Faster Real Feasibility via Circuit Discriminants

Frédéric Bihan UFR SFA,
Campus Scientifique
73376 Le Bourget-du-Lac
Cedex
France
Frederic.Bihan@univ-savoie.fr

J. Maurice Rojas*
TAMU 3368
Department of Mathematics
Texas A&M University
College Station, Texas
77843-3368
USA
rojas@math.tamu.edu

Casey E. Stella†
18409 Newell Road
Shaker Heights, OH 44122
USA
cestella@gmail.com

Rojas dedicates this paper to the memory of his dear friend, Richard Adolph Snavely, May 22, 1955 – December 21, 2005.

ABSTRACT

We show that detecting real roots for honestly n-variate $(n+2)$-nomials (with integer exponents and coefficients) can be done in time polynomial in the **sparse** encoding for any fixed n. The best previous complexity bounds were exponential in the sparse encoding, even for n fixed. We then give a characterization of those functions $k(n)$ such that the complexity of detecting real roots for n-variate $(n+k(n))$-nomials transitions from **P** to **NP**-hardness as $n \longrightarrow \infty$. Our proofs follow in large part from a new complexity threshold for deciding the vanishing of $\mathcal{A}$-discriminants of n-variate $(n+k(n))$-nomials. Diophantine approximation, through linear forms in logarithms, also arises as a key tool.

Keywords

sparse, real, feasibility, polynomial-time, discriminant chamber, linear forms in logarithms

1. INTRODUCTION AND MAIN RESULTS

Consider real feasibility: the problem of deciding the existence of real roots for systems of polynomial equations. In addition to having numerous practical applications (see, e.g., [BG-V03]), real feasibility is an important motivation behind effectivity estimates for the Real Nullstellensatz (e.g., [Ste74, Sch00]), the quantitative study of sums of squares [Ble04, RS09, BHPR09], and their connection to semi-definite programming and optimization [Par03, Las07]. In particular, real solving of sparse polynomial systems arises in concrete applications such as satellite orbit mechanics [AM09], and real solving clearly involves real feasibility as an initial step. We are thus inspired to derive new algorithms and complexity lower bounds for real feasibility, in the refined setting of sparse polynomials.

To state our results, let us first clarify some basic notation concerning sparse polynomials and some well-known complexity classes. Recall that R^* is the multiplicative group of nonzero elements in any ring R.

DEFINITION 1.1. *When $a_j \in \mathbb{R}^n$, the notations $a_j = (a_{1,j}, \ldots, a_{n,j})$, $x^{a_j} = x_1^{a_{1,j}} \cdots x_n^{a_{n,j}}$, and $x = (x_1, \ldots, x_n)$ will be understood. If $f(x) := \sum_{j=1}^m c_i x^{a_j}$ where $c_j \in \mathbb{R}^*$ for all j, and the a_j are pair-wise distinct, then we call f a* **(real) n-variate m-nomial***, and we define* $\mathrm{Supp}(f) := \{a_1, \ldots, a_m\}$ *to be the* **support** *of f. We also let $\mathcal{F}_{n,m}$ denote the set of all n-variate $\lfloor m \rfloor$-nomials within $\mathbb{Z}[x_1, \ldots, x_n]$. Finally, for any $m \geq n+1$, we let $\mathcal{F}^*_{n,m} \subseteq \mathcal{F}_{n,m}$ denote the subset consisting of those f with* $\mathrm{Supp}(f)$ **not** *contained in any $(n-1)$-flat. We also call any $f \in \mathcal{F}^*_{n,m}$ an* **honest n-variate m-nomial** *(or* **honestly n-variate***).* $\diamond$

For example, $1+7x_1^2x_2x_3^7x_4^3-43x_1^{198}x_2^{99}x_3^{693}x_4^{297}$ is a 4-variate trinomial with support contained in a line segment, but it has a real root $x \in \mathbb{R}^4$ iff the honestly univariate trinomial $1 + 7y_1 - 43y_1^{99}$ has a real root $y_1 \in \mathbb{R}$. More generally, via a monomial change of variables, it will be natural to restrict to $\mathcal{F}^*_{n,n+k}$ (with $k \geq 1$) to study the role of sparsity in algorithmic complexity over the real numbers.

We will work with some well-known complexity classes from the classical Turing model of computation (see, e.g., [Pap95].) In particular, our underlying notion of input size is clarified in Definition 2.1 of Section 2.1 below, and illustrated in Example 1.4, immediately following our first main theorem. So for now, let us just recall the basic inclusions $\mathbf{NC}^1 \subseteq \mathbf{P} \subseteq \mathbf{NP} \subseteq \mathbf{PSPACE}$. While it is known that $\mathbf{NC}^1 \neq \mathbf{PSPACE}$ the properness of each of the remaining inclusions above is a famous open problem.

1.1 Sparse Real Feasibility and $\mathcal{A}$-Discriminant Complexity

DEFINITION 1.2. *Let $\mathbb{R}_+$ denote the positive real numbers and let $\mathbf{FEAS}_{\mathbb{R}}$ (resp. $\mathbf{FEAS}_+$) denote the problem of deciding whether an arbitrary system of equations from $\bigcup_{n\in\mathbb{N}} \mathbb{Z}[x_1, \ldots, x_n]$ has a real root (resp. a root with all coordinates positive). Also, for any collection $\mathcal{F}$ of tuples chosen from $\bigcup_{k,n\in\mathbb{N}} (\mathbb{Z}[x_1, \ldots, x_n])^k$, we let $\mathbf{FEAS}_{\mathbb{R}}(\mathcal{F})$ (resp.*

*Partially supported by NSF individual grant DMS-0211458, NSF CAREER grant DMS-0349309, AIM, Sandia National Laboratories, and MSRI.

†Partially supported by NSF grant DMS-0211458.

ISSAC '09 Seoul, Korea

$\mathbf{FEAS}_+(\mathcal{F})$) *denote the natural restriction of* $\mathbf{FEAS}_{\mathbb{R}}$ *(resp.* $\mathbf{FEAS}_+$*) to inputs in* $\mathcal{F}$. $\diamond$

It has been known since the 1980s that $\mathbf{FEAS}_{\mathbb{R}} \in \mathbf{PSPACE}$ [Can88], and an **NP**-hardness lower bound was certainly known earlier. However, no sharper bounds in terms of sparsity were known earlier in the Turing model until our first main theorem.

THEOREM 1.3. *Let* $Z_+(f)$ *denote the zero set of* f *in* $\mathbb{R}^n_+$. *Then:*

0. $\mathbf{FEAS}_+\left(\bigcup_{n\in\mathbb{N}} \mathcal{F}^*_{n,n+1}\right)$ *and* $\mathbf{FEAS}_{\mathbb{R}}\left(\bigcup_{n\in\mathbb{N}} \mathcal{F}^*_{n,n+1}\right)$ *are in* $\mathbf{NC}^1$. *In particular, when* $f \in \mathcal{F}^*_{n,n+1}$, $Z_+(f)$ *is either empty or diffeotopic*[1] *to* $\mathbb{R}^{n-1}_+$, *with each case actually occuring.*

1. *For any fixed* n, $\mathbf{FEAS}_+(\mathcal{F}^*_{n,n+2})$ *and* $\mathbf{FEAS}_{\mathbb{R}}(\mathcal{F}^*_{n,n+2})$ *are in* $\mathbf{P}$.

2. *For any fixed* $\varepsilon > 0$, *both* $\mathbf{FEAS}_+\left(\bigcup_{n\in\mathbb{N}\ ,\ 0<\varepsilon'\leq\varepsilon} \mathcal{F}^*_{n,n+n^{\varepsilon'}}\right)$ *and* $\mathbf{FEAS}_{\mathbb{R}}\left(\bigcup_{n\in\mathbb{N}\ ,\ 0<\varepsilon'\leq\varepsilon} \mathcal{F}^*_{n,n+n^{\varepsilon'}}\right)$ *are* **NP***-hard.*

The recent paper [PRT09] proves, in the context of optimizing n-variate $(n+2)$-nomials, that sharper algorithmic complexity bounds hold when we instead work in the BSS model over $\mathbb{R}$ (thus counting arithmetic operations instead of bit operations). The latter paper also details how the results here can be extended to real exponents.

