Optimization and $\mathrm{NP}_{\mathbb{R}}$-Completeness of Certain Fewnomials

Philippe Pébay*
Sandia National Laboratories
PO Box 969, MS 9159
Livermore, CA 94551
USA
pppebay@sandia.gov

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

David C. Thompson‡
Sandia National Laboratories
PO Box 969, MS 9159
Livermore, CA 94551
USA
dcthomp@sandia.gov

We dedicate this paper to Tien-Yien Li on the occasion of his 65th birthday. Happy 65, TY!.

ABSTRACT

We give a high precision polynomial-time approximation scheme for the supremum of any honest n-variate $(n+2)$-nomial with a constant term, allowing real exponents as well as real coefficients. Our complexity bounds count field operations and inequality checks, and are polynomial in n and the **logarithm** of a certain condition number. For the special case of polynomials (i.e., integer exponents), the log of our condition number is sub-quadratic in the sparse size. The best previous complexity bounds were exponential in the sparse size, even for n fixed. Along the way, we partially extend the theory of $\mathcal{A}$-discriminants to real exponents and exponential sums, and find new and natural $\mathbf{NP}_{\mathbb{R}}$-complete problems.

Categories and Subject Descriptors: F.2.1 [Analysis of Algorithms and Problem Complexity]: Numerical Algorithms and Problems—*Computations on polynomials*; G.1.5 [Numerical Analysis] Roots of Nonlinear Equations—*Iterative methods*

General Terms: Algorithms, Performance, Theory

Keywords: optimizing, sparse, BSS model, real, exponential sum, polynomial-time, circuit, approximate, condition number

*Supported by the United States Department of Energy, Office of Defense Programs.

†Partially supported by NSF CAREER grant DMS-0349309, the United States Department of Energy (Office of Defense Programs), MSRI, and the Wenner Gren Foundation.

‡Supported by the United States Department of Energy, Office of Defense Programs.

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94AL85000.

SNC'09, August 3–5, 2009, Kyoto, Japan.

1. INTRODUCTION AND MAIN RESULTS

Maximizing or minimizing polynomial functions is a central problem in science and engineering. Typically, the polynomials have an underlying structure, e.g., sparsity, small expansion with respect to a particular basis, invariance with respect to a group action, etc. In the setting of sparsity, Fewnomial Theory [Kho91] has succeeded in establishing bounds for the number of real solutions (or real extrema) that depend just on the number of monomial terms. However, the current general complexity bounds for real solving and nonlinear optimization (see, e.g., [BPR06, S08, Par03]) are still stated in terms of degree and number of variables, and all but ignore any finer input structure. In this paper, we present new speed-ups for the optimization of certain sparse multivariate polynomials, extended to allow real exponents as well. Along the way, we also present two new families of problems that are $\mathbf{NP}_{\mathbb{R}}$-complete, i.e., the analogue of **NP**-complete for the **BSS model over** $\mathbb{R}$. (The BSS model, derived in the 1980s by Blum, Shub, and Smale [BSS89], is a generalization of the classical Turing model of computation with an eye toward unifying bit complexity and algebraic complexity.)

Our framework has both symbolic and numerical aspects in that (a) we deal with real number inputs and (b) our algorithms give either yes or no answers that are always correct, or numerically approximate answers whose precision can be efficiently tuned. Linear Programming (LP) forms an interesting parallel to the complexity issues we encounter. In particular, while LP admits polynomial-time algorithms relative to the Turing model, polynomial-time algorithms for linear programming relative to the BSS model over $\mathbb{R}$ (a.k.a. strongly polynomial-time algorithms or polynomial arithmetic complexity) remain unknown. Furthermore, the arithmetic complexity of LP appears to be linked with a fundamental invariant measuring the intrinsic complexity of numerical solutions: the **condition number** (see, e.g., [VY93, CCP03]). Our results reveal a class of non-linear problems where similar subtleties arise when comparing discrete and continuous complexity.

To state our results, let us first clarify some basic notation concerning sparse polynomials and complexity classes over $\mathbb{R}$. Recall that $\lfloor x \rfloor$ is the greatest integer not exceeding a real number x, and 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*[1] *and, 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***).* ⋄

For example, the dishonestly 4-variate trinomial

$$-1 + \sqrt{7} x_1^2 x_2 x_3^7 x_4^3 - e^{43} x_1^{198e^2} x_2^{99e^2} x_3^{693e^2} x_4^{297e^2}$$

(with support contained in a line segment) has the same supremum over $\mathbb{R}^4_+$ as the **honest uni**variate trinomial

$$-1 + \sqrt{7} y_1 - e^{43} y_1^{99e^2}$$

has over $\mathbb{R}_+$. 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 BSS model over $\mathbb{R}$ (treated fully in [BCSS98]), so we will only briefly review a few definitions, focusing on a particular extension we need. Our underlying notion of input size, including a variant of the condition number, is clarified in Definition 2.1 of Section 2.1 below, and illustrated in Example 1.6 immediately following our first main theorem.

So for now let us just recall the following basic inclusions of complexity classes: $\mathbf{NC}^1_{\mathbb{R}} \subsetneq \mathbf{P}_{\mathbb{R}} \subseteq \mathbf{NP}_{\mathbb{R}}$ [BCSS98, Ch. 19, Cor. 1, pg. 364]. (The properness of the latter inclusion remains a famous open problem, akin to the more famous classical $\mathbf{P} \stackrel{?}{=} \mathbf{NP}$ question.) Let us also recall that $\mathbf{NC}^k_{\mathbb{R}}$ is the family of real valued functions (with real inputs) computable by arithmetic circuits[2] with size polynomial in the input size and depth $O(\log^k(\text{Input Size}))$ (see [BCSS98, Ch. 18] for further discussion).

