in [x]$. Let S be a variable and consider polynomials

$$(2.5) \quad \chi(a + Sb, T) = \prod_{[x] \subseteq \overline{\mathbb{Z}}} (T - a(x) - Sb(x))^{\mu_{[x]}},$$

$$(2.6) \quad g(a, b, T) = \frac{\partial \chi(a + Sb, T)}{\partial S} \Big|_{S=0}.$$

Then

$$(2.7) \quad g(a, b, T) = - \sum_{[x] \subseteq \overline{\mathbb{Z}}} b(x) \mu_{[x]} (T - a(x))^{\mu_{[x]}-1} \prod_{[x] \neq [y] \subseteq \overline{\mathbb{Z}}} (T - a(y))^{\mu_{[y]}}.$$

Observe that

$$g^{(\mu_{[x]}-1)}(a, b, T) = -b(x)(\mu_{[x]})! \prod_{[x] \neq [y] \subseteq \overline{\mathbb{Z}}} (T - a(y))^{\mu_{[y]}} + (T - a(x))h(T).$$

for some polynomial $h(T)$. Therefore

$$(2.8) \quad g^{(\mu_{[x]}-1)}(a, b, a(x)) = -b(x)(\mu_{[x]})! \prod_{[x] \neq [y] \subseteq \overline{\mathbb{Z}}} (a(x) - a(y))^{\mu_{[y]}}.$$

In particular the following holds.

LEMMA 2.9. *For a P -separating a , any $0 \neq r \in \mathbb{F}\langle \varepsilon \rangle$ and $x \in \overline{\mathbb{Z}}$,*

$$(2.10) \quad P_i(x) = \frac{g^{(\mu_{[x]}-1)}(a, rP_i, a(x))}{g^{(\mu_{[x]}-1)}(a, r, a(x))} \quad \text{for each } 1 \leq i \leq m.$$

PROOF. As a is P -separating, $g^{(\mu_{[x]}-1)}(a, r, a(x))$ never vanishes, and $[x] \neq [y]$ as soon as $P_i(x) \neq P_i(y)$. Thus $P_i(X)$ is a a -class function, and (2.8) holds for $b = P_i$. Now (2.8) implies the statement of the lemma, as the terms $r(\mu_{[x]})! \prod_{[x] \neq [y]} (a(x) - a(y))^{\mu_{[y]}}$ occur in both numerator and denominator of the right-hand side of (2.10). $\square$

Note that $N = \sum_{x \in \overline{\mathbb{Z}}} \mu(x)$. To compute the coefficients of $\chi(a + Sb, T)$, it is convenient to write it as

$$\chi(a + Sb, T) = \sum_{j=0}^N b_j(S)T^j, \quad \text{where } b_j \in \overline{\mathbb{F}}[S], \quad \deg b_j = N - j.$$

We need to be able to compute the trace $\text{Tr}(f(S))$ of a linear transformation $f(S) \in A[S]$. By additivity of the trace, this is easy to do once a basis $\overline{\mathbf{U}}(\mathcal{B})$ of A , the multiplication table for A in $\overline{\mathbf{U}}(\mathcal{B})$, that is, a tensor $\lambda_{\alpha\omega}^\iota$ specifying linear combinations

$$\alpha\omega = \sum_{\iota \in \overline{\mathbf{U}}(\mathcal{B})} \lambda_{\alpha\omega}^\iota \iota \quad \text{for } \alpha, \omega \in \overline{\mathbf{U}}(\mathcal{B}),$$

and the expression of $f(S)$ as a linear combination of the elements of $\overline{\mathbf{U}}(\mathcal{B})$ with coefficients in $\overline{\mathbb{F}}[S]$ are known. Namely,

$$\text{Tr}(f(S)) = \text{Tr}\left(\sum_{\omega \in \overline{\mathbf{U}}(\mathcal{B})} f_\omega(S)\omega\right) = \sum_{\alpha \in \overline{\mathbf{U}}(\mathcal{B})} \sum_{\omega \in \overline{\mathbf{U}}(\mathcal{B})} f_\omega(S)\lambda_{\alpha\omega}^\alpha,$$

and the computation can be done separately for each coefficient of the polynomial $\text{Tr}(f(S))$.

It is well-known (cf. e.g. [5, Thm. 4.69]) that for $f \in A$

$$(2.11) \quad \text{Tr}(f^j) = \sum_{x \in \overline{\mathbb{Z}}} \mu(x) f(x)^j, \quad j \geq 0.$$

It follows that

$$(2.12) \quad \text{Tr}((a + Sb)^j) = \sum_{x \in \overline{\mathbb{Z}}} \mu(x) (a(x) + Sb(x))^j = \sum_{[x] \subseteq \overline{\mathbb{Z}}} \mu_{[x]} (a(x) + Sb(x))^j, \quad j \geq 0.$$

Then, the $b_j(S)$'s (that is, the *elementary symmetric functions* of the roots) can be computed knowing the *power symmetric functions* (also known as *Newton sums*) $\text{Tr}(a + Sb), \dots, \text{Tr}((a + Sb)^N)$ of the roots of $\chi(a + Sb, T)$. Namely, the following holds (cf. [5, (4.2), (11.8)]).

$$\frac{\partial \chi(a + Sb, T)}{\partial T} = \chi(a + Sb, T) \sum_{m \geq 0} \frac{\text{Tr}((a + Sb)^m)}{T^{m+1}}.$$

By equating the coefficients of T^j , for each j satisfying $-1 \leq j < N$, on the both sides of the latter, and recalling that $\chi(a + Sb, T)$ is monic in T , that is, $b_N(S) = 1$, one obtains the following.

LEMMA 2.13. *Let $\chi(a + Sb, T) = \sum_{j=0}^N b_j(S) T^j$ be the characteristic polynomial of a linear transformation $a + Sb \in A[S]$. Then*

$$(2.14) \quad b_i(S) = -\frac{1}{N-i} \sum_{j=1}^{N-i} b_{i+j}(S) \text{Tr}((a + Sb)^j), \quad 0 \leq i \leq N-1, \quad b_N = 1$$

gives a recurrence for $b_i(S)$'s, for $i = N-1, N-2, \dots, 0$. □

REMARK 2.15. *Formulae similar to (2.14) are known since long time, and attributed to [12]. An explicit expression for $b_i(S)$ in terms of a determinant of certain “almost Toeplitz” matrix with entries specified by $\text{Tr}((a + Sb)^j)$'s can be found by using [13, Ex. I.2.8]. See also [13, (2.14)].*

Using the latter lemma, we can construct $\chi(a, T)$ and $g(a, b, T)$ given $a, b \in A$. We do not need to compute $\chi(a + Sb, T)$ completely; namely, only S -linear parts of $b_i(S)$ and $\text{Tr}((a + bS)^j)$ need to be computed, in view of (2.14) and

(2.6). To avoid the necessity to handle rational expressions arising from the term $\frac{1}{N-i}$ in (2.14), compute $N! \chi(a, T)$ and $N! g(a, b, T)$ instead.

In what follows we restrict ourselves to separating elements of the form $a(P(X))$, for $a \in D[T_1, \dots, T_m]$. Note that a P -separating a exists and can be chosen as follows, for some $0 \leq j \leq (m-1) \binom{N}{2}$:

$$(2.16) \quad a(P(X)) = a(j, P(X)) = \sum_{i=1}^m j^{i-1} P_i(X).$$

To see this, one proceeds as in e.g. [5, Lemma 4.60]. Let $s \neq y \in P(\overline{\mathbb{Z}})$, and observe that the univariate polynomial $a(Y, s) - a(Y, y) = \sum_{i=1}^m (s_i - y_i) Y^{i-1}$ is not identically 0, and has at most $m-1$ roots. Thus by avoiding at most $m-1$ values of j , one can make sure that a separates s and y . As there are at most $\binom{N}{2}$ distinct pairs of s and y as above, the claim follows.

Thus we can construct univariate representations of the elements $s \in P(\overline{\mathbb{Z}})$ of multiplicity $\mu(s) = \mu + 1$ in the form

$$(2.17) \quad u(T) = N! (\chi(a, T), g^{(\mu)}(a, r, T), g^{(\mu)}(a, rP_1(X), T), \dots, g^{(\mu)}(a, rP_m(X), T)),$$

where $r \in D[\varepsilon]$ is chosen so that the functions $ra(X), rP_1(X), \dots, rP_m(X)$ of A are $D[\varepsilon]$ -linear combinations of the basis elements of A . The latter are chosen so that the entries of the multiplication table of A belong to $D[\varepsilon]$, as dictated in turn by the coefficients of the leading monomials in $\mathcal{B}$. Taking r to be the LCM $b_{\mathcal{B}}$ of the coefficients of the leading monomials in $\mathcal{B}$ suffices; the numbers $\mu(s) \leq N$ are not known *a priori*, thus we have roughly N -fold redundancy in the output.

To summarize, we have the following.

PROPOSITION 2.18. *Let $\mathcal{B}$, P be as in Theorem 2.1, $A = \overline{\mathbb{F}}\langle \varepsilon \rangle[X]/(\mathcal{B})$ be of dimension N , and $a \in \mathbb{Z}[T_1, \dots, T_m]$ define a P -separating element $a(P(X))$ given by (2.16) with coefficients in $\mathbb{Z}$ of size at most $O(\log mN)$.*

Then a set of at most $N^{O(1)}$ univariate representations $u(T) \in D[\varepsilon]^{m+2}$ of the form (2.17), containing for each $s \in P(\overline{\mathbb{Z}})$, a representation (2.17) with $\mu = \mu(s) - 1$, can be computed in $N^{O(\ell)}$ arithmetic operations in D . The degrees of the polynomials in $u(T)$ are at most N , and their coefficients are of degree at most $O(d^3)$ in ε .

When $D = \mathbb{Z}$, the bitsize of the coefficients of $u(T)$ and of the intermediate data is bounded as stated in Theorem 2.1.