EXAMPLE 1.4. *A very special case of Assertion (1) of Theorem 1.3 implies that one can decide — for any nonzero* $c_1, \ldots, c_5 \in \mathbb{Z}$ *and* $D \in \mathbb{N}$ *— whether*

$$c_1 + c_2 x_1^{999} + c_3 x_1^{73} x_3^{19} + c_4 x_2^{27D} + c_5 x_1^{74} x_2^D x_3$$

has a root in $\mathbb{R}^3$, *using a number of bit operations polynomial in*

$$\log(D) + \log\left[(|c_1|+1)\cdots(|c_5|+1)\right].$$

The best previous results (e.g., via the critical points method, infinitesimals, and rational univariate reduction, as detailed in [BPR06]) would yield a bound polynomial in $D + \log\left[(|c_1|+1)\cdots(|c_5|+1)\right]$ *instead.* $\diamond$

We thus see that for sparse polynomials, large degree can be far less of a complexity bottleneck over $\mathbb{R}$ than over $\mathbb{C}$. Theorem 1.3 is proved in Section 3.2 below. The underlying techniques include $\mathcal{A}$-discriminants (a.k.a. sparse discriminants) (cf. Section 2.3), Viro's Theorem from toric geometry (see, e.g., [GKZ94, Thm. 5.6]), and effective estimates on linear forms in logarithms [Bak77, Nes03].

In particular, for any collection $\mathcal{F}_{\mathcal{A}}$ of n-variate m-nomials with support $\mathcal{A}$, there is a polynomial $\Delta_{\mathcal{A}}$ in the coefficients (c_i) called the ***A*-discriminant**. Its real zero set partitions $\mathcal{F}_{\mathcal{A}}$ into **chambers** (connected components of the complement) on which the zero set of an $f \in \mathcal{F}_{\mathcal{A}}$ has constant topological type. A toric deformation argument employing Viro's Theorem enables us to decide whether a given chamber consists of f having empty or non-empty $Z_+(f)$. For any $\mathcal{A} \subset \mathbb{Z}^n$ of cardinality $n+2$ (in sufficiently general position), there is then a compact formula for the $\mathcal{A}$-discriminant that enables us to pick out which chamber contains a given f: one simply computes the sign of a linear combination of logarithms. Our resulting algorithms are thus quite implementable, requiring only fast approximation of logarithms and some basic triangulation combinatorics for $\mathrm{Supp}(f)$. (A preliminary `Matlab` implementation can be downloaded from `www.math.tamu.edu/~rojas/cktposfeas.m` .)

[1] See Definition 2.9 of Section 2.3 below.

EXAMPLE 1.5. *Consider* $\mathcal{A} := \{(0,0,0), (999,0,0), (73,0,19), (0,2009,0), (74,293,1)\}$, *which gives us the family of trivariate pentanomials*

$$\mathcal{F}_{\mathcal{A}} := \left\{ c_1 + c_2 x_1^{999} + c_3 x_1^{73} x_3^{19} + c_4 x_2^{2009} + c_5 x_1^{74} x_2^{293} x_3 \;\middle|\; c_i \in \mathbb{R}^* \right\}.$$

Suppose further that $f \in \mathcal{F}_{\mathcal{A}}$ *is an element satisfying* $c_1, c_2, c_3, c_4 > 0$ *and* $c_5 < 0$. *It then turns out via Lemma 2.12 (cf. Section 2.4 below) that* $Z_+(f)$ *has a degeneracy iff the* ***A*-discriminant**, $\Delta_{\mathcal{A}}(c) :=$

$$38132829^{38132829} c_1^{27886408} c_2^{2677997} c_3^{2006991} c_4^{5561433} - 27886408^{27886408} 2677997^{2677997} 2006991^{2006991} 5561433^{5561433} c_5^{38132829}$$

vanishes. In fact, via the techniques underlying Theorem 1.3, $Z_+(f)$ *is either empty, a point, or isotopic to a* 2*-sphere, according as* $\Delta_{\mathcal{A}}(c)$ *is positive, zero, or negative. Note in particular that determining the sign of* $\Delta_{\mathcal{A}}(c)$ *is equivalent to determining the sign of*

$$38132829\log(38132829) + 27886408\log(c_1) + 2677997\log(c_2) + 2006991\log(c_3) + 5561433\log(c_4)$$
$$-27886408\log(27886408) - 2677997\log(2677997) - 2006991\log(2006991) - 5561433\log(5561433) - 38132829\log(c_5).$$ $\diamond$

While we review $\mathcal{A}$-discriminants in Section 2.3 below, it is important to observe now how the computational complexity of $\mathcal{A}$-discriminants closely parallels that of $\mathbf{FEAS}_{\mathbb{R}}$: compare Theorem 1.3 above with Theorem 1.7 below.

DEFINITION 1.6. *Let* $\mathbf{ADISC}_=$ *(resp.* $\mathbf{ADISC}_>$*) denote the problem of deciding whether* $\Delta_{\mathcal{A}}(f)$ *vanishes (resp. determining the sign of* $\Delta_{\mathcal{A}}(f)$*) for an input polynomial* f *with integer coefficients, where* $\mathcal{A} = \mathrm{Supp}(f)$. *Finally, let* $\mathbf{ADISC}_=(\mathcal{F})$ *(resp.* $\mathbf{ADISC}_>(\mathcal{F})$*) be the natural restriction of* $\mathbf{ADISC}_=$ *(resp.* $\mathbf{ADISC}_>$*) to inputs in some family* $\mathcal{F}$. $\diamond$

THEOREM 1.7.

1. $\mathbf{ADISC}_=\left(\bigcup_{n\in\mathbb{N}} \mathcal{F}^*_{n,n+2}\right) \in \mathbf{P}$ *and, for any* **fixed** n, $\mathbf{ADISC}_>(\mathcal{F}^*_{n,n+2}) \in \mathbf{P}$.

2. *For any fixed* $\varepsilon > 0$, *both* $\mathbf{ADISC}_=\left(\bigcup_{\substack{n\in\mathbb{N}\\ 0<\varepsilon'\leq\varepsilon}} \mathcal{F}^*_{n,n+n^{\varepsilon'}}\right)$ *and* $\mathbf{ADISC}_>\left(\bigcup_{\substack{n\in\mathbb{N}\\ 0<\varepsilon'\leq\varepsilon}} \mathcal{F}^*_{n,n+n^{\varepsilon'}}\right)$ *are* **NP***-hard.*

Theorem 1.7 is proved in Section 3.1, after the development of some necessary theory in Section 2 below.

1.2 Related Work

Earlier work on algorithmic fewnomial theory has mainly gone in directions other than polynomial-time algorithms. For example, Gabrielov and Vorobjov have given singly exponential time algorithms for weak stratifications of semi-Pfaffian sets [GV04] — data from which one can compute homology groups of real zero sets of a class of functions more general than sparse polynomials. Our approach thus highlights a subproblem where faster and simpler algorithms are possible.

Focussing on feasibility, other than the elementary results $\mathbf{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,1}) \in \mathbf{NC}^0$ and $\mathbf{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,2}) \in \mathbf{NC}^0$, there appear to have been no earlier complexity upper bounds of the form $\mathbf{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,m}) \in \mathbf{P}$, or even $\mathbf{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,m}) \in \mathbf{NP}$, for $m \geq 3$. (With the exception of [RY05], algorithmic work on univariate real polynomials has focussed on algorithms that

are quasi-linear in the degree. See, e.g., [LM01].) Echoing the parallels between $\mathbf{FEAS}_{\mathbb{R}}$ and $\mathbf{ADISC}_{>}$ provided by Theorems 1.3 and 1.7, both $\mathbf{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,4}) \stackrel{?}{\in} \mathbf{P}$ and $\mathbf{ADISC}_{>}(\mathcal{F}_{1,4}) \stackrel{?}{\in} \mathbf{P}$ are open problems.

As for earlier complexity lower bounds for $\mathbf{FEAS}_{\mathbb{R}}$ in terms of sparsity, we are unaware of any. Indeed, it is not even known whether $\mathbf{FEAS}_{\mathbb{R}}(\mathbb{Z}[x_1, \ldots, x_n])$ is **NP**-hard for some fixed n. Also, complexity lower bounds for the vanishing of discriminants of n-variate $(n + k(n))$-nomials (with k a slowly growing function of n) appear to be new. However, recent work shows that the geometry of discriminants chambers can be quite intricate already for $f \in \mathcal{F}^*_{3,3+3}$ [DRRS07]. Also, it was known even earlier that deciding the vanishing of sparse discriminants of **uni**variate m-nomials (with m unbounded) is already **NP**-hard with respect to randomized reductions [KS99]. Considering Theorems 1.3 and 1.7, one may thus be inclined to conjecture that $\mathbf{FEAS}_{\mathbb{R}}(\mathbb{Z}[x_1])$ is **NP**-hard. Curiously, over a different family of complete fields (the ***p*-adic rationals**), one can already prove that detecting roots for univariate m-nomials (with m unbounded) is **NP**-hard, with respect to randomized reductions [IRR07]. Relative to the more efficient **SLP-encoding**, Peter Bürgisser found a short and elegant proof that $\mathbf{FEAS}_{\mathbb{R}}(\mathbb{Z}[x_1])$ is **NP**-hard (see [Per08] for an alternative proof).