To characterize a natural class of problems with efficiently computable numerical answers, we will define the notion of a **High Precision Polynomial Time Approximation Scheme**: We let $\mathbf{HPTAS}_{\mathbb{R}}$ denote the class of functions $\phi : \mathbb{R}^\infty \longrightarrow \mathbb{R} \cup \{+\infty\}$ such that, for any $\varepsilon > 0$, there is an algorithm guaranteed to approximate $\phi(x)$ to within a $1+\varepsilon$ factor, using a number of arithmetic operations polynomial in $\mathrm{size}(x)$ **and** $\log\log\frac{1}{\varepsilon}$.[3] Our notation is inspired by the more familiar classical family of problems **FPTAS** (i.e., those problems admitting a **Fully Polynomial Time Approximation Scheme**), where instead the input is Boolean and the complexity need only be polynomial in $\frac{1}{\varepsilon}$. The complexity class **FPTAS** was formulated in [ACGKM-SP99] and a highly-nontrivial example of a problem admitting a **FPTAS** is counting matchings in bounded degree graphs [BGKNT07].

[1] Here we allow real coefficients, unlike [BRS09] where the same notation included a restriction to integer coefficients.

[2] This is one of 2 times we will mention circuits in the sense of complexity theory: Everywhere else in this paper, our circuits will be **combinatorial** objects as in Definition 2.7 below.

[3] When $\phi(x) = 0$ we will instead require an **additive** error of ε or less. When $\phi(x) = +\infty$ we will require the approximation to be $+\infty$, regardless of ε.

REMARK 1.2. *For a vector function $\phi = (\phi_1, \ldots, \phi_k) : \mathbb{R}^\infty \longrightarrow (\mathbb{R} \cup \infty)^k$ it will be natural to say that ϕ admits an* **HPTAS** *iff each coordinate of ϕ_i admits an* **HPTAS***.* ⋄

1.1 Sparse Real Optimization

The main computational problems we address are the following.

DEFINITION 1.3. *Let $\mathbb{R}_+$ denote the positive real numbers, and let* **SUP** *denote the problem of deciding, for a given $(f, \lambda) \in \left(\bigcup_{n \in \mathbb{N}} \mathbb{R}[x^a \mid a \in \mathbb{R}^n]\right) \times \mathbb{R}$, whether $\sup_{x \in \mathbb{R}^n_+} f \geq \lambda$ or not. Also, for any subfamily $\mathcal{F} \subseteq \bigcup_{n \in \mathbb{N}} \mathbb{R}[x^a \mid a \in \mathbb{R}^n]$, we let $\mathbf{SUP}(\mathcal{F})$ denote the natural restriction of* **SUP** *to inputs in $\mathcal{F}$. Finally, we let* **FSUP** *(resp. $\mathbf{FSUP}(\mathcal{F})$) denote the obvious functional analogue of* **SUP** *(resp. $\mathbf{SUP}(\mathcal{F})$) where (a) the input is instead $(f, \varepsilon) \in \left(\bigcup_{n \in \mathbb{N}} \mathbb{R}[x^a \mid a \in \mathbb{R}^n]\right) \times \mathbb{R}_+$ and (b) the output is instead a pair*

$$(\bar{x}, \bar{\lambda}) \in (\mathbb{R}_+ \cup \{0, +\infty\})^n \times (\mathbb{R} \cup \{+\infty\})$$

with $\bar{x} = (\bar{x}_1, \ldots, \bar{x}_n)$ (resp. $\bar{\lambda}$) an **HPTAS** *for x^* (resp. λ^*) where $\lambda^* := \sup_{x \in \mathbb{R}^n_+} f = \lim_{x \to x^*} f(x)$ for some $x^* = (x^*_1, \ldots, x^*_n) \in (\mathbb{R}_+ \cup \{0, +\infty\})^n$.*

REMARK 1.4. *Taking logarithms, it is clear that our problems above are equivalent to maximizing a function of the form $g(y) = \sum_{i=1}^m c_i e^{a_i \cdot y}$ over $\mathbb{R}^n$. When convenient, we will use the latter notation but, to draw parallels with the algebraic case, we will usually speak of "polynomials" with real exponents.* ⋄

We will need to make one final restriction when optimizing n-variate m-nomials: we let $\mathcal{F}^{**}_{n,n+k}$ denote the subset of $\mathcal{F}^*_{n,n+k}$ consisting of those f with $\mathrm{Supp}(f) \ni \mathbf{O}$. While technically convenient, this restriction is also natural in that level sets of $(n+k)$-nomials in $\mathcal{F}^{**}_{n,n+k}$ become zero sets of $(n+k')$-nomials with $k' \leq k$.

We observe that checking whether the zero set of an $f \in \mathbb{R}[x_1, \ldots, x_n]$ is nonempty (a.k.a. the **real (algebraic) feasibility problem**) is equivalent to checking whether the maximum of $-f^2$ is 0 or greater. So it can be argued that the **NP**-hardness (and $\mathbf{NP}_{\mathbb{R}}$-hardness) of **SUP** has been known at least since the 1990s [BCSS98]. However, it appears that no sharper complexity upper bounds in terms of sparsity were known earlier.

THEOREM 1.5. *We can efficiently optimize n-variate $(n+k)$-nomials over $\mathbb{R}^n_+$ for $k \leq 2$. Also, for k a slowly growing function of n, optimizing n-variate $(n+k)$-nomials over $\mathbb{R}^n_+$ is* **NP***-hard. More precisely:*

0. Both $\mathbf{SUP}\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}^{**}_{n,n+1}\right)$ and $\mathbf{FSUP}\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}^{**}_{n,n+1}\right)$ are in $\mathbf{NC}^1_{\mathbb{R}}$.