The complexity analysis is very similar to algorithms in [5, Chapter 11]. The most expensive part is computing the appropriate multiplication table for A , see the exposition preceding Theorem 2.1 and [5, Alg. 11.22]), and this is identical to the special case considered in [loc.cit.]. Note that P_i 's are expressed as linear combinations of the basis elements of A and therefore our setting for complexity analysis of computation of $u(T)$ is essentially the same as in [loc.cit.].

2.2. Computation of the limit. We proceed to computing the limits of the points given by univariate representations $u(T)$ of Proposition 2.18. We show that the limits of points in $P(\overline{Z})$ correspond to the limits $S_{<\infty}$ of bounded roots of $\chi(a, T)$. Then we normalize the polynomials in $u(T)$ by a Puiseux monomial in ε that makes the coefficients of $\chi(a, T)$ and of the rational functions $P_i(T) = \frac{g^{(\mu)}(a, bX_i, T)}{g^{(\mu)}(a, b, T)}$, for each $1 \leq i \leq m$, bounded. At the same time the values of the limit of the normalized denominator $g^{(\mu)}(a, b, T)$ will be nonzero on $S_{<\infty}$. Thus $\lim_{\varepsilon} P_i(T)$ at $S_{<\infty}$ can be computed by taking the limits of the coefficients of $P_i(T)$, and then evaluating on the elements of S_{∞} . We will give explicit formulae for $\lim_{\varepsilon} P_i(T)$ in terms of the appropriate repeated derivatives of the numerator and of the denominator.

We need some further notation related to a real closed extension $\mathbb{F}\langle\zeta\rangle$ of a real closed field $\mathbb{F}$ by infinitesimals $\zeta_1 \gg \zeta_2 \gg \dots \gg \zeta_{\ell}$, and its algebraic closure $\overline{\mathbb{F}}\langle\zeta\rangle$. Let $0 \neq \tau \in \overline{\mathbb{F}}\langle\zeta\rangle$. Then τ can be written uniquely as $\tau = \zeta^{o(\tau)}(\text{in}(\tau) + \tau')$, with $0 \neq \text{in}(\tau) \in \overline{\mathbb{F}}$ and $o(\tau) \in \mathbb{Q}^{\ell}$ such that $\zeta^{o(\tau)}$ is the biggest, with respect to the order in $\mathbb{F}\langle\zeta\rangle$, ζ -monomial of τ , and τ' satisfying $\lim_{\zeta} \tau' = 0$. In particular τ' is bounded over $\mathbb{F}$, and $\lim_{\zeta} \tau = \text{in}(\tau)$ iff $o(\tau) = 0$. Further, for $0 \neq v \in \overline{\mathbb{F}}\langle\zeta\rangle^n$ define $o(v) = \max_{1 \leq j \leq n} o(v_j)$, with max taken in the sense of the ordering in $\mathbb{F}\langle\zeta\rangle$.

Note that boundedness of $\tau = \Re\tau + i\Im\tau \in \overline{\mathbb{F}}\langle\zeta\rangle$ is understood here and elsewhere in this section in the usual sense of the norm $\sqrt{(\Re\tau)^2 + (\Im\tau)^2}$ being bounded over $\mathbb{F}$. As well, $\lim_{\zeta} \tau$ is understood purely algebraically, that is, in the appropriate order setting to 0 the corresponding infinitesimals. In just introduced notation, τ is bounded if and only if either $o(\tau) = 0$, or the rightmost nonzero entry of $o(\tau)$ is positive.

Let $f(T) = \sum_{j=0}^m c_j T^j \in \mathbb{F}(\zeta)[T]$. Then $o(f)$ is defined to be such that $\zeta^{o(f)}$ is minimal, with respect to the order in $\mathbb{F}\langle\zeta\rangle$, monomial making $\zeta^{-o(f)}c_j$ for $0 \leq j \leq m$ bounded over $\mathbb{F}$. In fact $o(f) = o(c_j)$ for some j . Define

$$(2.19) \quad \hat{f}(T) = \lim_{\zeta} \zeta^{-o(f)} f(T),$$

where the limit is taken coefficient-wise. For $F(T) \in \mathbb{F}(\zeta)[T]^q$, we denote

$$(2.20) \quad o(F) = \max_{1 \leq i \leq k} o(F_i).$$

The following statement is an extension of Lemma 11.37 from [5] adjusted to the non-multiplicity-free situation, and will be used repeatedly.

LEMMA 2.21. *Let $f(T) \in \mathbb{F}(\zeta)[T]$ be monic. Denoting $Z_f = Z(f(T), \overline{\mathbb{F}}(\zeta))$, one has*

$$o(f) = \sum_{\substack{\tau \in Z_f, \\ \text{unbounded}}} \mu(\tau) o(\tau), \quad Z(\hat{f}(T), \overline{\mathbb{F}}) = \lim_{\zeta} Z_f.$$

Let $y \in Z(\hat{f}(T), \overline{\mathbb{F}})$. Then $\mu(y) = \sum_{\substack{\tau \in Z_f, \\ \lim_{\zeta} \tau = y}} \mu(\tau)$ equals the multiplicity of y as a root of $\hat{f}$. Here the summands $\mu(\tau)$ denote multiplicities of roots τ of f .

PROOF. (Sketch.) The coefficients f_i of $f = \sum_{i=0}^d f_i T^i$, where $f_d = 1$, are elementary symmetric functions of $x \in Z_f$ taken with multiplicities. Let Σ denote the multiset of roots of f . Then

$$(2.22) \quad f_{d-i} = \sum_{\substack{\Theta \subseteq \Sigma \\ |\Theta|=i}} \prod_{\tau \in \Theta} \tau,$$

implying

$$o(f_{d-i}) \leq \max_{\substack{\Theta \subseteq \Sigma \\ |\Theta|=i}} \sum_{\tau \in \Theta} o(\tau),$$

where the inequality might be strict due to a possible cancellation of higher order terms in the sum (2.22). When one of the multisets Θ equals Υ , the sub-multiset of unbounded roots of f , this inequality turns into equality $\bar{o} = o(f_{d-|\Upsilon|}) = \sum_{\tau \in \Upsilon} o(\tau)$, as the order of the remaining summands in (2.22) is strictly less than $\bar{o}$. As $\bar{o} \geq \sum_{\tau \in \Theta} o(\tau)$ for any $\Theta \subseteq \Sigma$, and we obtain $o(f) = \bar{o}$. Therefore

$$(2.23) \quad \hat{f}(T) = \lim_{\zeta} \varepsilon^{-o(f)} f(T) = \lim_{\zeta} \prod_{\tau \in \Sigma - \Upsilon} (T - \tau) \prod_{\tau \in \Upsilon} \varepsilon^{-o(\tau)} (T - \tau) =$$

$$(2.24) \quad = \prod_{\tau \in \Upsilon} (-\text{in}(\tau)) \prod_{\tau \in \Sigma - \Upsilon} (T - \lim_{\zeta} \tau),$$

and the first part of the lemma follows. The second part follows from (2.24). $\square$

A P -separating element a will be called *well- P -separating* (with respect to $\lim_\varepsilon$) if the following conditions hold:

1. for any $s, y \in P(\overline{Z})$ such that $\lim_\varepsilon s \neq \lim_\varepsilon y$ one has $a(\lim_\varepsilon s) \neq a(\lim_\varepsilon y)$;
2. $o(P(u)) = o(a(P(u)))$ for any $u \in \overline{Z}$.

In particular, 2 implies that if $s \in P(\overline{Z})$ is unbounded over $\mathbb{F}$ then $a(s)$ is also unbounded over $\mathbb{F}$.

LEMMA 2.25. *Let $a, b \in D[T_1, \dots, T_m]$ be linear. Let a be well- P -separating. Then for any $s \in \lim_\varepsilon P(\overline{Z})$ of multiplicity $\mu_{[s]}$*

$$b(s) = \frac{\hat{g}^{(\mu_{[s]}-1)}(a, b, a(s))}{\hat{g}^{(\mu_{[s]}-1)}(a, 1, a(s))}. \quad \text{In particular, } s_j = \frac{\hat{g}^{(\mu_{[s]}-1)}(a, P_j, a(s))}{\hat{g}^{(\mu_{[s]}-1)}(a, 1, a(s))},$$

for any $1 \leq j \leq m$.

PROOF. As a is well- P -separating, b is an a -class function on $P(\overline{Z})$ as well as on $\lim_\varepsilon P(\overline{Z})$. Note that b satisfies $\lim_\varepsilon b(P(x)) = b(\lim_\varepsilon P(x))$, for any $x \in \overline{Z}$ for which $\lim_\varepsilon P(x)$ is defined. We shall express $\lim_\varepsilon b(P(X))$ as a univariate rational function, that gives $\lim_\varepsilon b(P(y))$ when evaluated at $a(P(y))$. With $a = a(P(X))$ and $b = b(P(X))$, denote

$$\eta(a, b, S, T) = \lim_\varepsilon \varepsilon^{-o(\chi(a, T))} \chi(a + Sb, T).$$

Denoting by $\overline{Z}_{<\infty}$ the set of $x \in \overline{Z}$ such that $a(P(x))$ is bounded, using (2.5) and Lemma 2.21 (in particular (2.24)) one obtains

$$\begin{aligned} \eta(a, b, S, T) &= \lim_\varepsilon \varepsilon^{-o(\chi(a, T))} \prod_{[x] \subseteq \overline{Z}} (T - a(P(x)) - Sb(P(x)))^{\mu_{[x]}} = \\ &= \prod_{[x] \subseteq \overline{Z}_{<\infty}} \lim_\varepsilon (T - a(P(x)) - Sb(P(x)))^{\mu_{[x]}} \times \\ &\times \prod_{[x] \subseteq \overline{Z} - \overline{Z}_{<\infty}} \lim_\varepsilon (\varepsilon^{-o(a(P(x)))} (T - a(P(x)) - Sb(P(x))))^{\mu_{[x]}} = \\ &= G(S) \prod_{[y] \subseteq \lim_\varepsilon P(\overline{Z})} (T - a(y) - Sb(y))^{\mu_{[y]}}, \quad \text{where} \\ G(S) &= \prod_{[x] \subseteq \overline{Z} - \overline{Z}_{<\infty}} (-\ln(a(P(x))) - Sb(\lim_\varepsilon \varepsilon^{-o(P(x))} P(x)))^{\mu_{[x]}} \in \overline{\mathbb{F}}[S], \end{aligned}$$

as $o(a(P(x))) = o(P(x))$. Moreover, $G(0) \neq 0$, as $\text{in}(a(P(x))) \neq 0$ on unbounded $x \in P(\overline{\mathbb{Z}})$.