2. BACKGROUND AND ANCILLARY RESULTS

After recalling a basic complexity construction, we will present some tools for dealing with n-variate $(n+1)$-nomials, and then move on to n-variate $(n + k)$-nomials with $k \geq 2$.

2.1 A Key Reduction

To measure the complexity of our algorithms, let us fix the following definitions for input size.

DEFINITION 2.1. *For any $a \in \mathbb{Z}$, we define its size,* $\mathrm{size}(a)$*, to be* $1 + \log(1 + |a|)$*. More generally, we define the* **size** *of a matrix $U = [u_{i,j}] \in \mathbb{Z}^{m \times n}$ to be $\sum_{i,j} \mathrm{size}(u_{i,j})$. Also, for any $f(x) = \sum_{i=1}^{m} c_i x^{a_i} \in \mathbb{Z}[x_1, \ldots, x_n]$, we define* $\mathrm{size}(f)$ *to be $\sum_{i=1}^{m} [\mathrm{size}(c_i) + \mathrm{size}(a_i)]$. Finally, for $F = (f_1, \ldots, f_k) \in (\mathbb{Z}[x_1, \ldots, x_n])^k$, we define $\mathrm{size}(F) = \sum_{i=1}^{k} \mathrm{size}(f_i)$.* $\diamond$

A key construction we will use later in our **NP**-hardness proofs is a refinement of an old trick for embedding Boolean satisfiability (**3CNFSAT** specifically [Pap95]) into real/complex satisfiability.

PROPOSITION 2.2. *Given any* **3CNFSAT** *instance $B(X)$ with n variables, N clauses, and $N \geq 1 + \frac{n}{4}$, let W_B denote $((\{1\} \times \mathbb{P}^1_{\mathbb{C}}) \cup (\mathbb{P}^1_{\mathbb{C}} \times \{1\}))^{4N-n}$. Then there is an $(8N - n) \times (8N - n)$ polynomial system F_B with the following properties:*

1. *$B(X)$ is satisfiable iff F_B has a root in $\{1, 2\}^n \times W_B$.*
2. *F_B has no more than $33N - 4n$ monomial terms, $\mathrm{size}(F_B) = O(N)$, and every root of F_B in $(\mathbb{P}^1_{\mathbb{C}})^{8N-n}$ lies in $\{1, 2\}^n \times W_B$ and is degenerate.*

Also, if we define $t_M(z_1, \ldots, z_M)$ to be $1 + z_1^{M+1} + \cdots + z_M^{M+1} - (M + 1)z_1 \cdots z_M$, then

3. *t_M is nonnegative on $\mathbb{R}^M_+$, with a unique positive root at $(1, \ldots, 1)$ that happens to be the only degenerate root of t_M in $\mathbb{C}^M$.*
4. *If $\varepsilon > 0$, $f \in \mathcal{F}^*_{n,n+k}$, and $M := \left\lceil k^{1/\varepsilon} \right\rceil$, then $f(x) + t_M(z) \in \mathcal{F}^*_{\eta,\eta+\eta^\delta}$ for $\eta = n + M$ and some positive $\delta \leq \varepsilon$. In particular, $\mathrm{size}(f(x) + t_M(z)) = O\left(\mathrm{size}(f)^{1/\varepsilon}\right)$.* ■

The seemingly mysterious polynomial t_M defined above will be useful later when we will need to decrease the difference between the number of terms and variables in certain polynomials.

2.2 Efficient Linear Algebra on Exponents

A simple and useful change of variables is to use monomials in new variables.

DEFINITION 2.3. *For any ring R, let $R^{m \times n}$ denote the set of $m \times n$ matrices with entries in R. For any $M = [m_{ij}] \in \mathbb{R}^{n \times n}$ and $y = (y_1, \ldots, y_n)$, we define the formal expression $y^M := (y_1^{m_{1,1}} \cdots y_n^{m_{n,1}}, \ldots, y_1^{m_{1,n}} \cdots y_n^{m_{n,n}})$. We call the substitution $x := y^M$ a* **monomial change of variables**. *Also, for any $z := (z_1, \ldots, z_n)$, we let $xz := (x_1 z_1, \ldots, x_n z_n)$. Finally, let $\mathbb{GL}_n(\mathbb{Z})$ denote the group of all matrices in $\mathbb{Z}^{n \times n}$ with determinant ± 1 (the set of* **unimodular** *matrices).* $\diamond$

PROPOSITION 2.4. *(See, e.g., [LRW03, Prop. 2].) For any $U, V \in \mathbb{R}^{n \times n}$, we have the formal identity $(xy)^{UV} = (x^U)^V (y^U)^V$. Also, if $\det U \neq 0$, then the function $e_U(x) := x^U$ is an analytic automorphism of $\mathbb{R}^n_+$, and preserves smooth points and singular points of positive zero sets of analytic functions. Moreover, if $\det U > 0$, then e_U in fact induces a diffeotopy on any positive zero set of an analytic function. Finally, $U \in \mathbb{GL}_n(\mathbb{R})$ implies that $e_U^{-1}(\mathbb{R}^n_+) = \mathbb{R}^n_+$ and that e_U maps distinct open orthants of $\mathbb{R}^n$ to distinct open orthants of $\mathbb{R}^n$.* ■

Proposition 2.4, with minor variations, has been observed in many earlier works (see, e.g., [LRW03]). Perhaps the only new ingredient is the observation on diffeotopy, which follows easily from the fact that $\mathbb{GL}^+_n(\mathbb{R})$ (the set of all $n \times n$ real matrices with positive determinant) is a connected Lie group.

Recall that the **affine span** of a point set $\mathcal{A} \subset \mathbb{R}^n$, $\mathrm{Aff}\mathcal{A}$, is the set of real linear combinations $\sum_{a \in \mathcal{A}} c_a a$ satisfying $\sum_{a \in \mathcal{A}} c_a = 0$. Via the now well-studied algorithms for Smith normal form [Sto98], we can easily derive the following facts. (In what follows, we use # for set cardinality and e_i for the $i^{\underline{\text{th}}}$ standard basis vector of $\mathbb{R}^n$.)

LEMMA 2.5. *For any $f \in \mathcal{F}^*_{n,n+1}$ we can compute $\ell \in \{0, \ldots, n\}$ within* $\mathbf{NC}^1$ *and $\gamma \in \mathbb{R}_+$ such that $\bar{f}(x) := \gamma + x_1 + \cdots + x_\ell - x_{\ell+1} - \cdots - x_n$ satisfies: (1) either f or $-f$ has exactly $\ell + 1$ positive coefficients, and (2) $Z_+(\bar{f})$ and $Z_+(f)$ are diffeotopic.* ■

COROLLARY 2.6. *Suppose $f \in \mathcal{F}^*_{n,n+1}$ and $\mathrm{Supp}(f) = \{a_1, \ldots, a_{n+1}\} \subset \mathbb{R}^n$. Then*

1. *f has a root in $\mathbb{R}^n_+ \iff$ not all the coefficients of f have the same sign. In particular, $Z_+(f)$ is diffeotopic to either $\mathbb{R}^{n-1}_+$ or $\emptyset$.*
2. *If all the coefficients of f have the same sign, then f has a root in $(\mathbb{R}^*)^n \iff$ there are indices $i \in [n]$ and $j, j' \in [n + 1]$ with $a_{i,j} - a_{i,j'}$ odd.* ■

2.3 Combinatorics and Topology of Certain $\mathcal{A}$-Discriminants

The connection between topology of discriminant complements and computational complexity dates back to the late 1970s, having been observed relative to (a) the membership problem for semi-algebraic sets [DL79] and (b) the approximation of roots of univariate polynomials [Sma87]. Our goal here is a precise connection between $\mathbf{FEAS}_{\mathbb{R}}$ and $\mathcal{A}$-discriminant complements. (See also [DRRS07] for further results in this direction.)