1. $\mathbf{SUP}\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}^{**}_{n,n+2}\right) \in \mathbf{P}_{\mathbb{R}}$ and $\mathbf{FSUP}\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}^{**}_{n,n+2}\right) \in \mathbf{HPTAS}_{\mathbb{R}}$.

2. For any fixed $\delta > 0$, $\mathbf{SUP}\left(\bigcup_{\substack{n \in \mathbb{N} \\ 0 < \delta' < \delta}} \mathcal{F}^{**}_{n,n+n^{\delta'}} \cap \mathbb{R}[x_1, \ldots, x_n]\right)$ is $\mathbf{NP}_{\mathbb{R}}$-complete.

EXAMPLE 1.6. *Suppose $\varepsilon > 0$. A very special case of Assertion (1) of Theorem 1.5 then implies that we can approximate within a factor of $1+\varepsilon$ — for any real nonzero $c_1, \ldots, c_{n+2}$ and D — the maximum of the function $f(x)$ defined to be*

$$c_1 + c_2(x_1^D \cdots x_n^{D^n}) + c_3 x_1^{2D} \cdots x_n^{2^n D^n} + \cdots + c_{n+2} x_1^{(n+1)D} \cdots x_n^{(n+1)^n D^n},$$

using a number of arithmetic operations linear in

$$n^2(\log(n) + \log D) + \log\log\tfrac{1}{\varepsilon}.$$

The best previous results in the algebraic setting (e.g., the critical points method as detailed in [S08], or by combining [BPR06] and the efficient numerical approximation results of [MP98]) would yield a bound polynomial in

$$n^n D^n + \log\log\tfrac{1}{\varepsilon}$$

instead, and only under the assumption that $D \in \mathbb{N}$. Alternative approaches via semidefinite programming also appear to result in complexity bounds superlinear in $n^n D^n$ (see, e.g., [Par03, Las06, DN08, KM09]), and still require $D \in \mathbb{N}$. Moving to Pfaffian/Noetherian function techniques, [GV04] allows arbitrary real D but still yields an arithmetic complexity bound exponential in n. It should of course be pointed out that the results of [BPR06, MP98, S08, Par03, Las06, DN08, KM09, GV04] apply to real polynomials in complete generality. ⋄

We thus obtain a significant speed-up for a particular class of analytic functions, laying some preliminary groundwork for improved optimization of $(n+k)$-nomials with k arbitrary. Our advance is possible because, unlike earlier methods which essentially revolved around commutative algebra (and were more suited to complex algebraic geometry), we are addressing a real analytic problem with real analytic tools. Theorem 1.5 is proved in Section 3.2 below. Our main new technique, which may be of independent interest, is an extension of $\mathcal{A}$-discriminants (a.k.a. sparse discriminants) to real exponents (Theorem 2.9 of Section 2.3).

Our algorithms are quite implementable (see Algorithm 3.2 of Section 3.2) and derived via a combination of tropical geometric ideas and $\mathcal{A}$-discriminant theory, both extended to real exponents. In particular, for n-variate $(n+1)$-nomials, a simple change of variables essentially tells us that tropical geometry rules (in the form of **Viro diagrams** [GKZ94, Ch. 5, pp. 378–393], but extended to real exponents), and in the case at hand this means that one can compute extrema by checking inequalities involving the coefficients (and possibly an input λ). Tropical geometry still applies to the n-variate $(n+2)$-nomial case, but only after one evaluates the sign of a particular generalized $\mathcal{A}$-discriminant.[4] More precisely, an n-variate m-nomial f (considered as a function on $\mathbb{R}_+^n$) with bounded supremum λ^* must attain the value λ^* at a critical point of f in the nonnegative orthant. In particular, the nonnegative zero set of $f - \lambda^*$ must be degenerate, and thus we can attempt to solve for λ^* (and a corresponding maximizer) if we have a sufficiently tractable notion of discriminant to work with.

So our hardest case reduces to (a) finding efficient formulas for discriminants of n-variate m-nomials and (b) efficiently detecting unboundedness for n-variate m-nomials. When $m = n+2$, (a) fortuitously admits a solution, based on a nascent theory developed further in [CR09]. We can also reduce Problem (b) to Problem (a) via some tropical geometric tricks. So our development ultimately hinges deriving an efficient analogue of discriminant polynomials for discriminant varieties that are no longer algebraic.

[4]For n-variate $(n+3)$-nomials, knowing the sign of a discriminant is no longer sufficient, and efficient optimization still remains an open problem. Some of the intricacies are detailed in [DRRS07, BHPR09].

Example 1.7. *Consider the trivariate pentanomial $f := c_1 + c_2 x_1^{999} + c_3 x_1^{73} x_3^{\sqrt{363}} + c_4 x_2^{2009} + c_5 x_1^{74} x_2^{108e} x_3$, with $c_1, \ldots, c_4 < 0$ and $c_5 > 0$. Theorem 2.10 of Section 2.4 then easily implies that f attains a maximum of λ^* on $\mathbb{R}_+^3$ iff $f - \lambda^*$ has a degenerate root in $\mathbb{R}_+^3$. Via Theorem 2.9 of Section 2.3 below, the latter occurs iff*

$$b_5^{b_5}(c_1 - \lambda^*)^{b_1} c_2^{b_2} c_3^{b_3} c_4^{b_4} - b_1^{b_1} b_2^{b_2} b_3^{b_3} b_4^{b_4} c_5^{b_5}$$

vanishes, where $b := (b_1, b_2, b_3, b_4, -b_5)$ is any generator of the kernel of the map $\varphi : \mathbb{R}^5 \longrightarrow \mathbb{R}^4$ defined by the matrix

$$\begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & 999 & 73 & 0 & 74 \\ 0 & 0 & 0 & 2009 & 108e \\ 0 & \sqrt{363} & 0 & 0 & 1 \end{bmatrix},$$

normalized so that $b_5 > 0$. In particular, such a b can be computed easily via 5 determinants of 4×4 submatrices (via Cramer's Rule), and we thus see that λ^* is nothing more than c_1 minus a monomial (involving real exponents) in $c_2, \ldots, c_5$. Via the now classical fast algorithms for approximating log *and* exp *[Bre76], real powers of real numbers (and thus λ^*) can be efficiently approximated. Similarly, deciding whether λ^* exceeds a given λ reduces to checking an inequality involving real powers of positive numbers.* ⋄