In view of (2.24) one obtains in particular

$$(2.26) \quad \hat{\chi}(a, T) = C \prod_{[y] \subseteq \lim_{\varepsilon} P(\overline{\mathbb{Z}})} (T - a(y))^{\mu_{[y]}}, \quad \text{with } 0 \neq C \in \overline{\mathbb{F}},$$

where $\mu_{[y]}$ denotes the multiplicity of $a(y)$ as a root of $\hat{\chi}(a, T)$.

Next, we compute, for $g(a, b, T)$ defined by (2.6), $\lim_{\varepsilon} \varepsilon^{-o(\chi(a, T))} g(a, b, T)$. We see that it equals to $\hat{g}(a, b, T)$ (if it would not be the case, it had to vanish identically, as $o(\chi(a, T)) \geq o(g(a, b, T))$), as defined in (2.19), so we get

$$\begin{aligned} \hat{g}(a, b, T) &= \frac{\partial}{\partial S} \eta(a, b, S, T)|_{S=0} = -G'(0) \prod_{[y] \subseteq \lim_{\varepsilon} P(\overline{\mathbb{Z}})} (T - a(y))^{\mu_{[y]} - 1} \\ &\quad - G(0) \sum_{[y] \subseteq \lim_{\varepsilon} P(\overline{\mathbb{Z}})} \mu_{[y]} b(y) (T - a(y))^{\mu_{[y]} - 1} \prod_{[y] \neq [s] \subseteq \lim_{\varepsilon} P(\overline{\mathbb{Z}})} (T - a(s))^{\mu_{[s]}}. \end{aligned}$$

Then, for any $y \in \lim_{\varepsilon} P(\overline{\mathbb{Z}})$,

$$\hat{g}^{(\mu_{[y]} - 1)}(a, b, T) = -G(0) b(y) (\mu_{[y]})! \prod_{\substack{[s] \subseteq \lim_{\varepsilon} P(\overline{\mathbb{Z}} \\ [y] \neq [s]}} (T - a(s))^{\mu_{[s]}} + (T - a(y)) H(T),$$

for $H(T) \in \overline{\mathbb{F}}[T]$. Hence

$$\hat{g}^{(\mu_{[y]} - 1)}(a, b, a(y)) = -G(0) b(y) (\mu_{[y]})! \prod_{[y] \neq [s] \subseteq \lim_{\varepsilon} P(\overline{\mathbb{Z}})} (a(y) - a(s))^{\mu_{[s]}}$$

and we obtain, in view of (2.26), the statement of the lemma. $\square$

Let us show that a can be taken to be $a = a(j, P(X))$ for a certain j as in (2.16). The only difference with the argument above is that we have to avoid more “wrong” values of j . Let $x, y \in P(\overline{\mathbb{Z}})$ be such that $\lim_{\varepsilon} x$ and $\lim_{\varepsilon} y$ exist and are not equal. Then the polynomials $a(Y, \lim_{\varepsilon} x) - a(Y, \lim_{\varepsilon} y)$ and $a(Y, x) - a(Y, y)$ are not identically 0 and each of them has at most $m - 1$ roots. Thus by avoiding at most $2(m - 1)$ values of j , one can make sure that a separates $\lim_{\varepsilon} x$ and $\lim_{\varepsilon} y$, as well as x and y .

To ensure the remaining condition in the definition of well- P -separating element, consider $W = \{\lim_{\varepsilon} \varepsilon^{-o(s)} s \mid 0 \neq s \in P(\overline{\mathbb{Z}})\}$. Choose j such that $a(j, w) \neq 0$ for any $w \in W$. Such a choice is always possible: $a(Y, w) \in \mathbb{F}[Y]$

has at most $m - 1$ roots; thus avoiding $(m - 1)|W|$ values of j achieves the required. Then $o(a(s)) = o(s)$ for all $s \in P(\overline{Z})$, implying condition 2 of the definition.

As there are at most $\binom{N}{2}$ distinct pairs of s and y as above, and since $|W| \leq N$, we obtain

LEMMA 2.27. *There exists an integer $0 \leq j \leq (m - 1)N^2$ such that $a(P(X)) = a(j, P(X))$ is well- P -separating. $\square$*

Combining Lemmas 2.25 and 2.27 gives for an appropriate a at most $O(\sqrt{N})$ candidates for univariate representations $u(T)$ for the points in $\lim_\varepsilon P(\overline{Z})$, as $\chi(a, T)$ has at most $O(\sqrt{N})$ different root multiplicities. We outline now how $u(T)$ are actually computed. Let $b = P_i$ for some $1 \leq i \leq m$. We loop through $1 \leq j \leq (m - 1)N^2$ in order to be sure to find an appropriate well- P -separating $a(P(X)) = a(j, P(X))$. This means that we will return candidate representations for each such j .

A further technical point is that we operate in the ring $D[\varepsilon]$ rather than in a field. We utilize the idea of [5, Remark 11.44] for $r = b_{\mathcal{B}}$:

$$(2.28) \quad \chi(ra + Srb, rT) = r^N \chi(a + Sb, T)$$

and proceed similarly to the procedure of [5, Alg. 11.45].

REMARK 2.29. *Compared to [loc.cit.], our simplification is that we do not try to filter out “wrong” a ; such a check would require finding $|P(\overline{Z})|$ to be able to verify that a is P -separating, and then proceeding similarly to remarks on [5, p.398]. To obtain “good” a , one would first select a ’s with the biggest degree of $\chi(a, T)$; among the latter select a ’s that are P -separating, by choosing a ’s with minimal degree of $\gcd(\chi(a, T), \frac{d\chi(a, T)}{dT})$. To ensure that a is well- P -separating, one would select a ’s with minimal degree of $\gcd(\hat{\chi}(a, T), \frac{d\hat{\chi}(a, T)}{dT})$. Finally, to make sure that $o(a(x)) = o(x)$ for any $x \in P(\overline{Z})$, one would check that $o(\chi(a, T)) \geq o(g(a, P_i, T))$ for all $1 \leq i \leq m$.*

Anyhow, this shortcut does not worsen the asymptotic running time.

Thus we compute $\chi(ra + Srb, T)$, as already outlined in the first part of this section, and then operate with $\varepsilon^{-o(\chi(ra, rT))} \chi(ra + Srb, rT)$ to obtain the remaining data for the limit computation. We have

$$\hat{\chi}(ra + Srb, rT) = \lim_{\varepsilon} \varepsilon^{-o(\chi(ra, rT))} \chi(ra + Srb, rT).$$

Using the latter to compute $\hat{g}(ra, rb, bT)$ and $\hat{g}(ra, r, bT)$ as in the proof of Lemma 2.25, by (2.28) we obtain $b(x) = \frac{\hat{g}^{(\mu)}(ra, rb, ra(x))}{\hat{g}^{(\mu)}(ra, r, ra(x))}$ for each $x \in \lim_\varepsilon P(\overline{Z})$

of multiplicity $\mu + 1$, as required. For each j we loop through all the possible values of μ . Thus in total we will have no more than $(m - 1)N^3$ univariate representations, as stated in the theorem.

The complexity analysis for a given j runs parallel to the analysis given in [5, Sect. 11.5] (see also a remark following Proposition 2.18 above) and is omitted. This completes the proof of Theorem 2.1.

3. Limits of solution images: dimension > 0

In Section 2 we described a procedure to compute limits of images of finite algebraic sets under polynomial mappings. To complete the proof of Theorem 1.10, we proceed to computing limits, with respect to $\varepsilon \rightarrow 0$, of rational images of samples of connected components of real algebraic sets, by reducing to the 0-dimensional case.

For $F_0 \in D[\varepsilon][Y_1, \dots, Y_{q-1}]$, let $Z_0 = Z(F_0(Y), \mathbb{F}\langle\varepsilon\rangle^{q-1})$, and let Ψ be as in (1.9). We want to find points in each connected component of $\lim_{\varepsilon} \Psi(Z_0) \subseteq \mathbb{F}^m$.

THEOREM 3.1. *Let F_0 and Ψ be as above, with the Y -degree of F_0 at most d and the Y -degrees of Ω_i and Λ less than $d - 1$, and their ε -degrees at most d . Then one can construct a $(q + 2)$ -variate 0-dimensional polynomial system $\mathcal{B}$ over the real closed transcendental extension $\mathbb{F}\langle\varepsilon_1, \dots, \varepsilon_{\ell+2}\rangle$ of $\mathbb{F}\langle\varepsilon\rangle$ by two infinitesimals $0 < \varepsilon_{\ell+2} \ll \varepsilon_{\ell+1} \ll \varepsilon_{\ell}$, with zero set $Z(\mathcal{B})$ such that*

$$\mathcal{L} = \lim_{\varepsilon_1, \dots, \varepsilon_{\ell+2}} \Psi(\pi(Z(\mathcal{B}))),$$

where π is the orthogonal projection to $Y_1, \dots, Y_{q-1}$, intersects each connected component of $\lim_{\varepsilon_1, \dots, \varepsilon_{\ell}} \Psi(Z_0)$.

The system $\mathcal{B}$ can be constructed in $(m + d)^{O(q\ell)}$ operations over D , and it satisfies the properties required by Theorem 2.1; the set $\mathcal{L}$ can be computed using the procedure of Theorem 2.1.

In the proof we will denote $\mu = \varepsilon_{\ell+1}$ and $\zeta = \varepsilon_{\ell+2}$.