DEFINITION 2.7. *[GKZ94, Ch. 1 & 9–11] Given any $\mathcal{A}=\{a_1,\ldots,a_m\}\subset\mathbb{Z}^n$ of cardinality m and $c_1,\ldots,c_m\in\mathbb{C}^*$, we define $\nabla_{\mathcal{A}}\subset\mathbb{P}^{m-1}_{\mathbb{C}}$ — the* **$\mathcal{A}$-discriminant variety** *— to be the closure of the set of all $[c_1:\cdots:c_m]\in\mathbb{P}^{m-1}_{\mathbb{C}}$ such that $f(x)=\sum_{i=1}^m c_i x^{a_i}$ has a degenerate root in $\mathbb{C}^n$. We then define $\Delta_{\mathcal{A}}\in\mathbb{Z}[c_1,\ldots,c_m]\setminus\{0\}$ — the* **$\mathcal{A}$-discriminant** *— to be the unique (up to sign) irreducible defining polynomial of $\nabla_{\mathcal{A}}$. Also, when $\nabla_{\mathcal{A}}$ has complex codimension at least 2, we set $\Delta_{\mathcal{A}}$ to the constant 1. For convenience, we will sometimes write $\Delta_{\mathcal{A}}(f)$ in place of $\Delta_{\mathcal{A}}(c_1,\ldots,c_m)$.* $\diamond$

To prove our results, it will actually suffice to deal with a small subclass of $\mathcal{A}$-discriminants.

DEFINITION 2.8. *We call $\mathcal{A}\subset\mathbb{R}^n$ a* **(non-degenerate) circuit**[2] *iff $\mathcal{A}$ is affinely dependent, but every proper subset of $\mathcal{A}$ is affinely independent. Also, we say that $\mathcal{A}$ is a* **degenerate circuit** *iff $\mathcal{A}$ contains a point a and a proper subset $\mathcal{B}$ such that $a\in B$, $\mathcal{A}\setminus a$ is affinely independent, and $\mathcal{B}$ is a non-degenerate circuit.* $\diamond$

For instance, both  and  are circuits, but  is a degenerate circuit. In general, for any degenerate circuit $\mathcal{A}$, the subset $\mathcal{B}$ named above is always unique.

The relevance of $\mathcal{A}$-discriminants to m-nomial zero sets can be summarized as follows.

DEFINITION 2.9. *Following the notation of Definition 2.7, we call any connected component of $\mathbb{P}^{m-1}_{\mathbb{R}}\setminus(\nabla_{\mathcal{A}}\cup\{[x_1:\cdots:x_m]\mid x_1\cdots x_m=0\})$ a* **(real) $\mathcal{A}$-discriminant chamber**. *Also, given any subsets $X,Y\subseteq\mathbb{R}^n_+$, we say that they are* **isotopic** *(resp.* **diffeotopic***) iff there is a continuous (resp. differentiable) function $H:[0,1]\times X\longrightarrow\mathbb{R}^n_+$ such that $H(t,\cdot)$ is a homeomorphism (resp. diffeomorphism) for all $t\in[0,1]$, $H(0,\cdot)$ is the identity on X, and $H(1,X)=Y$. Finally, for any $\mathcal{A}\subset\mathbb{R}^n$ of cardinality m, let $\mathcal{F}_{\mathcal{A}}$ denote the set of all n-variate m-nomials with support $\mathcal{A}$.* $\diamond$

REMARK 2.10. *Note that when $\mathcal{A}$ has cardinality m, we may naturally identify elements of $\mathbb{P}^{m-1}_{\mathbb{C}}$ (resp. $\mathbb{P}^{m-1}_{\mathbb{R}}$) with equivalence classes determined by nonzero complex (resp. real) multiples of elements of $\mathcal{F}_{\mathcal{A}}$.* $\diamond$

The topology of toric real zero sets is known to be constant on discriminant chambers (see, e.g., [GKZ94, Ch. 11, Sec. 5A, Prop. 5.2, pg. 382]). However, we will need a refinement of this fact to positive zero sets: When $\mathcal{A}$ is in sufficiently general position — a non-degenerate circuit, for instance — one can derive the following statement via basic toric geometry (see, e.g., Lemma 15 of [LRW03]).

[2]This terminology comes from matroid theory and has nothing to do with circuits from complexity theory.

LEMMA 2.11. *Following the notation above, suppose $\mathcal{A}\subset\mathbb{R}^n$ is such that the minimum of any linear form on $\mathcal{A}$ is minimized at no more than $n+1$ points. Also let $\mathcal{C}$ be any $\mathcal{A}$-discriminant chamber. Then $f,g\in\mathcal{C}\Longrightarrow Z_+(f)$ and $Z_+(g)$ are diffeotopic.* ■

There is then a very compact description for $\nabla_{\mathcal{A}}$ when $\mathcal{A}$ is a circuit.

LEMMA 2.12. *Suppose $\mathcal{A}=\{a_1,\ldots,a_{n+2}\}\subset\mathbb{Z}^n$ is a non-degenerate circuit, f is a polynomial with support $\mathcal{A}$, $\hat{\mathcal{A}}$ is the $(n+1)\times(n+2)$ matrix whose $j^{\underline{th}}$ column is $\{1\}\times a_j$, $\hat{\mathcal{A}}_j$ is the submatrix of $\hat{\mathcal{A}}$ obtained by deleting the $j^{\underline{th}}$ column, and $b_j:=(-1)^j\det\hat{\mathcal{A}}_j/\beta$ where $\beta=\gcd(\det\hat{\mathcal{A}}_1,\ldots,\det\hat{\mathcal{A}}_{n+2})$. Then:*

1. *$\Delta(c_1,\ldots,c_{n+2})$ is, up to a multiple by a nonzero monomial term, $\prod_{i=1}^{n+2}\left(\frac{c_i}{b_i}\right)^{b_i}-1$. Also, $(b_1,\ldots,b_{n+2})$ can be computed in $\mathbf{P}$.*
2. *For all $[c_1:\cdots:c_{n+2}]\in\mathbb{P}^{n+1}_{\mathbb{R}}$ we have the equivalence $\prod_{i=1}^{n+2}(\mathrm{sign}(b_ic_i)c_i/b_i)^{\mathrm{sign}(b_ic_i)b_i}=1$ for some $[c_1:\cdots:c_{n+2}]\in\mathbb{P}^{n+1}_{\mathbb{R}}$ with $\mathrm{sign}(c_1b_1)=\cdots=\mathrm{sign}(c_{n+2}b_{n+2})\iff Z_+\left(\sum_{i=1}^{n+2}c_ix^{a_i}\right)$ contains a degenerate point ζ. In particular, $Z_+(f)$ has at most one degenerate point.*
3. *$\mathcal{A}$ has exactly two triangulations: one with simplices $\{\mathrm{Conv}(\mathcal{A}\setminus\{b_i\})\mid\mathrm{sign}(b_i)>0\}$, and the other with simplices $\{\mathrm{Conv}(\mathcal{A}\setminus\{b_i\})\mid\mathrm{sign}(b_i)<0\}$. Moreover, the preceding description also holds when $\mathcal{A}$ is a degenerate circuit.*

Proof of Lemma 2.12: With the exception of the assertion on complexity, Lemma 2.12 follows directly from [GKZ94, Prop. 1.8, Pg. 274], [GKZ94, Prop. 1.2, pg. 217], and the discussion following up to the end of Section B on page 218 of [GKZ94]. In particular, the factor β takes into account that $\mathcal{A}$ may not affinely generate $\mathbb{Z}^n$, but is always the integral affine image of an $\mathcal{A}'$ that is. So the sign condition arises simply from a binomial system (with odd determinant) that ζ must satisfy.

The assertion on the complexity of computing $(b_1,\ldots,b_{n+2})$ follows immediately upon consider the Smith factorization and employing Csanky's famous parallel algorithm for the determinant [Csa76]. Indeed, were it not for the gcd computation for β, we could instead assert an $\mathbf{NC}^2$ complexity bound. ■

2.4 Complexity of Circuit Discriminants and Linear Forms in Logarithms

Theorem 1.7 is a central tool behind the upper bounds and lower bounds of Theorem 1.3, and is precisely where diophantine approximation enters our scenery. To wit, the proof of Assertion (1) of Theorem 1.7 makes use of the following powerful result.

NESTERENKO-MATVEEV THEOREM. *[Nes03, Thm. 2.1, Pg. 55] For any integers $c_1,\alpha_1,\ldots,c_N,\alpha_N$ with $\alpha_i\geq 2$ for all i, define $\Lambda(c,\alpha):=c_1\log(\alpha_1)+\cdots+c_N\log(\alpha_N)$. Then $\Lambda(c,\alpha)\neq 0\Longrightarrow\log\left|\frac{1}{\Lambda(c,\alpha)}\right|$ is bounded above by*

$$2.9(N+2)^{9/2}(2e)^{2N+6}(2+\log\max_j|c_j|)\prod_{j=1}^{N}\log|\alpha_j|.$$ ■

Assertion (1) of Theorem 1.7 will follow easily from the two algorithms we state below, once we prove their correctness and verify their efficiency. However, we will first need to recall the concept of a gcd-free basis. In essence, a gcd-free basis is nearly as powerful as factorization into primes, but is far easier to compute.