1.2 Related Work

The computational complexity of numerical analysis continues to be an active area of research, both in theory and in practice. On the theoretical side, the BSS model over $\mathbb{R}$ has proven quite useful for setting a rigourous foundation. While this model involves exact arithmetic and field operations, there are many results building upon this model that elegantly capture round-off error and numerical conditioning (see, e.g., [CS99, ABKM09]). Furthermore, results on $\mathbf{P}_\mathbb{R}$ and $\mathbf{NP}_\mathbb{R}$ do ultimately impact classical complexity classes. For instance, the respective **Boolean parts** of these complexity classes, $\mathbf{BP}(\mathbf{P}_\mathbb{R})$ and $\mathbf{BP}(\mathbf{NP}_\mathbb{R})$, are defined as the respective restrictions of $\mathbf{P}_\mathbb{R}$ and $\mathbf{NP}_\mathbb{R}$ to integer inputs. While the best known bounds for these Boolean parts are still rather loose —

$\mathbf{P/Poly} \subseteq \mathbf{BP}(\mathbf{P}_\mathbb{R}) \subseteq \mathbf{PSPACE/Poly}$ [ABKM09],

$\mathbf{NP/Poly} \subseteq \mathbf{BP}(\mathbf{NP}_\mathbb{R}) \subseteq \mathbf{CH}$ [ABKM09],

— good algorithms for the BSS model and good algorithms for the Turing model frequently inspire one another, e.g., [Koi99, BPR06]. We recall that $\mathbf{P/Poly}$, referred to as **non-uniform polynomial-time**, consists of those decision problems solvable by a non-uniform family of circuits[5] of size polynomial in the input. $\mathbf{CH}$ is the **counting hierarchy** $\mathbf{PP} \cup \mathbf{PP}^{\mathbf{PP}} \cup \mathbf{PP}^{\mathbf{PP}^{\mathbf{PP}}} \cup \cdots$, which happens to be contained in **PSPACE** (see [ABKM09] and the references therein).

Let us also point out that the number of natural problems known to be $\mathbf{NP}_\mathbb{R}$-complete remains much smaller than the number of natural problems known to be **NP**-complete: deciding the existence of a real roots for multivariate polynomials (and various subcases involving quadratic systems or single quartic polynomials) [BCSS98, Ch. 5], linear programming feasibility [BCSS98, Ch. 5], and bounding the real dimension of algebraic sets [Koi99] are the main representative $\mathbf{NP}_\mathbb{R}$-complete problems. Optimizing n-variate $(n+n^\delta)$-nomials (with $\delta > 0$ fixed and n unbounded), and the corresponding feasibility problem (cf. Corollary 1.10 below), now join this short list.

While sparsity has been profitably explored in the context of interpolation (see, e.g., [KY07, GLL09]) and factoriza-

[5]i.e., there is no restriction on the power of the algorithm specifying the circuit for a given input size

tion over number fields [Len99, KK06, AKS07], it has been mostly ignored in numerical analysis (for nonlinear polynomials) and the study of the BSS model over $\mathbb{C}$ and $\mathbb{R}$. For instance, there appear to be no earlier published complexity upper bounds of the form $\mathbf{SUP}\,(\mathcal{F}_{1,m})\in\mathbf{P}_{\mathbb{R}}$ (relative to the sparse encoding) for any $m\geq 3$, in spite of beautiful recent work in semi-definite programming (see, e.g., [Las06, DN08, KM09]) that begins to address the optimization of sparse multivariate polynomials over the real numbers. In particular, while the latter papers give significant practical speed-ups over older techniques such as resultants and Gröbner bases, the published complexity bounds are still exponential (relative to the sparse encoding) for n-variate $(n+2)$-nomials, and require the assumption of integer exponents.

We can at least obtain a glimpse of sparse optimization beyond n-variate $(n+2)$-nomials by combining our framework with an earlier result from [RY05]. The proof is in Section 3.3.

Corollary 1.8.
(0) Using the same notion of input size as for $\mathbf{FSUP}$ *(cf. Definition 2.1 below), the positive roots of any real trinomial in* $\mathcal{F}_{1,3}\cap\mathbb{R}[x_1]$ *admit an* $\mathbf{HPTAS}$.
(1) $\mathbf{SUP}(\mathcal{F}^{**}_{1,4}\cap\mathbb{R}[x_1])\in\mathbf{P}_{\mathbb{R}}$ *and* $\mathbf{FSUP}(\mathcal{F}^{**}_{1,4}\cap\mathbb{R}[x_1])\in\mathbf{HPTAS}_{\mathbb{R}}$.

As for earlier complexity lower bounds for $\mathbf{SUP}$ in terms of sparsity, we are unaware of any. For instance, it is not even known whether $\mathbf{SUP}(\mathbb{R}[x_1,\ldots,x_n])$ is $\mathbf{NP}_{\mathbb{R}}$-hard for some fixed n (relative to the sparse encoding).