First of all, let us reduce the setting to the case when $\Lambda = 1$. Assume that $F_0(Y) \geq 0$, otherwise replace F_0 with F_0^2 . Consider

$$F(Y_1, \dots, Y_q) = F_0(Y) + (1 - Y_q \Lambda(Y))^2, \quad Z = Z(F(Y), \mathbb{F}\langle\varepsilon\rangle^q).$$

Then $\Psi(Z_0) = P(Z)$, where $P \in D[\varepsilon][Y_1, \dots, Y_q]^m$ is given by

$$(Y_1, \dots, Y_q) \mapsto (Y_q \Omega_1(Y), \dots, Y_q \Omega_m(Y)).$$

Thus from now on we consider the problem of computing representatives of connected components of the set $\lim_\varepsilon P(Z) \subseteq \mathbb{F}^m$. Note that the map P is as required by Theorem 2.1.

The idea is to look for points of minimal norm in the connected components of $P(Z)$. We will repeatedly use the following statement.

PROPOSITION 3.2. (*[5, Prop. 11.56]*) *Let $S \subseteq \mathbb{K}\langle\varepsilon\rangle^n$ be a semialgebraic set, for $\mathbb{K}\langle\varepsilon\rangle$ a real closed extension of a real closed field $\mathbb{K}$ by an infinitesimal ε . Then $\lim_\varepsilon S \subseteq \mathbb{K}^n$ is a closed semialgebraic set. If S is in addition connected and bounded over $\mathbb{K}$ then $\lim_\varepsilon S$ is connected.* $\square$

Introduce another variable Y_0 that will help us in this task; we will write down, after necessary preparations, equations that specify the critical points of the projection $Y \mapsto Y_0$. Replace F by

$$F'(Y_0, \dots, Y_q) = F(Y) + (Y_0 - \sum_{i=1}^m P_i(Y)^2)^2.$$

We extend P to the zero set Z' of F' by “ignoring” Y_0 . Obviously $P(Z) = P(Z')$.

Next, we need to make the set Z' bounded, in the standard way of e.g. [5, Chapt. 11]. Introduce an infinitesimal $0 < \mu \ll \varepsilon_\ell$, a variable Y_{q+1} and define

$$F_\mu(Y_0, \dots, Y_{q+1}) = F'(Y) + (1 - \mu^2 \sum_{j=0}^{q+1} Y_j^2)^2, \quad Z_\mu^{<\infty} = Z(F_\mu, \mathbb{F}\langle\varepsilon, \mu\rangle^{q+2}).$$

Again, extend P onto $Z_\mu^{<\infty}$ by “ignoring” Y_{q+1} . Define

$$Z_\mu = Z(F', \mathbb{F}\langle\varepsilon, \mu\rangle^{q+1})$$

LEMMA 3.3. *The following holds:*

- (i) $P(Z_\mu^{<\infty})$ is closed;
- (ii) $P(Z_\mu^{<\infty}) \subseteq P(Z_\mu)$;
- (iii) $\lim_\mu P(Z_\mu^{<\infty})$ equals the closure of $P(Z)$.

PROOF. The first part of the lemma follows from [5, Thm. 3.20], as $P(Z_\mu^{<\infty})$ is the image of a closed and bounded semialgebraic set under a continuous semialgebraic function.

The second part is straightforward, by observing that the projection of $Z_\mu^{<\infty}$ on the first $q+1$ coordinates is a subset of Z_μ .

Then, $\lim_\mu P(Z_\mu^{<\infty})$ is closed by Proposition 3.2. Next, $\lim_\mu P(Z_\mu^{<\infty}) \supseteq P(Z)$. Indeed, if $u = P(y)$ with $y \in Z$ then there exists $y_{q+1} \in \mathbb{F}\langle \varepsilon, \mu \rangle$ such that $\mu^2 \|y_\mu\|^2 = 1$, where we denoted $y_\mu = (y_0, y_1, \dots, y_{q+1}) \in Z_\mu^{<\infty}$, implying $F_\mu(y_\mu) = 0$ and $u = P(y_\mu) \in P(Z_\mu^{<\infty})$.

Finally, the inclusion of $\lim_\mu P(Z_\mu^{<\infty})$ in the closure of $P(Z)$ follows from the second part. $\square$

The next step is to introduce a deformation, same as in [5, Sect. 11.6], that ensures smoothness. As in [5, (11.13)-(11.14)], introduce an infinitesimal $0 < \zeta \ll \mu$, and set

$$(3.4) \quad F_\zeta(Y) = \zeta \left(\mu^{\bar{d}} \left(Y_0^{\bar{d}} + \sum_{j=1}^{q+1} (Y_j^{\bar{d}} + Y_j^2) \right) - 2q - 3 \right) + (1 - \zeta) F_\mu(Y).$$

Note that the difference between [5, (11.13)-(11.14)] and (3.4) is purely notational: variables and infinitesimals have different names, and they range slightly differently (i.e. there are $q+2$ variables Y instead of q variables X). Here $\bar{d}$ is the minimal positive even number strictly bigger than the degree of F_μ ; this simplification of [5, (11.13)-(11.14)] is explained in [5, Rem. 11.49].

By [5, 11.50-11.51], the set $Z_\zeta = Z(F_\zeta, \mathbb{F}\langle \varepsilon, \mu, \zeta \rangle^{q+2})$ is a nonsingular algebraic hypersurface contained in the ball of radius $\frac{1}{\mu}$, and such that $\lim_\zeta Z_\zeta = Z_\mu^{<\infty}$. In particular $\lim_\zeta u$ exists for any $u \in Z_\zeta$.

Consider the set K_ζ of the critical points of the projection $Y \mapsto Y_0$ on Z_ζ . They satisfy the system of equations

$$(3.5) \quad 0 = F_\zeta(Y) = \frac{\partial F_\zeta}{\partial Y_1} = \dots = \frac{\partial F_\zeta}{\partial Y_{q+1}}.$$

The zero set $\overline{K}_\zeta \supseteq K_\zeta$ of (3.5) in the algebraic closure of $\mathbb{F}\langle \varepsilon, \mu, \zeta \rangle$ is analysed in detail in [5, Sect. 11.6]. In particular, $|\overline{K}_\zeta| \leq d^{O(q)}$ (cf. [5, Prop. 11.57]). Note that $F_\zeta(Y)$ can be reduced (cf. [5, Notation 11.59]) using the relations $\frac{\partial F_\zeta}{\partial Y_j} = 0$, $j \geq 1$ so that the degree in Y_j of the result is strictly less than $\bar{d}$. The result of these reductions, together with the $\frac{\partial F_\zeta}{\partial Y_j}$, $j \geq 1$, forms the polynomial system $\mathcal{B}$ of Theorem 3.1 we are proving.

To complete the proof that $\mathcal{L}$ is as claimed, it remains to prove the following lemma. It will make sure that we recover in $\mathcal{L}$ a representative of each component of $\lim_{\varepsilon} P(Z)$.

LEMMA 3.6. *The set $\lim_{\zeta, \mu} P(K_{\zeta})$ intersects each connected component of $P(Z)$ in a point of minimal norm.*

PROOF. By construction, $P(K_{\zeta})$ contains the local minima of $Y \mapsto Y_0$ on $P(Z_{\zeta})$.

By [5, 11.55], the set $K_{\mu} = \lim_{\zeta} K_{\zeta}$ (the set of *pseudo-critical points*) intersects each connected component of $Z_{\mu}^{<\infty}$.

As P does not depend upon ζ , we have that

$$\lim_{\zeta} P(Z_{\zeta}) = P(\lim_{\zeta} Z_{\zeta}) = P(Z_{\mu}^{<\infty}).$$

Due to closedness and boundedness over $\mathbb{F}\langle\varepsilon, \mu\rangle$ of $P(Z_{\mu}^{<\infty})$ (cf. Lemma 3.3), $Y \mapsto Y_0$ reaches its minimum on each connected component $C_{\mu} \subseteq P(Z_{\mu}^{<\infty})$. Note that by Proposition 3.2 we have $C_{\mu} = \lim_{\zeta} C_{\zeta}$, for a connected $C_{\zeta} \subseteq P(Z_{\zeta})$. Consider $P^{-1}(C_{\mu}) \subseteq Z_{\mu}^{<\infty}$. This set is a union of connected components of $Z_{\mu}^{<\infty}$, and K_{μ} intersects each of them; among these intersection points there are ones with minimum value of Y_0 . Therefore $P(K_{\mu})$ contains representatives of the local minima (that need not be singletons any more) of $Y \mapsto Y_0$ on $P(Z_{\mu})$.

As $P(Z_{\mu}^{<\infty})$ is the intersection of $P(Z_{\mu})$ with $P(B_{1/\mu})$, where $B_{1/\mu}$ denotes the ball of radius $\frac{1}{\mu}$ and centre at the origin, each connected component C of $P(Z_{\mu})$ that intersects $P(B_{1/\mu})$ will contain a connected component C_{μ} of $P(Z_{\mu}^{<\infty})$.

We assume now, without loss of generality, that C contains a point of $\mathbb{F}\langle\varepsilon\rangle$ -finite norm. Consider the intersection C_r of C and the (smaller) set $P(B_r)$, where $1 + \min_{u \in C} \|u\| < r \in \mathbb{F}\langle\varepsilon\rangle$. This is a bounded over $\mathbb{F}\langle\varepsilon\rangle$ semialgebraic set, with finitely many connected components C_r^i . By Proposition 3.2 the set $\lim_{\mu} C_r^i$ is connected, for any i . As $C_r^i \subseteq C_{\mu}$, and as Y_0 has a local minimum on C_r^i , each $C_r^i \cap P(K_{\mu}) \neq \emptyset$. Thus $\lim_{\mu} P(K_{\mu})$ intersects each connected component of $P(Z)$ in a point of minimal norm, as required. $\square$

A straightforward count of the number of operations needed to construct $\mathcal{B}$ completes the proof of Theorem 3.1.