DEFINITION 2.13. *[BS96, Sec. 8.4] For any subset* $\{\alpha_1, \ldots, \alpha_N\} \subset \mathbb{N}$, *a* **gcd-free basis for** $\{\alpha_1, \ldots, \alpha_N\}$ *is a pair of sets* $\left(\{\gamma_i\}_{i=1}^{\eta}, \{e_{ij}\}_{(i,j)\in[N]\times[\eta]}\right)$ *such that (1)* $\gcd(\gamma_i, \gamma_j) = 1$ *for all* $i \neq j$, *and (2)* $\alpha_i = \prod_{j=1}^{\eta} \gamma_j^{e_{ij}}$ *for all* i. $\diamond$

ALGORITHM 2.14.
Input: *Integers* $\alpha_1, u_1, \ldots, \alpha_M, u_M$ *and* $\beta_1, v_1, \ldots, \beta_N, v_N$.
Output: *A true declaration as to whether* $\alpha_1^{u_1} \cdots \alpha_M^{u_M} = \beta_1^{v_1} \cdots \beta_N^{v_N}$.
Description:

0. *If* $\prod_{i=1}^{M} (\operatorname{sign} \alpha_i)^{u_i \bmod 2} \neq \prod_{i=1}^{N} (\operatorname{sign} \beta_i)^{v_i \bmod 2}$ *then output "*`They are not equal.`*" and* `STOP`.

1. *Replace the* α_i *and* β_i *by their absolute values and then construct a gcd-free basis* $\left(\{\gamma_i\}_{i=1}^{\eta}, \{e_{ij}\}_{(i,j)\in[M+N]\times[\eta]}\right)$ *for* $\{\alpha_1, \ldots, \alpha_M, \beta_1, \ldots, \beta_N\}$.

2. *If* $\sum_{i=1}^{M} e_{ij} u_i = \sum_{i=M+1}^{M+N} e_{ij} v_i$ *for all* $j \in [\eta]$ *then output "*`They are equal.`*" and* `STOP`.

3. *Output "*`They are not equal.`*"*

ALGORITHM 2.15.
Input: *Positive integers* $\alpha_1, u_1, \ldots, \alpha_M, u_M$ *and* $\beta_1, v_1, \ldots, \beta_N, v_N$ *with* $\alpha_i, \beta_i \geq 2$ *for all* i.
Output: *The sign of* $\alpha_1^{u_1} \cdots \alpha_M^{u_M} - \beta_1^{v_1} \cdots \beta_N^{v_N}$.
Description:

0. *Check via Algorithm 2.14 whether* $\alpha_1^{u_1} \cdots \alpha_M^{u_M} = \beta_1^{v_1} \cdots \beta_N^{v_N}$. *If so, output "*`They are equal.`*" and* `STOP`.

1. *Let* $U := \max\{u_1, \ldots, u_M, v_1, \ldots, v_N\}$, *and*
$$E := \tfrac{2.9}{\log 2}(2e)^{2M+2N+6}(1 + \log U) \times \left(\prod_{i=1}^{M} \log|\alpha_i|\right)\left(\prod_{i=1}^{N} \log|\beta_i|\right).$$

2. *For all* $i \in [M]$ *(resp.* $i \in [N]$*), let* A_i *(resp.* B_i*) be a rational number agreeing with* $\log \alpha_i$ *(resp.* $\log \beta_i$*) in its first* $2 + E + \log_2 M$ *(resp.* $2 + E + \log_2 N$*) leading bits.*[3]

3. *Output the sign of* $\left(\sum_{i=1}^{M} u_i A_i\right) - \left(\sum_{i=1}^{N} v_i B_i\right)$ *and* `STOP`.

LEMMA 2.16. *Algorithms 2.14 and 2.15 are both correct. Moreover, following the preceding notation, Algorithms 2.14 and 2.15 run within a number of bit operations asymptotically linear in, respectively,*
$$\left(\sum_{i=1}^{M} (1 + \log(u_i)\log(\alpha_i))^2\right) + \left(\sum_{i=1}^{N} (1 + \log(v_i)\log(\beta_i))^2\right)$$
and
$$(M+N)(30)^{M+N} L(\log U)\left(\prod_{i=1}^{M} L(\log(\alpha_i))\right)\left(\prod_{i=1}^{N} L(\log(\beta_i))\right),$$
where $L(x) := x \log^2(x) \log\log(x)$. ■

[3]For definiteness, let us use Arithmetic-Geometric Mean Iteration as in [Ber03] to find these approximations.

That Algorithm 2.14 is correct and runs in the time stated follows immediately from [BS96, Thm. 4.8.7, Sec. 4.8] and the naive complexity bounds for integer multiplication. The remainder of Lemma 2.16 then follows routinely from the Nesterenko-Matveev Theorem and the refined bit complexity estimates for fast multiplication of [BS96, Table 3.1, pg. 43].

2.5 Positive Feasibility for Circuits

For a real polynomial supported on a non-degenerate circuit, there are just two ways it can fail to have a positive root: a simple way and a subtle way. This is summarized below. Recall that the **Newton polytope of** f is simply $\operatorname{Newt}(f) := \operatorname{Conv}(\operatorname{Supp}(f))$, where $\operatorname{Conv}(S)$ denotes the **convex hull** (smallest convex set) containing S.

THEOREM 2.17. *Suppose* $f(x) = \sum_{i=1}^{n+2} c_i x^{a_i} \in \mathcal{F}^*_{n,n+2}$, $\operatorname{Supp}(f)$ *is a non-degenerate circuit, and* b *is the vector from Lemma 2.12. Then* $Z_+(f)$ *is empty iff one of the following conditions holds:*
1. *All the* c_i *have the same sign.*
2. $\operatorname{Newt}(f)$ *is an* n*-simplex and, assuming* $a_{j'}$ *is the unique element of* $\mathcal{A}$ *lying in the interior of* $\operatorname{Newt}(f)$, *we have* $-\operatorname{sign}(c_{j'}) = \operatorname{sign}(c_i)$ *for all* $i \neq j'$ *and*
$$\prod_{i=1}^{n+2} \left(\operatorname{sign}(b_{j'} c_{j'}) \tfrac{c_i}{b_i}\right)^{\operatorname{sign}(b_{j'} c_{j'}) b_i} > 1.$$

Proof of Theorem 2.17: First note that Condition (1) implies that f maintains the same (non-zero) sign throughout $\mathbb{R}^n_+$. So Condition (1) trivially implies that $Z_+(f) = \emptyset$, and we may assume henceforth that not all the coefficients of f have the same sign.

Now, if $\operatorname{Newt}(f)$ is not a simplex, then every point of $\mathcal{A}$ is a vertex of $\operatorname{Newt}(f)$ and thus, independent of the triangulation, any Viro diagram for $\mathcal{A}$ must be non-empty. Since there are only two discriminant chambers (by Lemma 2.12), and thus just two possible Viro diagrams, one can then show that $Z_+(f)$ must be non-empty, assuming $\Delta_{\mathcal{A}}(f) \neq 0$. Lemma 2.12 also tells us that $Z_+(f)$ must be non-empty if $\Delta_{\mathcal{A}}(f) = 0$. So we may assume henceforth that $\operatorname{Newt}(f)$ is a simplex and that $\Delta_{\mathcal{A}}(f) \neq 0$.

Continuing our focus on Condition (2), note that if the sign equalities from Condition (2) fail, then there must exist coefficients c_i and $c_{i'}$ of opposite sign such that a_i and $a_{i'}$ vertices of $\operatorname{Newt}(f)$. So, again, independent of the triangulation, any Viro diagram for $\mathcal{A}$ must be non-empty and thus (just as in the preceding paragraph) $Z_+(f)$ must again be non-empty. So we may assume henceforth that the sign equalities from Condition (2) hold.

At this point, it is clear that we need only show that (under our current assumptions) $Z_+(f) = \emptyset \iff$ the discriminant inequality from Condition (2) holds. Toward this end, observe that the lifting that assigns $a_{j'} \mapsto 1$ and $a_i \mapsto 0$ for all $i \neq j'$ induces the unique triangulation of $\mathcal{A}$ consisting of a single simplex. In particular, the underlying Viro diagram is empty, due to the sign equalities. So by Viro's Theorem, $Z_+(f)$ is empty for $|c_{j'}|$ sufficiently small. By Lemmata 2.11 and 2.12, this topology persists for $|c_{j'}|$ just small enough to enforce the discriminant sign stated in Condition (2), so we are done. ■

Positive feasibility for polynomials supported on degenerate circuits can then essentially be reduced to the non-degenerate circuit case in some lower dimension. An additional twist arises from the fact that the zero sets of polynomials supported on degenerate circuits are, up to a monomial

change of variables, the graphs of polynomials supported on non-degenerate circuits.