The paper [BRS09], which deals exclusively with decision problems (i.e., yes/no answers) and bit complexity (as opposed to arithmetic complexity), is an important precursor to the present work. Here, we thus expand the context to real coefficient and real exponents, work in the distinct setting of optimization, and derive (and make critical use of) a new tool: generalized $\mathcal{A}$-discriminants for exponential sums. As a consequence, we are also able to extend some of the complexity lower bounds from [BRS09] as follows. (See Section 3.2 for the proof.)

Definition 1.9. *Let* $\mathbf{FEAS}_{\mathbb{R}}$ *(resp.* $\mathbf{FEAS}_{+}$*) denote the problem of deciding whether an arbitrary system of equations from* $\bigcup_{n\in\mathbb{N}}\mathbb{R}[x^a\mid a\in\mathbb{R}^n]$ *has a real root (resp. root with all coordinates positive). Also, for any collection* $\mathcal{F}$ *of tuples chosen from* $\bigcup_{k,n\in\mathbb{N}}(\mathbb{R}[x^a\mid a\in\mathbb{R}^n])^k$*, we let* $\mathbf{FEAS}_{\mathbb{R}}(\mathcal{F})$ *(resp.* $\mathbf{FEAS}_{+}(\mathcal{F})$*) denote the natural restriction of* $\mathbf{FEAS}_{\mathbb{R}}$ *(resp.* $\mathbf{FEAS}_{+}$*) to inputs in* $\mathcal{F}$. ⋄

Corollary 1.10. *For any* $\delta>0$,

$$\mathbf{FEAS}_{\mathbb{R}}\left(\bigcup_{\substack{n\in\mathbb{N}\\0<\delta'<\delta}}\mathcal{F}^{**}_{n,n+n^{\delta'}}\cap\mathbb{R}[x_1,\ldots,x_n]\right)$$

and

$$\mathbf{FEAS}_{+}\left(\bigcup_{\substack{n\in\mathbb{N}\\0<\delta'<\delta}}\mathcal{F}^{**}_{n,n+n^{\delta'}}\cap\mathbb{R}[x_1,\ldots,x_n]\right)$$

are each $\mathbf{NP}_{\mathbb{R}}$*-complete.*

2. BACKGROUND

2.1 Input Size

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

Definition 2.1. *Given any subset* $\mathcal{A}=\{a_1,\ldots,a_m\}\subset\mathbb{R}^n$ *of cardinality* m*, let us define* $\hat{\mathcal{A}}$ *to be the* $(n+1)\times m$ *matrix whose* $j^{\underline{th}}$ *column is* $\{1\}\times a_j$*, and* β_J *the absolute value of the determinant of the submatrix of* $\hat{\mathcal{A}}$ *consisting of those columns of* $\hat{\mathcal{A}}$ *with index in a subset* $J\subseteq\{1,\ldots,m\}$ *of cardinality* $n+1$*. Then, given any* $f\in\mathcal{F}^{*}_{n,m}$ *written* $f(x)=\sum_{i=1}^m c_i x^{a_i}$*, we define its* **condition number***,* $\mathcal{C}(f)$*, to be*

$$\left(\prod_{i=1}^m \max\left\{3,|c_i|,\tfrac{1}{|c_i|}\right\}\right)\times\prod_{\substack{J\subseteq\{1,\ldots,m\}\\ \#J=n+1}}\max{}^{*}\left(3,|\beta_J|,\tfrac{1}{|\beta_J|}\right),$$

where $\max^*(a,b,c)$ *is* $\max\{a,b,c\}$ *or* a*, according as* $\max\{b,c\}$ *is finite or not.*

Throughout this paper, we will use the following notions of input size for $\mathbf{SUP}$ *and* $\mathbf{FSUP}$*: The size of any instance* (f,λ) *of* $\mathbf{SUP}$ *(resp. an instance* (f,ε) *of* $\mathbf{FSUP}$*) is* $\log\left(\max^*\left(3,|\lambda|,\frac{1}{|\lambda|}\right)\right)+\log\mathcal{C}(f)$ *(resp.* $\log\mathcal{C}(f)$*).* ⋄

While our definition of condition number may appear unusual, it is meant to concisely arrive at two important properties: (1) $\log\mathcal{C}(f)$ is polynomial in $n\log\deg f$ when $f\in\mathcal{F}_{n,n+k}\cap\mathbb{R}[x_1,\ldots,x_n]$ and k is fixed, (2) $\mathcal{C}(f)$ is closely related to an underlying discriminant (see Theorem 2.9 below) that dictates how much numerical accuracy we will need to solve $\mathbf{FSUP}$. We also point out that for $f\in\mathbb{Z}[x_1,\ldots,x_n]$, it is easy to show that $\log\mathcal{C}(f)=O(nS(f))$ where $S(f)$ is the **sparse size** of f, i.e., $S(f)$ is the number of bits needed to write down the monomial term expansion of f. For sufficiently sparse polynomials, algorithms with complexity polynomial in $S(f)$ are much faster than those with complexity polynomial in n and $\deg(f)$. [Len99, KK06, AKS07, KY07, GLL09, BRS09] provide other interesting examples of algorithms with complexity polynomial in $S(f)$.

2.2 Tricks with Exponents

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

Definition 2.2. *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* $xz:=(x_1z_1,\ldots,x_nz_n)$*. Finally, let* $\mathbb{GL}_n(\mathbb{R})$ *denote the group of all invertible matrices in* $\mathbb{R}^{n\times n}$*.* ⋄

Proposition 2.3. *(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. 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* $\mathbb{R}^n$*.* ■

A consequence follows: 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 $\sum_{a\in\mathcal{A}}c_a=0$. To optimize an $f\in\mathcal{F}^{**}_{n,n+1}$ it will help to have a much simpler canonical form. In what follows, we use $\#$ for set cardinality and e_i for the $i^{\underline{\mathrm{th}}}$ standard basis vector of $\mathbb{R}^n$.