To complete the proof of Theorem 1.10, apply Theorem 3.1 (with $\ell = 2$) to $Z_0 := Z_F$ and $F_0 := F$. We obtain a 0-dimensional system $\mathcal{B}$ of equations in the real closed extension of $\mathbb{F}\langle\varepsilon\rangle$ by two extra infinitesimals $\mu = \varepsilon_3$ and

$\zeta = \varepsilon_4$, with the set of solutions $\overline{K}_\zeta$ such that $\lim_{\mu, \zeta} \Psi(\pi(\overline{K}_\zeta))$ intersects each connected component of $\Psi(Z_0)$ in a point of minimal norm. Hence applying to $\mathcal{B}$ the procedure of Theorem 2.1 will produce a set $\lim_{\varepsilon, \mu, \zeta} \Psi(\pi(\overline{K}_\zeta))$ intersecting each connected component of $\lim_\varepsilon \Psi(Z_F)$ in required running time, as $N \leq d^q$. $\square$

4. Pieces of extrema of a bounded level set over a quadratic map

Let the quadratic map $Q : \mathbb{K}^n \rightarrow \mathbb{K}^k$ be given by

$$(4.1) \quad \begin{aligned} Q_j : X &\mapsto \frac{1}{2} X^T H_j X + b_j^T X + c_j, \quad 1 \leq j \leq k \\ c_j &\in D, \quad b_j \in D^n, \quad H_j = H_j^T \in D^{n \times n}. \end{aligned}$$

Let $\zeta \in \mathbb{K}$. We consider the ζ -level set $V = Z(p(Q(X)) - \zeta, \mathbb{K}^n)$ of a polynomial $p \in D[Y_1, \dots, Y_k]$ of degree d over a quadratic map $Q : \mathbb{K}^n \rightarrow \mathbb{K}^k$ given by (4.1).

We assume that D is a computable subring of a real closed field $\mathbb{K}$ and V is bounded over $\mathbb{K}$. The latter assumption is in fact technical. It can (and will) be ensured in Section 5 by introducing an extra infinitesimal ε_0 and two extra variables, to make V the image of an algebraic set on the sphere of radius $1/\varepsilon_0$ under projection onto the first n coordinates, and subsequently removing ε_0 by restricted elimination.

As well, we assume that ζ is a *regular value* of $p(Q(X))$ and of $p(Y)$.

Let $V_c \subseteq V$ denote the set of critical points of the projection map $X \mapsto X_1$ from V to $\mathbb{K}$. Then, as ζ is not a critical value of $p(Q(X))$, V_c is an algebraic set defined by

$$(4.2) \quad p(Q(X)) - \zeta = 0,$$

$$(4.3) \quad \frac{\partial p(Q(X))}{\partial X_2} = \dots = \frac{\partial p(Q(X))}{\partial X_n} = 0.$$

Due to the assumption that V is bounded, V_c intersects nontrivially each connected component of V , see [10], [5, Prop. 7.6]. Thus a set S_c of representatives of each connected component of V_c will also intersect each component of V . A useful property of S_c is that it will contain points in V with minimal value of X_1 .

Define

$$p_j(Y) = \frac{\partial p(Y)}{\partial Y_j}, \quad \text{for } 1 \leq j \leq k.$$

The set of $m \times n$ matrices $A_1, \dots, A_k$ over $\mathbb{K}$, $m \leq n$, is said to be in *r-general position* with respect to V if

$$(4.4) \quad \text{rk}\left(\sum_{i=1}^k t_i A_i\right) \geq r \quad \text{for all } t \in (p_1(Q(V)), \dots, p_k(Q(V))) \subseteq \mathbb{K}^k.$$

Note that t is never 0 here, as ζ is not a critical value of $p(Y)$. This is a weaker notion of general position than the one used in [2], where t was allowed to range over $\overline{\mathbb{K}}^k$, instead of belonging to the image of V under the polynomial mapping $X \mapsto (p_1(Q(X)), \dots, p_k(Q(X)))$.

Note that all the nonzero $\overline{\mathbb{K}}$ -linear combinations of *almost* any k -set of $m \times n$, $m \leq n$, matrices over $\mathbb{K}$ have rank at least $m + 1 - k$.

Below we assume that the matrices $\hat{H}_j$'s, that are obtained from H_j 's by removing the first row, are in *r-general position* with respect to V . For $U \subseteq \{1, \dots, m\}$ we denote $\overline{U} = \{1, \dots, m\} - U$. For $U \subseteq \{1, \dots, m\}$ and $W \subseteq \{1, \dots, n\}$ we denote by A_{UW} the submatrix of A obtained by removing all the rows in $\overline{U}$ and all the columns in $\overline{W}$.

To compute S_c in less than exponential in n time, or even just to give an upper bound on its size of order less than exponential in n , standard methods such as one in [10], [5, Chapter 11] that treat the system (4.2)-(4.3) directly do not suffice.

Indeed, in these methods the number of variables, n in this case, appears unavoidably in the exponent of the bounds. In contrast, we are able to get better results for r being close to n . We cover V_c by (at most $n^{O(n-r)}$) semialgebraic sets we call *pieces*, each of them isomorphic to a semialgebraic set lying in $\mathbb{K}^t$ with $t \leq k + n - r$. (Here covering means simply that the union of the pieces is V_c ; they in general intersect, and can even be equal one to another.) Each of the latter is defined by at most $O(dn^2)$ polynomials of degree at most $O(nd)$. To them, one can apply the standard technique, see e.g. [5], to bound the number of its connected components, so there will be at most $(nd)^{O(k+n-r)}$ of them, and to find their (perhaps non-unique) representatives. However, for the latter, for a technical reason, namely the necessity to take limits with respect to certain infinitesimals (in particular ζ will be treated as such), we shall use the approach presented in Sections 2 and 3.

The main results of this section are as follows.

THEOREM 4.5. *Let ζ be a regular value of $p(Q(X))$ and $p(Y)$, and the level set $V = Z(p(Q(X)) - \zeta, \mathbb{K}^n)$ be bounded over $\mathbb{K}$. Further, let the matrices $\hat{H}_j$'s be in *r-general position* with respect to V . Then one can construct a covering of*

the set of critical points $V_c \subseteq V$ of the projection map $X \mapsto X_1$ by semialgebraic sets $V_c(U, W)$ indexed by invertible submatrices of $\sum_j p_j(Q(X))H_j$ with row sets $U \subseteq \{2, \dots, n\}$ and column sets $W \subset \{1, \dots, n\}$ that satisfy $r \leq |U| = |W| \leq n - 1$. For each such W , the polynomial mapping

$$(4.6) \quad \begin{aligned} \phi_W : \mathbb{K}^n &\rightarrow \mathbb{K}^{k+n-|W|} \\ X &\mapsto \begin{pmatrix} Q(X) \\ X_{\overline{W}} \end{pmatrix} \end{aligned}$$

on each $\phi_W(V_c(U, W))$ has explicitly given inverse

$$\begin{aligned} \phi_{UW}^{-1} : \mathbb{K}^{k+n-|W|} &\rightarrow \mathbb{K}^n \\ \begin{pmatrix} Y \\ T \end{pmatrix} &\mapsto \begin{pmatrix} \Theta_{UW}(p_1(Y), \dots, p_k(Y), T) \\ T \end{pmatrix}, \end{aligned}$$

where Θ_{UW} is a vector of rational functions in $p_j(Y)$'s and T , all with the same denominator. The degrees of the latter and of the numerators are at most $|W|$. The set $\phi(V_c(U, W))$ is defined by $p(Y) = \zeta$ and $(n - |W|)^2 + k + 1 + n - r$ polynomial equations and one un-equation in the p_j 's, Y , and T , of degree $O(|W|)$. In total there are at most $n^{O(n-r)}$ pieces.

REMARK 4.7. The statement of Theorem 4.5 holds un-amended in the more general setting of p being a differentiable function, and “semialgebraic set” replaced by “joint zeros of equations and non-zeros of inequations”. The proof we give goes through in this case, too. We chose to remain in semialgebraic setting for the sake of clarity only.

The following summarizes, in the setting of the paper, the necessary complexity estimates for the procedure outlined in the course of proving Theorem 4.5.

PROPOSITION 4.8. In notation of Theorem 4.5, let $\mathbb{K} = \mathbb{L}\langle \varepsilon \rangle$, with $\varepsilon = (\varepsilon_1, \dots, \varepsilon_\ell)$ infinitesimals, $\ell \geq 1$, $D = D_{\mathbb{L}}[\varepsilon]$, and $p \in D[Y]$ of degree d in Y . Let the degrees of Q and p in ε be at most d' . The map ϕ_{UW}^{-1} and the definition for $\phi(V_c(U, W))$ can be computed within $(n(d + d'))^{O(\ell + k + n - r)}$ arithmetic operations in $D_{\mathbb{L}}$. In the case $D_{\mathbb{L}} = \mathbb{Z}$, the bitsizes of them and the intermediate data are bounded by the bitsize of Q and p times a polynomial in n , k , $\log d$, $\log(1 + d')$, and ℓ .

In the course of the proof of Theorem 4.5 we construct the explicit maps ϕ_{UW}^{-1} and semialgebraic definitions for $V_c(U, W)$. Equations (4.3) can be written as

$$\frac{\partial p(Q(X))}{\partial X_i}(X) = \sum_{j=1}^k p_j(Q(X)) e_i^T (H_j X + b_j) = 0, \quad 2 \leq i \leq k,$$

where e_i is the standard i -th coordinate vector. Thus, denoting

$$\Phi(Y) = \sum_{j=1}^k p_j(Y) \hat{H}_j, \quad b(Y) = - \sum_{j=1}^k p_j(Y) \hat{b}_j,$$

where $\hat{b}_j$ denotes b_j with the first coordinate removed, one can write (4.3) in the matrix form, as follows.

$$(4.9) \quad \Phi(Q(X))X = b(Q(X)).$$

At this point we would like to give an outline of the remainder of the proof of Theorem 4.5. We will compute a set of solutions of the system of equations (4.2)-(4.9) that intersects each connected component of V . Doing this in the standard way would take an exponential in n number of operations, and this is exactly what we want to avoid.