THEOREM 2.18. *Suppose $f(x)=\sum_{i=1}^{n+2} c_i x^{a_i} \in \mathcal{F}^*_{n,n+2}$ has support $\mathcal{A}\subset\mathbb{R}^n$ that is a degenerate circuit with non-degenerate subcircuit $\mathcal{B}=\{a_1,\ldots,a_{j'}\}$, and b is the vector defined in Lemma 2.12 (ignoring the non-degeneracy assumption for $\mathcal{A}$). Then, when not all the coefficients of f have the same sign, $Z_+(f)$ is empty iff both the following conditions hold:*

a. $\mathrm{Conv}(\mathcal{B})$ *is a* $(j'-2)$*-simplex and, permuting indices so that* $a_{j'}$ *is the unique element of* $\mathcal{B}$ *lying in the relative interior of* $\mathrm{Conv}(\mathcal{B})$*, we have* $-\mathrm{sign}(c_{j'})=\mathrm{sign}(c_i)$ *for all* $i\neq j'$.

b. $\prod_{i=1}^{j'}\left(\mathrm{sign}(b_{j'}c_{j'})\frac{c_i}{b_i}\right)^{\mathrm{sign}(b_{j'}c_{j'})b_i}\geq 1$.

Sketch of Proof of Theorem 2.18: Similar to Lemma 2.5 we can easily find a monomial change of variables (and multiply by a suitable monomial term) so that $\bar{f}(x):=x^v f(x^M)$ is of the form

$c_1+c_2x_1^{u_1}+\cdots+c_{j'-1}x_{j'-2}^{u_{j'-2}}+c_{j'}x^\alpha+c_{j'+1}x_{j'-1}^{u_{j'-1}}+\cdots+c_{n+2}x_n^{u_n}$,

where $u_1,\ldots,u_n\in\mathbb{N}$ and $\alpha\in\mathbb{N}^{j'-2}\times\{0\}^{n-j'+2}$. In particular, defining $\bar{f}_{\mathcal{B}}(x)=c_1+c_2x_1^{u_1}+\cdots+c_{j'-1}x_{j'-2}^{u_{j'-2}}+c_{j'}x^\alpha$, it is clear that $Z_+(\bar{f})$ is nothing more than the intersection of the graph of $\bar{f}_{\mathcal{B}}$ (which is analytic on $\mathbb{R}_+^{j'-2}$) with an orthant. So $Z_+\left(\bar{f}\right)$ is smooth and thus, by Proposition 2.4, we obtain that $Z_+(f)$ is diffeotopic to $Z_+\left(\bar{f}\right)$. It thus suffices to prove our theorem for $\bar{f}$. Note also that the conditions on $Z_+(f)$ allegedly characterizing $Z_+(f)=\emptyset$ are preserved under monomial multiples and monomial changes of variables. The remainder of the proof is then a case by case analysis, depending on whether $Z_+\left(\bar{f}\right)$ consists of 0, 1, or ≥ 2 points. ■

3. THE PROOFS OF OUR MAIN RESULTS: THEOREMS 1.7 AND 1.3

We go in increasing order of proof length.

3.1 Proving Theorem 1.7

Assertion (1): First note that any input f must have support $\mathcal{A}=\{a_1,\ldots,a_{n+2}\}$ equal to either a degenerate circuit or a non-degenerate circuit. Recalling Assertion (1) of Lemma 2.12, observe then that the vector $b:=(b_1,\ldots,b_{n+2})$ has a zero coordinate iff $\mathcal{A}$ is a degenerate circuit, and b can be computed in time polynomial in size($\mathcal{A}$). If $\mathcal{A}$ is a degenerate circuit then (following easily from the definition) $\Delta_{\mathcal{A}}$ must be identically 1, thus leaving Assertion (1) of our present theorem trivially true. So let us assume henceforth that $\mathcal{A}$ is a non-degenerate circuit, and that c_j is the coefficient of x^{a_j} in f for all j.

Via Assertion (1) of Lemma 2.12 once again, Assertion (1) of Theorem 1.7 follows routinely from the complexity bounds from Lemma 2.16. In particular, the latter lemma tells us that the bit complexity of $\mathbf{ADISC}_=$, for input coefficients $(c_1,\ldots,c_{n+2})$, is polynomial in $\sum_{i=1}^{n+2}\log(c_ib_i)$ (following the notation of Lemma 2.12); and the same is true for $\mathbf{ADISC}_>$ provided n is fixed. The classical Hadamard inequality then tells us that size$(b_i)=O(n\log(n\max_{j,k}\{a_{jk}\}))$. So the complexity of $\mathbf{ADISC}_=$ is indeed polynomial in size(f); and the same holds for $\mathbf{ADISC}_>$ when n is fixed. ■

Assertion (2): We will construct an explicit reduction of **3CNFSAT** to $\mathbf{ADISC}_=\left(\bigcup_{n\in\mathbb{N}\ ,\ 0<\varepsilon'\leq\varepsilon}\mathcal{F}^*_{n,n+n^{\varepsilon'}}\right)$. In particular, to any **3CNFSAT** instance $B(X)$ with N clauses and n variables with $N\geq n/3$, let us first consider $F_B=(f_1,\ldots,f_{8N-n})$ — the associated $(8N-n)\times(8N-n)$ polynomial system from Proposition 2.2 of Section 2.1. (By renaming variables, it is easy to see that a polynomial-time algorithm for **3CNFSAT** in the special case $N\geq n/3$ implies **3CNFSAT**$\in$**P**.)

Let us then set $M:=\left\lceil(17N-2n+2)^{1/\varepsilon}\right\rceil$ and define the **single** polynomial f_B to be

$f_1+\lambda_1f_2+\cdots+\lambda_{8N-n-1}f_{8N-n}+\lambda_{8N-n}t_M(z_1,\ldots,z_M)$.

Letting $\mathcal{A}$ be the support of f_B, it is then easily checked (from Proposition 2.2) that $\mathcal{A}$ is affinely independent and f_B is in $\mathcal{F}^*_{16N-2n+M,N'}$ for some $N'\leq 33N-4n+M+2$.

By the **Cayley Trick** [GKZ94, Prop. 1.7, pp. 274] we then obtain that $\Delta_{\mathcal{A}}(f_B)=0$ iff

($\star$) F_B has a degenerate root in $(\mathbb{P}^1_{\mathbb{C}})^{2N-n}$ **and** t_M has a degenerate root in $(\mathbb{C}^*)^M$.

(Since Newt(t_M) is a simplex, it is easily checked that t_M has no complex degenerate roots at infinity.) By Proposition 2.2, the degenerate roots of F_B are exactly $\{1,2\}^n\times W_B$, and t_M has a unique degenerate root by construction. So ($\star$) holds iff $B(X)$ has a satisfying assignment. We have thus reduced **3CNFSAT** to detecting the vanishing of a particular $\mathcal{A}$-discriminant.

To conclude, observe that the number of terms of f_B is only slightly larger than its number of variables, thanks to Proposition 2.2. In particular, size$(f_B)=O(\mathrm{size}(B)^{1/\varepsilon})$ and $f_B\in\bigcup_{n\in\mathbb{N}}\mathcal{F}^*_{n,n+n^\delta}$ for some $\delta\in(0,\varepsilon]$. Clearly then, $\mathbf{ADISC}_=\left(\bigcup_{n\in\mathbb{N}\ ,\ 0<\varepsilon'\leq\varepsilon}\mathcal{F}^*_{n,n+n^{\varepsilon'}}\right)\in\mathbf{P}\implies\mathbf{P}=\mathbf{NP}$, thus proving our first desired **NP**-hardness lower bound.

The **NP**-hardness of our remaining problem, being a subcase of a problem now shown to be **NP**-hard, then follows immediately. ■

3.2 Proving Theorem 1.3

Assertion (2): We will give an explicit reduction of **3CNFSAT** to $\mathbf{FEAS}_+\left(\bigcup_{n\in\mathbb{N}\ ,\ 0<\varepsilon'\leq\varepsilon}\mathcal{F}^*_{n,n+n^{\varepsilon'}}\right)$. Attaining such a reduction will require little effort, thanks to our earlier reduction used to prove Assertion (2) of Theorem 1.7.