Corollary 2.4. *For any* $f \in \mathcal{F}^{**}_{n,n+1}$ *we can compute* $c \in \mathbb{R}$ *and* $\ell \in \{0,\ldots,n\}$ *within* $\mathbf{NC}^1_{\mathbb{R}}$ *such that*

$$\bar{f}(x) := c + x_1 + \cdots + x_\ell - x_{\ell+1} - \cdots - x_n$$

satisfies:

(1) $\bar{f}$ *and* f *have exactly the same number of positive coefficients, and*

(2) $\bar{f}(\mathbb{R}^n_+) = f(\mathbb{R}^n_+)$.

Proof: Suppose f has support $\mathcal{A} = \{0, a_2, \ldots, a_{n+1}\}$ and corresponding coefficients $c_1, \ldots, c_{n+2}$. Letting B denote the $n \times n$ matrix whose $i^{\underline{\text{th}}}$ column is a_{i+1}, Proposition 2.3, via the substitution $x = y^{B^{-1}}$, tells us that we may assume that f is of the form $c_1 + c_2x_1 + \cdots + c_{n+1}x_n$. Moreover, to obtain $\bar{f}$, we need only perform a suitable positive rescaling and reordering of the variables. In summary, c is simply the constant term of f and ℓ is the number of positive coefficients not belonging to the constant term — both of which can be computed simply by a search and a sort clearly belonging to $\mathbf{NC}^1_{\mathbb{R}}$. ■

Note that we don't actually need to compute B^{-1} to obtain ℓ: B^{-1} is needed only for the proof of our corollary.

A final construction we will need is the notion of a **generalized** Viro diagram. Recall that a **triangulation** of a point set $\mathcal{A}$ is simply a simplicial complex Σ whose vertices lie in $\mathcal{A}$. We say that a triangulation of $\mathcal{A}$ is **induced by a lifting** iff it its simplices are exactly the domains of linearity for some function that is convex, continuous, and piecewise linear on the convex hull of[6] $\mathcal{A}$.

Definition 2.5. *Suppose* $\mathcal{A} \subset \mathbb{R}^n$ *is finite,* $\dim \operatorname{Aff}\mathcal{A} = n$, *and* $\mathcal{A}$ *is equipped with a triangulation* Σ *induced by a lifting* **and** *a function* $s : \mathcal{A} \longrightarrow \{\pm\}$ *which we will call a* **distribution of signs for** $\mathcal{A}$. *We then define a piecewise linear manifold — the* **generalized Viro diagram** $\mathcal{V}_{\mathcal{A}}(\Sigma, s)$ *— in the following local manner: For any* n*-cell* $C \in \Sigma$, *let* L_C *be the convex hull of the set of midpoints of edges of* C *with vertices of opposite sign, and then define* $\mathcal{V}_{\mathcal{A}}(\Sigma, s) := \bigcup_{C \text{ an } n\text{-cell}} L_C$. *When* $\mathcal{A} = \operatorname{Supp}(f)$ *and* s *is the corresponding sequence of coefficient signs, then we also call* $\mathcal{V}_\Sigma(f) := \mathcal{V}_{\mathcal{A}}(\Sigma, s)$ *the* **(generalized) Viro diagram of** f. ⋄

We use the appelation "generalized" since, to the best of our knowledge, Viro diagrams have only been used in the special case $\mathcal{A} \subset \mathbb{Z}^n$ (see, e.g., Proposition 5.2 and Theorem 5.6 of [GKZ94, Ch. 5, pp. 378–393]). We give examples of Viro diagrams in Section 2.4 below.

2.3 Generalized Circuit Discriminants and Efficient Approximations

Our goal here is to extract an extension of $\mathcal{A}$-discriminant theory sufficiently strong to prove our main results.

Definition 2.6. *Given any* $\mathcal{A} = \{a_1, \ldots, a_m\} \subset \mathbb{R}^n$ *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* **generalized** $\mathcal{A}$**-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* $g(x) = \sum_{i=1}^m c_i e^{a_i \cdot y}$ *has a degenerate root in* $\mathbb{C}^n$. *In particular, we call* f *an* n**-variate exponential** m**-sum**. ⋄

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

[6] i.e., smallest convex set containing...

Definition 2.7. *We call* $\mathcal{A} \subset \mathbb{R}^n$ *a* **(non-degenerate) circuit**[7] *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* $a \in \mathcal{B}$, $\mathcal{A} \setminus a$ *is affinely independent, and* $\mathcal{B}$ *is a non-degenerate circuit.* ⋄

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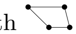 and 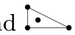 are circuits, but 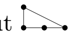 is a degenerate circuit. In general, for any degenerate circuit $\mathcal{A}$, the subset $\mathcal{B}$ named above is always unique.

Definition 2.8. *For any* $\mathcal{A} \subset \mathbb{R}^n$ *of cardinality* m, *let* $\mathcal{G}_{\mathcal{A}}$ *denote the set of all* n*-variate exponential* m*-sums with support* $\mathcal{A}$. ⋄

There is then a surprisingly succinct description for $\nabla_{\mathcal{A}}$ when $\mathcal{A}$ is a non-degenerate circuit. The theorem below is inspired by [GKZ94, Prop. 1.2, pg. 217] and [GKZ94, Prop. 1.8, Pg. 274] — important precursors that covered the special case of integral exponents.