The structure of the equations (4.2)-(4.9) suggests the substitution $Y = Q(X)$. It turns (4.9) into a system of linear equations $\Phi(Q(Y))X = b(Y)$ with respect to X . We cannot simply “invert” $\Phi(Q(Y))$, as it need not be of full rank. However, $\text{rk}(\Phi(Q(Y)))$ will always be at least r , allowing us to “split” the solving into $n^{O(n-r)}$ cases, one for each maximal invertible submatrix (parametrized by U and W), that will be inverted. This gives (for each such case) rational expressions ϕ_{UW}^{-1} for r of the X ’s in terms of Y and the remaining $n - r$ of X ’s (that we will denote T), as well as giving extra (in)equations involving Y and T . The latter define the sets $V_c(U, W)$, essentially completing the proof.

We proceed with the detailed proof now, preparing the ground for introducing in (4.18) the variables Y . The r -general position assumption implies

$$\text{rk}(\Phi(Q(x))) \geq r \quad \text{for any } x \in V.$$

Thus there are at most $n^{O(n-r)}$ maximal by inclusion invertible submatrices of $\Phi(Q(X))$. Indeed, there are at most $s(r) = (\sum_{m=r}^n \binom{n}{m})^2$ of them, by counting number of pairs (U, W) of subsets $U \subseteq \{2, \dots, n\}$, $W \subseteq \{1, \dots, n\}$ satisfying $r \leq |U| = |W|$. If $r \leq n/2$ then $s(r) < 2^{O(n)} < n^{O(n-r)}$. Otherwise one has $s(r) < (n-r)^2 \binom{n}{n-r}^2 < n^{O(n-r)}$, again as required.

As well, at least one $\Phi(Q(X))_{UW}$ will be invertible. Hence

$$(4.10) \quad \begin{aligned} \det \Phi(Q(X))_{U'W'} &= 0 \quad \text{for all } |U'| = |W'| = |U| + 1 \\ U \subset U' \subseteq \{2, \dots, n\}, \quad W \subset W' \subseteq \{1, \dots, n\}, \\ \det \Phi(Q(X))_{UW} &\neq 0. \end{aligned}$$

Noting that (4.10) implies that $\det \Phi(Q(X))_{U'W'} = 0$ for all U' and W' of size bigger than $|U|$, one obtains the following.

LEMMA 4.11.

$$(4.12) \quad V_c = \bigcup_{\substack{U \subseteq \{2, \dots, n\} \\ W \subseteq \{1, \dots, n\} \\ r \leq |U| = |W| \leq n-1}} V_c(U, W),$$

where $V_c(U, W)$ is defined by the equations (4.2)-(4.3) and by (4.10). The number of elements in the union in (4.12) is at most $n^{O(n-r)}$.

Without loss in generality $U = \{2, \dots, r+1\}$, $W = \{1, \dots, r\}$. Invert the matrix $\Psi = \Phi(Q(X))_{UW}$ using the Cramer rule:

$$(4.13) \quad \Psi_{ij}^{-1} = \frac{(-1)^{i+j} \det \Psi(i, j)}{\det \Psi}, \quad \text{for } 1 \leq i, j \leq r,$$

where $\Psi(i, j)$ is the matrix obtained from Ψ by removing i -th row and j -th column. Then the system $\Phi(Q(X))X = b(Q(X))$ can be rewritten in the block form as

$$(4.14) \quad \begin{pmatrix} \Psi & \Phi_{U, \overline{W}} \\ \Phi_{\overline{U}, W} & \Phi_{\overline{U} \overline{W}} \end{pmatrix} \begin{bmatrix} X_W \\ X_{\overline{W}} \end{bmatrix} = \begin{bmatrix} b_U \\ b_W \end{bmatrix},$$

where the common “ $(Q(X))$ ” is dropped for the sake of readability. Applying $\begin{pmatrix} \Psi^{-1} & 0 \\ -\Phi_{\overline{U}, W} \Psi^{-1} & I \end{pmatrix}$ to both sides of (4.14), one obtains

$$(4.15) \quad \begin{pmatrix} I & \Psi^{-1} \Phi_{U, \overline{W}} \\ 0 & 0 \end{pmatrix} \begin{bmatrix} X_W \\ X_{\overline{W}} \end{bmatrix} = \begin{bmatrix} \Psi^{-1} b_U \\ b_W - \Phi_{\overline{U}, W} \Psi^{-1} b_U \end{bmatrix}.$$

Thus the following, together with (4.10) and (4.2), provides another definition of $V_c(U, W)$.

$$(4.16) \quad X_W = \Phi(Q(X))_{UW}^{-1} \cdot (b(Q(X))_U - \Phi(Q(X))_{U\overline{W}} X_{\overline{W}}),$$

$$(4.17) \quad b(Q(X))_{\overline{U}} = \Phi(Q(X))_{\overline{U}W} \Phi(Q(X))_{UW}^{-1} \cdot b(Q(X))_U.$$

Observe that in the latter definition of $V_c(U, W)$ the only place where X_W appears other than as an argument to Q is the left-hand side of (4.16). Set up the mapping ϕ_{UW}^{-1} as follows.

$$(4.18) \quad \phi_{UW}^{-1} : \mathbb{K}^{k+n-r} \rightarrow \mathbb{K}^n$$

$$\begin{pmatrix} Y \\ T \end{pmatrix} \mapsto \begin{pmatrix} \Phi(Y)_{UW}^{-1} \cdot (b(Y)_U - \Phi(Y)_{U\overline{W}}T) \\ T \end{pmatrix}.$$

One establishes that ϕ_{UW}^{-1} acts as claimed in the statement of the theorem.

LEMMA 4.19. *The mapping ϕ_W restricted onto $V_c(U, W)$ has inverse ϕ_{UW}^{-1} , that is,*

$$\phi_{UW}^{-1}(\phi_W(V_c(U, W))) = V_c(U, W).$$

PROOF. We have to check that $\phi_{UW}^{-1}(\phi_W(x)) = x$ for any $x \in V_c(U, W)$. As neither ϕ_W nor ϕ_{UW}^{-1} change $x_{\overline{W}}$, in view of (4.18) it suffices to check that $\Phi(Q(x))_{UW}^{-1}(b(Q(x))_U - \Phi(Q(x))_{U\overline{W}}x_{\overline{W}}) = x_W$. But the latter holds as it is nothing but (4.16), a part of a definition for $V_c(U, W)$. $\square$

We proceed to write down an explicit definition for $\phi_W(V_c(U, W))$ in terms of variables Y and T used in (4.18). Denoting $\Omega = \det \Phi(Y)_{UW}$, one has the following polynomial equations

$$(4.20) \quad p(Y) = \zeta$$

$$(4.21) \quad \Omega^2 Y = \Omega^2 Q(\phi_{UW}^{-1}(Y, T))$$

$$(4.22) \quad \Omega b(Y)_{\overline{U}} = \Omega \Phi(Y)_{\overline{U}, W} \Phi(Y)_{UW}^{-1} \cdot b(Y)_U,$$

where puzzlingly looking multiplication of both sides of (4.21) and (4.22) by Ω clears denominators coming from (4.13) in the right-hand sides, and

$$(4.23) \quad \begin{aligned} \det \Phi(Y)_{U'W'} &= 0 \quad \text{for all } |U'| = |W'| = |U| + 1 \\ U \subset U' \subseteq \{2, \dots, n\}, \quad W \subset W' \subseteq \{1, \dots, n\}, \\ \det \Phi(Y)_{UW} &\neq 0. \end{aligned}$$

Apart from (4.21), that “bootstraps” Q , these (in)equations already appeared above, with Y substituted for $Q(X)$ and T substituted for $X_{\overline{W}}$.

LEMMA 4.24. *The relations (4.20)-(4.23) provide a semialgebraic definition of $\phi_W(V_c(U, W))$.*

PROOF. Let (y, t) belong to the semialgebraic set defined by (4.20)-(4.23), and $x = \phi_{UW}^{-1}(y, t)$. We shall check that $x \in V_c(U, W)$. Due to (4.21), the equation (4.2) holds for $X = x$. Similarly, the remaining (in)equations (4.10), (4.16)-(4.17) defining $V_c(U, W)$ hold for $X = x$.

By inspection, any $x \in V_c(U, W)$ gives rise to $(y, t) = \phi_W(x)$ in the set defined by (4.20)-(4.23). $\square$

The entries of the matrix $\Phi(Y)$ are linear polynomials in $p_1, \dots, p_k$. Thus the determinants of its $m \times m$ -submatrices, that come via (4.13) into the definition of $V_c(U, W)$ by (4.20)-(4.23), will be polynomials of degree at most m . Similar straightforward degree counts complete the proof of Theorem 4.5.

We proceed to prove Proposition 4.8. The only nontrivial part concerns the complexity of computing the map $\phi^{-1}(U, W)$ using the Cramer rule (4.13).

To count the number of arithmetic operations in D_L required to compute the determinants, one can either slightly extend [2, (2.8)], or refer to [5, Alg. 8.38, Rem. 8.39(b)], to obtain that computing the determinant of a submatrix of $\Phi(Y)$ can be done in $n^{O(1)}$ arithmetic operations in $D_L[\varepsilon, Y]$, and then refer to [5, Alg. 8.38].

Using the latter source, one sees that for $D_L = \mathbb{Z}$, the bitsizes in the answer and in the intermediate data will be bounded by $(\tau + \log n)n + (k+1) \log(n(d+d') + 1)$, with τ a bound on the bitsize of coefficients in the entries of $\Phi(Y)$. Noting that τ is bounded by $\log d$ times the bitsize of p and Q completes the computation of the bitsize bound for determinants. The remainder of the proof is a straightforward use of the complexity analysis of arithmetic operations in polynomial rings in [5, Algs. 8.8, 8.10].