In particular, for any **3CNFSAT** instance B with N clauses and n variables with $N\geq n/3$, let us recall the system $F_B=(f_1,\ldots,f_{8N-n})$ from Proposition 2.2. Let us then define M to be $\left\lceil(42N-n+2)^{1/\varepsilon}\right\rceil$ and define $g_B(x,z)$ to be $f_1^2(x)+\cdots+f_{4N}^2(x)+t_M(z_1,\ldots,z_M)$. It is then easily checked that $f_B\in\mathcal{F}^*_{n+M,N'}$ for some $N'\leq 42N+M+2$. Moreover, B has a satisfying assignment iff g_B has a positive root. (Indeed, any root of g_B clearly lies in $\{1,2\}^n\times\{1\}^M$.) We have thus reduced **3CNFSAT** to a special case of $\mathbf{FEAS}_+$.

Now observe that the number of terms of g_B is only slightly larger than its number of variables, thanks to Proposition 2.2. In particular, size$(g_B)=O\left(\mathrm{size}(B)^{1/\varepsilon}\right)$ and g_B is in $\bigcup_{n\in\mathbb{N}}\mathcal{F}^*_{n,n+n^\delta}$ for some $\delta\in(0,\varepsilon]$. Clearly then, $\mathbf{FEAS}_+\left(\bigcup_{n\in\mathbb{N}\ ,\ 0<\varepsilon'\leq\varepsilon}\mathcal{F}^*_{n,n+n^{\varepsilon'}}\right)\in\mathbf{P}\implies\mathbf{P}=\mathbf{NP}$, thus proving one of our desired **NP**-hardness lower bounds.

The **NP**-hardness of our remaining problem can be proved by employing our preceding argument almost verbatim. The only difference is that we instead use the polynomial

$h_B(x,z) := f_1^2(x) + \cdots + f_{4N}^2(x) + t_M(z_1^2,\ldots,z_M^2)$, and observe that $t_M(z_1^2,\ldots,z_M^2)$ is nonnegative on all of $\mathbb{R}^n$. So we are done. ■

Assertion (0): Our topological assertion follows immediately from Lemma 2.5 and Corollary 2.6.

To obtain our algorithmic assertions, simply note that by Assertion (1) of Corollary 2.6, detecting positive roots for f reduces to checking whether all the coefficients have the same sign. This can clearly be done by n sign evaluations and $n-1$ comparisons, doable in logarithmic parallel time. So the inclusion involving $\mathbf{FEAS}_+$ is proved.

Let us now show that we can detect roots in $(\mathbb{R}^*)^n$ within $\mathbf{NC}^1$: Employing our algorithm from the last paragraph, we can clearly assume the signs of the coefficients of f are all identical (for otherwise, we would have detected a root in $\mathbb{R}^n_+$ and finished). So then, by Assertion (2) of Corollary 2.6, we can simply do a parity check (trivially doable in $\mathbf{NC}^1$) of the entries of $[a_2-a_1,\ldots,a_{n+1}-a_1]$.

To conclude, we simply observe that our algorithm for detecting roots in $(\mathbb{R}^*)^n$ trivially extends to root detection in $\mathbb{R}^n$: Any root of f in $\mathbb{R}^n$ must lie in some coordinate subspace L of minimal positive dimension. So, on L, the honest n-variate $(n+1)$-nomial f will restrict to an $f' \in \mathcal{F}^*_{n',n'+1}$ with $n' \leq n$ and support a subset of the columns of a submatrix of $\mathcal{A}$. So then, we must check whether (a) all the coefficients of f' have the same sign or (if not), (b) a **sub**matrix of $[a_2-a_1,\ldots,a_{n+1}-a_1]$ has an odd entry. In other words, f has a root in $\mathbb{R}^n \iff f$ has a root in $(\mathbb{R}^*)^n \cup \{\mathbf{O}\}$. Since checking whether f vanishes at $\mathbf{O}$ is the same as checking whether f is missing a constant term, checking for roots in $\mathbb{R}^n$ is thus also in $\mathbf{NC}^1$. ■

REMARK 3.1. *Note that checking whether a given $f \in \mathcal{F}_{n,n+1}$ lies in $\mathcal{F}^*_{n,n+1}$ can be done within $\mathbf{NC}^2$: one simply finds $d = \dim \mathrm{Supp}(f)$ in $\mathbf{NC}^2$ by computing the rank of the matrix whose columns are $a_2-a_1,\ldots,a_m-a_1$ (via the parallel algorithm of Csanky [Csa76]), and then checks whether $d=n$.* ⋄

Assertion (1): The algorithm we use to prove $\mathbf{FEAS}_+(\mathcal{F}^*_{n,n+2}) \in \mathbf{P}$ for fixed n is described just below. Note also that once we have $\mathbf{FEAS}_+(\mathcal{F}^*_{n,n+2}) \in \mathbf{P}$ for fixed n, it easily follows that $\mathbf{FEAS}_{\mathbb{R}}(\mathcal{F}^*_{n,n+2}) \in \mathbf{P}$: The polynomial obtained from an $f \in \mathcal{F}^*_{n,n+2}$ by setting any non-empty subset of its variables to 0 clearly lies in $\mathcal{F}^*_{n',n'+2}$ for some $n'<n$ (modulo a permutation of variables). Thus, since we can apply changes of variables like $x_i \mapsto -x_i$ in $\mathbf{P}$, and since there are exactly 3^n sequences of the form $(\varepsilon_1,\ldots,\varepsilon_n)$ with $\varepsilon_i \in \{0,\pm 1\}$ for all i, it thus clearly suffices to show that $\mathbf{FEAS}_+(\mathcal{F}^*_{n,n+2}) \in \mathbf{P}$ for **fixed** n.

We thus need only prove correctness, and a suitable complexity bound, for the following algorithm:

ALGORITHM 3.2.
Input: *A coefficient vector $c := (c_1,\ldots,c_{n+2})$ and a (possibly degenerate) circuit $\mathcal{A}=\{a_1,\ldots,a_{n+2}\}$ of cardinality $n+2$.*
Output: *A true declaration as to whether $Z_+(f)$ is empty or not, where $f(x) := \sum_{i=1}^{n+2} c_i x^{a_i}$.*
Description:

1. *If all the c_i have the same sign then output "$Z_+(f)=\emptyset$" and* `STOP`.
2. *Let $b=(b_1,\ldots,b_{n+2}) \in \mathbb{Z}^n$ be the vector obtained by applying Lemma 2.12 to $\mathcal{A}$. If b or $-b$ has a unique negative coordinate $b_{j'}$, and $c_{j'}$ is the unique negative coordinate of c or $-c$, then do the following:*
   (a) *Replace b by $-\mathrm{sign}(b_{j'})b$, replace c by $-\mathrm{sign}(c_{j'})c$, and then reorder b, c, and $\mathcal{A}$ by the same permutation so that $b_{j'}<0$ and [$b_i>0$ iff $i<j'$].*
   (b) *If $j'<n+2$ and $(-b_{j'})^{-b_{j'}} \prod_{i=1}^{j'-1} c_i^{b_i} = (-c_{j'})^{-b_{j'}} \prod_{i=1}^{j'-1} b_i^{b_i}$ (the latter decided via Algorithm 2.14) then output "$Z_+(f)=\emptyset$" and* `STOP`.
   (c) *Decide via Algorithm 2.15 whether $(-b_{j'})^{-b_{j'}} \prod_{i=1}^{j'-1} c_i^{b_i} \stackrel{?}{>} (-c'_j)^{-b_{j'}} \prod_{i=1}^{j'-1} b_i^{b_i}$. If so, output "$Z_+(f)=\emptyset$" and* `STOP`.
3. *Output "$Z_+(f)$* `is non-empty!` *" and* `STOP`.

The correctness of Algorithm 3.2 follows directly from Theorems 2.18 and 2.17. In particular, note that b_i is simply the signed volume of $\mathrm{Conv}(\mathcal{A}\backslash\{a_i\})$. So the geometric interpretation b or $-b$ having a unique negative coordinate is that the convex hull of the unique non-degenerate subcircuit of $\mathcal{A}$ is a simplex, with $a_{j'}$ lying in its relative interior. Similarly, the geometric interpretation of $j'<n+2$ is that $\mathcal{A}$ is a degenerate circuit. Finally, the product comparisons from Steps (b) and (c) simply decide the product inequalities stated in Theorem 2.17 and Theorem 2.18.