Theorem 2.9. *Suppose* $\mathcal{A} = \{a_1, \ldots, a_{n+2}\} \subset \mathbb{R}^n$ *is a non-degenerate circuit, and let* $b := (b_1, \ldots, b_{n+2})$ *where* b_i *is* $(-1)^i$ *times the determinant of the matrix with columns* $1 \times a_1, \ldots, \widehat{1 \times a_i}, \ldots, a_{n+2}$ *(*$\widehat{(\cdot)}$ *denoting omission). Then:*

1. $\nabla_{\mathcal{A}} \subseteq \left\{ [c_1 : \cdots : c_{n+2}] \in \mathbb{P}^{n+1}_{\mathbb{C}} \ : \ \prod_{i=1}^{n+2} \left| \frac{c_i}{b_i} \right|^{b_i} = 1 \right\}$. *Also,* $(b_1, \ldots, b_{n+2})$ *can be computed in* $\mathbf{NC}^2_{\mathbb{R}}$.
2. *There is a* $[c_1 : \cdots : c_{n+2}] \in \mathbb{P}^{n+1}_{\mathbb{R}}$ *with*

   (i) $\operatorname{sign}(c_1b_1) = \cdots = \operatorname{sign}(c_{n+2}b_{n+2})$

   and

   (ii) $\prod_{i=1}^{n+2} (\operatorname{sign}(b_ic_i)c_i/b_i)^{\operatorname{sign}(b_ic_i)b_i} = 1$

   iff the real zero set of $g(y) := \sum_{i=1}^{n+2} c_i e^{a_i \cdot y}$ *contains a degenerate point* ζ. *In particular, any such* ζ *satisfies* $e^{a_i \cdot \zeta} = \operatorname{sign}(b_1c_1)b_i/c_i$ *for all* i, *and thus the real zero set of* g *has at most one degenerate point.*

Theorem 2.9 is proved in Section 3 below.

We will also need a variant of a family of fast algorithms discovered independently by Brent and Salamin.

Brent-Salamin Theorem. *[Bre76, Sal76] Given any positive* $x, \varepsilon > 0$, *we can approximate* $\log x$ *and* $\exp(x)$ *within a factor of* $1 + \varepsilon$ *using just* $O\left(|\log x| + \log\log \frac{1}{\varepsilon}\right)$ *arithmetic operations.* ■

While Brent's paper [Bre76] does not explicitly mention general real numbers, he works with a model of floating point number from which it is routine to derive the statement above.

2.4 Unboundedness and Sign Checks

Optimizing an $f \in \mathcal{F}^{**}_{n,n+1}$ will ultimately reduce to checking simple inequalities involving just the coefficients of f. The optimum will then in fact be either $+\infty$ or the constant term of f. Optimizing an $f \in \mathcal{F}^{**}_{n,n+2}$ would be as easy were it not for two additional difficulties: deciding unboundedness already entails checking the sign of a generalized $\mathcal{A}$-discriminant, and the optimum can be a transcendental function of the coefficients.

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

To formalize the harder case, let us now work at the level of exponential sums: let us define $\mathcal{G}_{n,m}$, $\mathcal{G}^*_{n,m}$, and $\mathcal{G}^{**}_{n,m}$ to be the obvious respective exponential m-sum analogues of $\mathcal{F}_{n,m}$, $\mathcal{F}^*_{n,m}$, and $\mathcal{F}^{**}_{n,m}$. Recall that $\mathrm{Conv}\mathcal{A}$ is the convex hull of $\mathcal{A}$.

Theorem 2.10. *Suppose we write* $g\in\mathcal{G}^{**}_{n,n+2}$ *in the form* $g(y)=\sum_{i=1}^{n+2} c_i e^{a_i\cdot y}$ *with* $\mathcal{A}=\{a_1,\ldots,a_{n+2}\}$*. Let us also order the monomials of* f *so that* $\mathcal{B}:=\{a_1,\ldots,a_{j'}\}$ *is the unique non-degenerate sub-circuit of* $\mathcal{A}$ *and let* $b:=(b_1,\ldots,b_{n+2})$ *where* b_i *is* $(-1)^i$ *times the determinant of the matrix with columns* $1\times a_1,\ldots,\widehat{1\times a_i},\ldots,a_{n+2}$ *(*$\widehat{(\cdot)}$ *denoting omission). Then* $\sup_{y\in\mathbb{R}^n} g(y)=+\infty \iff$ *one of the following 2 conditions hold:*

1. $c_j>0$ *for some vertex* a_j *of* $\mathrm{Conv}\mathcal{A}$ *not equal to* $\mathbf{O}$.
2. $\mathbf{O}\notin\mathcal{B}$*, we can further order the monomials of* f *so that* $a_{j'}$ *is the unique point of* $\mathcal{B}$ *in the relative interior of* $\mathcal{B}$*,* $c_{j'}>0$*, and* $\prod_{i=1}^{j'}\left(\mathrm{sign}(b_{j'})\frac{c_i}{b_i}\right)^{\mathrm{sign}(b_{j'})b_i}<1$.

Finally, if $\sup_{y\in\mathbb{R}^n} g(y)=\lambda^*<+\infty$ *and* $a_j=\mathbf{O}$*, then* $\lambda^*=c_j$*, or* λ^* *is the unique solution to*

$$\left(\mathrm{sign}(b_{j'})\frac{c_j-\lambda^*}{b_j}\right)^{\mathrm{sign}(b_{j'})b_j}\times\prod_{i\in\{1,\ldots,j'\}\setminus\{j\}}\left(\mathrm{sign}(b_{j'})\frac{c_i}{b_i}\right)^{\mathrm{sign}(b_{j'})b_i}=1$$

with $(c_j-\lambda^*)b_jb_{j'}>0$*; where the equation for* λ^* *holds iff:*

3. $\mathbf{O}\in\mathcal{B}$*, we can further order the monomials of* f *so that* $a_{j'}$ *is the unique point of* $\mathcal{B}$ *in the relative interior of* $\mathcal{B}$*, and* $c_{j'}>0$.