5. Proof of Theorem 1.2

Let $\varepsilon_0 > 0$ be an infinitesimal over $\mathbb{K}$. We use it to deform $Z = Z(p(Q(X)), \mathbb{K}^n)$ so that it becomes bounded. (*A priori* this deformation is not necessary if Z is known to be bounded in the first place.) We can assume that $p(Y) \geq 0$ for all Y , otherwise we can replace p by p^2 . Introduce extra variables X_0 and Y_0 , and set

$$\tilde{p}(Y) = Y_0^2 + p(Y).$$

Set

$$Q_0(X) = 1 - \varepsilon_0^2 \sum_{i=0}^n X_i^2$$

and abuse slightly the notation by setting $Q = (Q_0, Q_1, \dots, Q_k)$. Further, set $\tilde{k} = k + 1$ and $\tilde{n} = n + 1$. Then for $\tilde{Z} = Z(\tilde{p}(Q(X)), \mathbb{K}\langle\varepsilon_0\rangle^{\tilde{n}})$ one sees, by

using [5, Prop. 11.47], that the projection to the last n coordinates of any set S meeting every connected component of $\hat{Z}$ meets every connected component of $Z(p(Q(X)), \mathbb{K}\langle\varepsilon_0\rangle^n)$. Assuming one can compute S as a set of univariate representations in $D[\varepsilon_0, T]$, one then can use *restricted elimination* [5, Alg. 12.43] to replace ε_0 by a sufficiently small element of the field of fractions of D to obtain univariate representations of points of Z .

Next, we deform Q by defining $\tilde{Q}(t) : \mathbb{K}^{\tilde{n}} \rightarrow \mathbb{K}[t]^{\tilde{k}}$ as follows

$$(5.1) \quad \tilde{Q}_j(t, X) = Q_j(X) + \frac{t}{2} X^T \text{diag}(1^j, 2^j, \dots, \tilde{n}^j) X, \quad 0 \leq j \leq k.$$

Obviously $\tilde{Q}(t)$ defines, as well, a quadratic map $\mathbb{F}^{\tilde{n}} \rightarrow \mathbb{F}^{\tilde{k}}$ for any field $\mathbb{F} \supseteq \mathbb{K}\langle t \rangle$. The following lemma states that the Hessians of the $\tilde{Q}_j(t)$'s are in general position, in sense that their nonzero linear combinations (that will be the matrices $A(Y)$ mentioned in the discussion in the beginning of this section) are of maximal possible rank. The statement is similar to [2, (3.6)].

LEMMA 5.2. *Let the matrix $A(Y, T)$ with the entries in $\mathbb{K}[Y, T]$ be defined by*

$$A(Y, T) = \sum_{j=1}^k Y_j (H_j + T \text{diag}(1^{j-1}, 2^{j-1}, \dots, n^{j-1})).$$

Let t be transcendental over $\overline{\mathbb{K}}$. Then for any field $\mathbb{F} \supseteq \mathbb{K}\langle t \rangle$ and $0 \neq y \in \overline{\mathbb{F}}^k$, the rank of $A(y, t)$ is at least $n - k + 1$.

There exists $0 < \iota' \in \mathbb{K}$ such that for any $\iota \in \mathbb{K}$ satisfying $0 < \iota < \iota'$ and $0 \neq y \in \overline{\mathbb{K}}$ the rank of $A(y, \iota)$ is at least $n - k + 1$.

PROOF. Consider

$$B = B(Y, T, \mu) = \sum_{j=1}^k Y_j (\mu H_j + T \text{diag}(1^{j-1}, 2^{j-1}, \dots, n^{j-1}))$$

and the homogeneous, with respect to Y , as well as with respect to $\{\mu, T\}$, ideal $J = (\det B_{UW} \mid U, W \subset \{1, \dots, n\}, |U| = |W| = n - k)$ in the ring $D[Y, T, \mu] \subset \overline{\mathbb{K}}[Y, T, \mu]$.

Note that every $(y^*, t^*) \neq 0$ satisfying $\text{rk}(A(y^*, t^*)) < n - k + 1$ gives rise to $0 \neq (y^*, t^*, 1) \in Z(J) = Z(J, \overline{\mathbb{K}}^{k+2})$. Vice versa, $0 \neq (y^*, t^*, 1) \in Z(J)$ obviously implies $\text{rk}(A(y^*, z^*)) < n - k + 1$.

The idea is now to show that there exists a nonzero polynomial $f(T) \in \overline{\mathbb{K}}[T]$ such that $f(t) = 0$ for all $(y^*, t) \neq 0$ with $\text{rk}(A(y^*, t)) < n - k + 1$. As t is

transcendental over $\mathbb{K}$, that is, it cannot be a root of a polynomial in $\overline{\mathbb{K}}[T]$, this will imply the statement of the lemma.

The ideal $J' = (J : (Y)^\infty) \cap \overline{\mathbb{K}}[T, \mu]$, obtained by “projectively” eliminating Y from J , is homogeneous. By the “Main Theorem of Elimination Theory”, see e.g. [7, Theorem 14.1], the image of $Z(J)$ under the corresponding projection is Zariski-closed. Hence it coincides with $Z(J') = Z(J', \overline{\mathbb{K}}^2)$. Moreover, J' is nonempty, as $Z(J)$, and hence $Z(J')$, do not contain elements with $\mu = 0$ and $T = 1$, as follows immediately from the properties of the diagonal matrices $\text{diag}(1^j, 2^j, \dots, n^j)$. Thus J' contains a homogeneous polynomial $f(T, \mu) \in \overline{\mathbb{K}}[T, \mu]$ that is not divisible by μ .

Hence any t for which there exists $y^* \neq 0$ satisfying $\text{rk}(A(y^*, t)) < n - k + 1$, satisfies $f(t, 1) = 0$, that is, t is algebraic over $\overline{\mathbb{K}}$, a contradiction showing the first part of the lemma. To obtain the second part, observe that $f(T, 1)$ vanishes only on finitely many elements of $\mathbb{K}$. Choose ι' to be the closest to 0 root of $f(T, 1)$ among the positive elements of $\mathbb{K}$, if such a root exists. Otherwise choose $\iota' = 1$. $\square$

Next, we shall perturb $\tilde{p}(\tilde{Q}(X))$ so that 0 is not its (and neither that of $\tilde{p}(Y)$) critical value by subtracting an appropriate constant τ from it. (Such τ is called a *regular value* of $\tilde{p}(\tilde{Q}(X))$ and of $\tilde{p}(Y)$). We will talk about the τ -level set of $\tilde{p}(\tilde{Q}(X))$, that is just $Z(\tilde{p}(\tilde{Q}(X)) - \tau, \mathbb{K}^n)$. The following is an immediate consequence of the semialgebraic Sard’s theorem [6, Thm. 9.6.2], [5, Thm. 5.57].

LEMMA 5.3. *Let $\mathbb{F}$ be a real closed field and τ a transcendental infinitesimal over $\mathbb{F}$ (respectively, $\tau > 0$ a sufficiently close to 0 element of $\mathbb{F}$). Then τ (and any ι satisfying $0 < \iota < \tau$) is a regular value of any nonzero $f(Y) \in \mathbb{F}[Y]$. In particular, provided that $\mathbb{F}$ contains the field generated by the coefficients of p and Q , one has that τ (and any ι as above) is a regular value of $\tilde{p}(\tilde{Q}(X))$ and of $\tilde{p}(Y)$. $\square$*

Let $\varepsilon_0 \gg \varepsilon_1 \gg \varepsilon_2 > 0$ be two more extra infinitesimals over $\mathbb{K}$, and denote $\tilde{Q} = \tilde{Q}(\varepsilon_2)$. We deform $\tilde{Z}$ as follows:

$$\tilde{Z} = Z(\tilde{p}(\tilde{Q}(X)) - \varepsilon_1, \mathbb{K}\langle \varepsilon_0, \varepsilon_1, \varepsilon_2 \rangle^{\tilde{n}}).$$

At this point we are ready to use the tool from Section 4, where it is described in slightly greater generality. According to Theorem 4.5 we have a covering of the set V_c of the critical points of $X \mapsto X_0$ on $\tilde{Z}$ by $n^{O(k)}$ semialgebraic sets $V_c(U, W)$. Moreover, Theorem 4.5 gives us for each $V_c(U, W)$ an isomorphism ϕ_W (given by polynomials in $D[X]$ of degree at most $2d$) so that

$\phi_W(V_c(U, W)) \subseteq \mathbb{K}\langle \varepsilon_0, \varepsilon_1, \varepsilon_2 \rangle^{O(k)}$, as well as its inverse ϕ_{UW}^{-1} (given by rational functions, with common denominator, with coefficients in $D[\varepsilon]$, of degrees at most $\tilde{n}$). By Proposition 4.8, this data can be computed by $(n(d + d'))^{O(k)}$ arithmetic operations in D .

The sets $\phi_W(V_c(U, W))$ and $V_c(U, W)$ are both defined by equations and one inequation $\Lambda \neq 0$, with $\Lambda \in D[\varepsilon][Y]$ (respectively, $\Lambda \in D[\varepsilon][X]$). By adding one extra variable as in the beginning of Section 3 we convert each of them into a real algebraic set: add equation $Y_{\tilde{k}+1}\Lambda = 1$ (respectively, $X_{\tilde{n}+1}\Lambda = 1$) and extend the maps ϕ_W and ϕ_{UW}^{-1} by “ignoring” these extra variables. Apply to $Z_F := \phi_W(V_c(U, W))$ and $\Psi := \phi_{UW}^{-1}$ the procedure of Theorem 1.10. It will produce a set $R(U, W)$ of univariate representations of points intersecting each connected component of $\lim_\varepsilon V_c(U, W)$.

By the following lemma, the union of the $R(U, W)$ ’s over U, W will intersect each connected component C of Z , as by Proposition 3.2 one has $C = \lim_{\varepsilon_1, \varepsilon_2} C_\varepsilon$, where C_ε is a connected component of $\tilde{Z}$, and C_ε intersects some $V_c(U, W)$.

LEMMA 5.4. $\hat{Z} = \lim_{\varepsilon_1, \varepsilon_2} \tilde{Z}$.