So now we need only bound complexity, and this follows immediately from Lemma 2.16 (assuming we use Algorithm 2.14 for Step (b)). ■

It is worth noting that we need to compute the sign of a linear combination of logarithms only when the unique non-degenerate subcircuit $\mathcal{B}$ of $\mathcal{A}$ is a simplex, and all "vertex" coefficients have sign opposite from the "internal" coefficient. Also, just as in Remark 3.1, checking whether a given $f \in \mathcal{F}_{n,n+2}$ lies in $\mathcal{F}^*_{n,n+2}$ can be done within $\mathbf{NC}^2$ by computing $d = \dim \mathrm{Supp}(f)$ efficiently. Moreover, from our preceding proof, we see that deciding whether a circuit is degenerate (and extracting $\mathcal{B}$ from $\mathcal{A}$ when $\mathcal{A}$ is degenerate) can be done in $\mathbf{NC}^2$ as well, since we can set $\beta=1$ if we only want the signs of $(b_1,\ldots,b_{n+2})$.

Acknowledgements

The authors thank Francisco Santos for earlier discussions on counting regular triangulations, and Frank Sottile for inviting the second author to an April 2008 meeting at the Institute Henri Poincare where a version of these results was presented. Thanks also to Dima Pasechnik for discussions, and Sue Geller and Bruce Reznick for detailed commentary, on earlier versions of this work. We also thank AIM and IMA for their hospitality and support while this paper was being developed at respective workshops on Random Analytic Surfaces and Complexity, Coding, and Communication. Finally, we thank MSRI, and Philippe Pébay and David C. Thompson at Sandia National Laboratories, for their great hospitality during the completion of this paper.

Dedication Richard A. Snavely won the Allen G. Marr Prize for his 2005 UC Davis Ph.D. thesis on the physics of high-power, ultra-short pulse duration lasers. However, even more than for his diverse talents (including being a Sensei in Kendo), he will be remembered as a gentle, kind, and noble soul.

4. REFERENCES

[AM09] Avendaño, Martin and Mortari, Daniele, *"The Multi-Impulse Orbit Transfer Problem,"* preprint, Texas A&M University, 2009.

[BS96] Bach, Eric and Shallit, Jeff, *Algorithmic Number Theory, Vol. I: Efficient Algorithms,* MIT Press, Cambridge, MA, 1996.

[Bak77] Baker, Alan, *"The Theory of Linear Forms in Logarithms,"* in Transcendence Theory: Advances and Applications: proceedings of a conference held at the University of Cambridge, Cambridge, Jan.–Feb., 1976, Academic Press, London, 1977.

[BHPR09] Bastani, Osbert; Hillar, Chris; Popov, Dimitar; and Rojas, J. Maurice, *"Sums of Squares, Randomization, and Sparse Polynomials,"* in preparation, 2009.

[BG-V03] Basu, Saugata and Gonzalez-Vega, Laureano, *Algorithmic and Quantitative Real Algebraic Geometry,* Papers from the DIMACS Workshop on Algorithmic and Quantitative Aspects of Real Algebraic Geometry in Mathematics and Computer Science held at Rutgers University, Piscataway, NJ (March 12–16, 2001), DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 60.

[BPR06] Basu, Saugata; Pollack, Ricky; and Roy, Marie-Francoise, *Algorithms in Real Algebraic Geometry,* Algorithms and Computation in Mathematics, vol. 10, Springer-Verlag, 2006.

[Ber03] Bernstein, Daniel J., *"Computing Logarithm Intervals with the Arithmetic-Geometric Mean Iterations,"* available from `http://cr.yp.to/papers.html` .

[Ble04] Blekherman, Grigoriy, *"Convexity properties of the cone of nonnegative polynomials,"* Discrete Comput. Geom. 32 (2004), no. 3, pp. 345–371.

[Can88] Canny, John F., *"Some Algebraic and Geometric Computations in PSPACE,"* Proc. $20^{\underline{\text{th}}}$ ACM Symp. Theory of Computing, Chicago (1988), ACM Press.

[Csa76] Csanky, L., *"Fast Parallel Matrix Inversion Algorithms,"* SIAM J. Comput. **5** (1976), no. 4, pp. 618–623.

[DRRS07] Dickenstein, Alicia; Rojas, J. Maurice; Rusek, Korben; and Shih, Justin, *"$\mathcal{A}$-Discriminants and Extremal Real Algebraic Geometry,"* Moscow Mathematical Journal, vol. 7, no. 3, (July–September, 2007).

[DL79] Dobkin, David and Lipton, Richard, *"On the Complexity of Computations Under Varying Sets of Primitives,"* J. of Computer and System Sciences 18, pp. 86–91, 1979.

[GV04] Gabrielov, Andrei and Vorobjov, Nicolai, *"Complexity of computations with Pfaffian and Noetherian functions,"* Normal Forms, Bifurcations and Finiteness Problems in Differential Equations, pp. 211–250, Kluwer, 2004.

[GKZ94] Gel'fand, Israel Moseyevitch; Kapranov, Misha M.; and Zelevinsky, Andrei V.; *Discriminants, Resultants and Multidimensional Determinants,* Birkhäuser, Boston, 1994.

[IRR07] Ibrahim, Ashraf; Rojas, J. Maurice; and Rusek, Korben, *"First Steps in Algorithmic Arithmetic Fewnomial Theory,"* Math ArXiV preprint `0711.2562` .

[KS99] Karpinski, Marek and Shparlinski, Igor, *"On the computational hardness of testing square-freeness of sparse polynomials,"* Applied algebra, algebraic algorithms and error-correcting codes (Honolulu, HI, 1999), pp. 492–497, Lecture Notes in Comput. Sci., 1719, Springer, Berlin, 1999.

[Kho91] Khovanski, Askold, *Fewnomials,* AMS Press, Providence, Rhode Island, 1991.

[Las07] Lassere, Jean B., *"A Sum of Squares Approximation of Nonnegative Polynomials,"* SIAM Review, Vol. 49, No. 4, pp. 651–669.

[LM01] Lickteig, Thomas and Roy, Marie-Francoise, *"Sylvester-Habicht Sequences and Fast Cauchy Index Computation,"* J. Symbolic Computation (2001) **31**, pp. 315–341.

[LRW03] Li, Tien-Yien; Rojas, J. Maurice; and Wang, Xiaoshen, *"Counting Real Connected Components of Trinomial Curve Intersections and m-nomial Hypersurfaces,"* Discrete and Computational Geometry, 30 (2003), no. 3, pp. 379–414.

[Nes03] Nesterenko, Yuri, *"Linear forms in logarithms of rational numbers,"* Diophantine approximation (Cetraro, 2000), pp. 53–106, Lecture Notes in Math., 1819, Springer, Berlin, 2003.

[Pap95] Papadimitriou, Christos H., *Computational Complexity,* Addison-Wesley, 1995.

[Par03] Parrilo, Pablo A., *"Semidefinite programming relaxations for semialgebraic problems,"* Algebraic and geometric methods in discrete optimization, Math. Program. 96 (2003), no. 2, Ser. B, pp. 293–320.

[PRT09] Pébay, Philippe; Rojas, J. Maurice; and Thompson, David C., *"$\mathbf{NP}_{\mathbb{R}}$-Completeness and Sparse Polynomials,"* in preparation, 2009.

[Per08] Perrucci, Daniel, *"Algorithmic Aspects of Semialgebraic Geometry,"* Ph.D. Thesis (in Spanish), Mathematics Department, University of Buenos Aires, Argentina, 2008.

[RY05] Rojas, J. Maurice and Ye, Yinyu, *"On Solving Sparse Polynomials in Logarithmic Time,"* Journal of Complexity, special issue for the 2002 Foundations of Computation Mathematics (FOCM) meeting, February 2005, pp. 87–110.

[RS09] Rojas, J. Maurice and Sethuraman, Swaminathan, *"Refined Asymptotics for Sparse Sums of Squares,"* extended abstract, submitted for publication. Also available as Math Arxiv preprint `0901.3786` .

[Sch00] Schmid, Joachim, *"On the Complexity of the Real Nullstellensatz in the* 0*-Dimensional Case,"* J. Pure Appl. Algebra **151** (2000), no. 3, pp. 301–308.

[Sma87] Smale, Steve, *"On the Topology of Algorithms I,"* Journal of Complexity **3** (1987), no. 2, pp. 81–89.

[Sto98] Storjohann, Arne, *"Computing Hermite and Smith normal forms of triangular integer matrices,"* Linear Algebra Appl. 282 (1998), no. 1–3, pp. 25–45.

[Ste74] Stengle, G., *"A nullstellensatz and a positivstellensatz in semialgebraic geometry,"* Math. Ann. **207** (1974) pp. 87–97.