PROOF. Denote $\varepsilon = (\varepsilon_2, \varepsilon_1)$. As $\lim_\varepsilon$ is a ring homomorphism from $\mathbb{K}\langle \varepsilon \rangle_b$ to $\mathbb{K}$, certainly $\lim_\varepsilon \tilde{Z} \subseteq \hat{Z}$. We shall show the reverse inclusion. Let $x \in \hat{Z}$. We find a point $\tilde{x} \in \tilde{Z}$ satisfying $\lim_\varepsilon \tilde{x} = x$. Note that $\tilde{p}(\tilde{Q}(x)) \in \varepsilon_2 \mathbb{K}$ and $\tilde{p}(\tilde{Q}(x)) - \varepsilon_1 < 0$, as $\varepsilon_1 \gg \varepsilon_2 > 0$. On the other hand, as $\tilde{p}(Q(X))$ is not identically 0, for any $0 < r \in \mathbb{K}$ the ball of radius r around x in $\mathbb{K}^{\tilde{n}}$ contains a point y such that $\tilde{p}(Q(y)) > 0$. As x lies in the closure of the semialgebraic set F_+ defined by $\tilde{p}(Q(X)) > 0$, there exists a semialgebraic path $\gamma : [0, 1] \rightarrow \mathbb{K}^{\tilde{n}}$ such that $\gamma(0) = x$ and $\gamma((0, 1]) \subseteq F_+$, cf. Curve selection lemma [5, Thm. 3.19]. As the image of a closed and bounded semialgebraic set under a continuous semialgebraic function on it is bounded, cf. [5, Thm. 3.20], $\gamma([0, 1])$ is bounded over $\mathbb{K}$.

Let $\bar{\gamma}$ denote the extension of γ to $\mathbb{K}\langle \varepsilon \rangle$. Then by [5, Prop. 2.84] the set $\bar{\gamma}([0, 1])$ is bounded over $\mathbb{K}\langle \varepsilon \rangle$. By the semialgebraic intermediate value theorem [5, Prop. 3.4] the set $\mathcal{I}(\tau_0)$ of all $\tau \in (0, \tau_0) \subset \mathbb{K}\langle \varepsilon \rangle$ satisfying $\tilde{p}(\tilde{Q}(\bar{\gamma}(\tau))) = 0$ is a nonempty closed semialgebraic set. Choose τ in the closest to 0 interval of $\mathcal{I}(\tau_0)$. Then $\lim_\varepsilon \tau = 0$, as τ_0 is arbitrary close to 0. As $\lim_\varepsilon$ is a ring homomorphism, we have $\tilde{p}(Q(\lim_\varepsilon \bar{\gamma}(\tau))) = 0$.

It remains to show that $x = \lim_\varepsilon \bar{\gamma}(\tau)$. Identify τ with the corresponding (representative of the) germ of semialgebraic continuous functions on $\mathbb{K}_{>0}$ and think of $\bar{\gamma}(\tau)$ as of composition $\gamma \circ \tau$. To complete the proof, apply [5, Lemma 3.21] that states that in this setting (γ a semialgebraic continuous

function on a closed bounded semialgebraic set over $\mathbb{K}$ and τ an element of the extension of this set to $\mathbb{K}\langle\varepsilon\rangle$) one has $\gamma(\lim_\varepsilon \tau) = \lim_\varepsilon(\gamma \circ \tau)$. $\square$

At this point we obtained a set of univariate representations $u(\varepsilon_0, T) \in D[\varepsilon_0, T]^{n+3}$ (see (1.1)) for points in each connected component of $\hat{Z}$. Now we get rid of the infinitesimal ε_0 . Remove from u the polynomial g_1 responsible for X_0 -coordinate, that is no longer needed, and apply [5, Alg. 12.46] (Removal of Infinitesimals), that consists of two steps.

The first step is running the restricted elimination algorithm [5, Alg. 12.43] with the input consisting of the polynomial f and the following polynomials:

$$(5.5) \quad f', f'', \dots, f^{(\deg(f)-1)}, \quad g_0^{\deg(p(Q))} p(Q(\frac{g_2}{g_0}, \dots, \frac{g_{n+1}}{g_0})).$$

It outputs a finite set $\mathcal{S} \subset D[\varepsilon_0]$ such that the degree of f , the number of roots of f in $\mathbb{K}$, the number of common roots of f and h in $\mathbb{K}$, and the signs of h at the roots of f , for all h in (5.5), are fixed on each connected component of the realization of any (realizable) sign condition on $\mathcal{S}$.

The second step computes, for each polynomial

$$h(\varepsilon_0) = h_\ell \varepsilon_0^\ell + h_{\ell+1} \varepsilon_0^{\ell+1} + \dots + h_{\ell+\omega} \varepsilon_0^{\ell+\omega} \in \mathcal{S}, \quad h_\ell \neq 0,$$

the Cauchy lower bound $\frac{|h_\ell|}{\sum_m |h_m|}$ on the absolute value of its *nonzero* roots (see [5, Lemma 10.3]) and substitutes the minimum $\frac{a}{b}$, for $a, b \in D$, of these bounds for ε_0 . The following remain unchanged upon substituting $\varepsilon_0 = \frac{a}{b}$:

- the number of real roots of $f(\frac{a}{b}, T)$ and their Thom encodings;
- the signs of the polynomials in (5.5) at these roots. In particular the Thom encodings of these roots will remain the same.

By construction, the set of points represented by the $u(\frac{a}{b}, T)$'s intersects each connected component of Z .

The number of arithmetic operations in D for these two steps is $(dn)^{O(k)}$, according to [5, p. 462].

It remains to convert the $u(\frac{a}{b}, T) = (f, g_0, g_2, \dots, g_{n+1})$'s into real univariate representations by computing the Thom encodings σ for each real root of f using [5, Alg. 10.64] and [5, Rem. 10.66]. The complexity of this procedure is $(dn)^{O(k)}$, and for the case $D = \mathbb{Z}$ the bitsizes of the intermediate data are bounded by $(dn)^{O(k)}$ times the bitsize of the input, by [loc.cit.].

This completes the proof of Theorem 1.2.

Acknowledgements

The second author is supported by NWO Grant 613.000.214. Parts of the work were completed while he held a position at THI/FB20 Informatik, University of Frankfurt (Germany) supported by DFG Grant SCHN-503/2-1. As well, he was partially supported by Université de Rennes I (France) and the Mathematical Sciences Research Institute (USA).

The authors thank Saugata Basu, Marie-Françoise Roy, Nikolaj Vorobjov, and two anonymous referees for useful remarks, and Michael Kettner for carefully reading a preliminary version of the text and tireless help in improving it. Sergey Fomin kindly provided references in Remark 2.15.

References

- [1] M.-E. Alonso, E. Becker, M.-F. Roy, and T. Wörmann. Zeros, multiplicities, and idempotents for zero-dimensional systems. In *Algorithms in algebraic geometry and applications (Santander, 1994)*, volume 143 of *Progr. Math.*, pages 1–15. Birkhäuser, Basel, 1996.
- [2] A. I. Barvinok. Feasibility testing for systems of real quadratic equations. *Discrete Comput. Geom.*, 10(1):1–13, 1993.
- [3] A. I. Barvinok. On the Betti numbers of semialgebraic sets defined by few quadratic inequalities. *Math. Z.*, 225(2):231–244, 1997.
- [4] S. Basu, R. Pollack, and M.-F. Roy. On the combinatorial and algebraic complexity of quantifier elimination. *J. ACM*, 43(6):1002–1045, 1996.
- [5] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*. Springer-Verlag, 2003.
- [6] J. Bochnak, M. Coste, and M.-F. Roy. *Real algebraic geometry*. Springer-Verlag, Berlin, 1998. Translated from the 1987 French original, Revised by the authors.
- [7] D. Eisenbud. *Commutative Algebra with a View Toward Algebraic Geometry*. Springer-Verlag, New York, 1995.
- [8] D. Grigoriev. Complexity of deciding Tarski algebra. *J. Symbolic Comput.*, 5(1-2):65–108, 1988.

- [9] D. Grigoriev and N. Vorobjov. Complexity lower bounds for computation trees with elementary transcendental function gates. *Theoret. Comput. Sci.*, 157(2):185–214, 1996.
- [10] D. Grigoriev and N. N. Vorobjov, Jr. Solving systems of polynomial inequalities in subexponential time. *J. Symbolic Comput.*, 5(1-2):37–64, 1988.
- [11] L. E. Heindel. Integer arithmetic algorithms for polynomial real zero determination. *J. Assoc. Comput. Mach.*, 18:533–548, 1971.
- [12] U. J. J. Le Verrier. Sur les variations séculaires des éléments elliptiques des sept planètes principales: Mercure, vénus, la terre, mars, jupiter, saturne et uranus. *J. Math. Pures Appl.*, 5:220–254, 1840.
- [13] I. G. Macdonald. *Symmetric functions and Hall polynomials*. Oxford Mathematical Monographs. The Clarendon Press, Oxford University Press, New York, second edition, 1995. With contributions by A. Zelevinsky, Oxford Science Publications.
- [14] E. W. Mayr and A. R. Meyer. The complexity of the word problems for commutative semigroups and polynomial ideals. *Adv. in Math.*, 46(3):305–329, 1982.
- [15] J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. I-III. *J. Symbolic Comput.*, 13(3):255–352, 1992.
- [16] F. Rouillier, M.-F. Roy, and M. Safey El Din. Finding at least one point in each connected component of a real algebraic set defined by a single equation. *J. Complexity*, 16:716–750, 2000.
- [17] J. von zur Gathen and J. Gerhard. *Modern computer algebra*. Cambridge University Press, New York, 1999.

DIMA GRIGORIEV
IRMAR
Université de Rennes I
Campus de Beaulieu
35042 Rennes cedex
France
dima@math.univ-rennes1.fr
<http://name.math.univ-rennes1.fr/dimitri.grigoriev/>

DMITRII V. PASECHNIK
Dept. E & OR and CentER
Tilburg University
P.O. Box 90153
5000 LE Tilburg
The Netherlands
d.v.pasechnik@uvt.nl
<http://center.uvt.nl/staff/pasechnik/>