It is easily checked that $c_1b_1b_{j'},\ldots,c_{j'-1}b_{j'-1}b_{j'}>0$ when Conditions 2 or 3 hold. While the 3 cases above may appear complicated, they are easily understood from a tropical perspective: our cases above correspond to 4 different families of generalized Viro diagrams that characterize how the function g can be bounded from above (or not) on $\mathbb{R}^n$. Some representative examples are illustrated below:

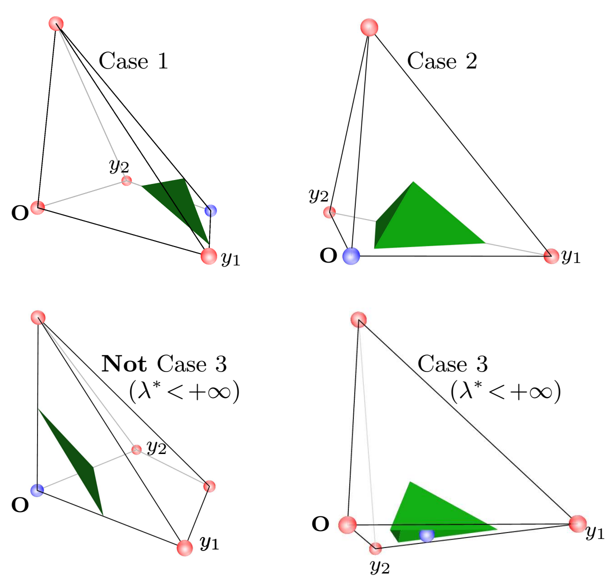

For example, the first two illustrations are meant to encode the fact that there exist directions in the positive quadrant along which g increases without bound. Similarly, the last 2 illustrations respectively show cases where g either approaches a supremum as some $y_i\longrightarrow-\infty$ or has a unique maximum in the real plane.

Sketch of Proof of Theorem 2.10: First, we identify the graph of g over $\mathbb{R}^n$ with the real zero set Z of $z-g(y)$. Since the supremum of g is unaffected by a linear change of variables, we can then assume (analogous to Corollary 2.4) that g is of the form

$$c+e^{y_1}+\cdots+e^{y_\ell}-e^{y_{\ell+1}}-\cdots-e^{y_n}+c'e^{\alpha\cdot y}.$$

(Note in particular that a linear change of variables for an exponential sum is, modulo applications of exp and log, the same as a monomial change of variables.) Note also that the classical Hadamard bound for the determinant guarantees that $\log\mathcal{C}(g)$ increases by at worst a factor of n after our change of variables. Let P denote the convex hull of $\{\mathbf{O},e_1,\ldots,e_n,e_{n+1},\alpha\}$.

Via a minor variation of the **moment map** (see, e.g., [Ful93]) one can then give a homeomorphism $\varphi:\mathbb{R}^{n+1}\longrightarrow\mathrm{Int}(P)$ that extends to a map $\bar\varphi$ encoding the "limits at toric infinity" of Z in terms of data involving P. (See also [LRW03, Sec. 6].) In particular, $\bar\varphi(Z)$ intersects the facet of P parallel to the y_i coordinate hyperplane iff Z contains points with y_i coordinates approaching $-\infty$. Similarly, the function g is unbounded iff $\bar\varphi(Z)$ intersects a face of P incident to e_{n+1} and some point in $\{e_1,\ldots,e_n,\alpha\}$. This correspondence immediately accounts for Condition 1.

This correspondence also accounts for Condition 2, but in a more subtle manner. In particular, Z has topology depending exactly on which connected component of the complement of $\nabla_\mathcal{A}$ contains g. Thanks to Theorem 2.9, this can be decided by determining the sign of expression involving powers of ratios of c_i and b_i. In particular, Condition 2 is nothing more than an appropriate accounting of when $\bar\varphi(Z)$ intersects a face of P incident to e_{n+1} and some point in $\{e_1,\ldots,e_n,\alpha\}$.

To conclude, one merely observes that Condition 3 corresponds to $\bar\varphi(Z)$ intersecting a face of P incident to $\mathbf{O}$ and e_{n+1}. In particular, the sign conditions merely guarantee that g has a unique maximum as some y_i tend to $-\infty$. ■

3. THE PROOFS OF OUR MAIN RESULTS: THEOREMS 2.9 AND 1.5, AND COROLLARIES 1.10 AND 1.8

We go in increasing order of proof length.

3.1 The Proof of Theorem 2.9

Assertion (1): It is easily checked that $Z_\mathbb{C}(f)$ has a degenerate point ζ iff

$$\hat{\mathcal{A}}\begin{bmatrix}c_1e^{a_1\cdot\zeta}\\ \vdots\\ c_{n+2}e^{a_{n+2}\cdot\zeta}\end{bmatrix}=\begin{bmatrix}0\\ \vdots\\ 0\end{bmatrix}.$$

In which case, $(c_1e^{a_1\cdot y},\ldots,c_{n+2}e^{a_{n+2}\cdot y})^T$ must be a generator of the right null space of $\hat{\mathcal{A}}$. On the other hand, by Cramer's Rule, one sees that $(b_1,\ldots,b_{n+2})^T$ is also a generator of the right null space of $\hat{\mathcal{A}}$. In particular, $\mathcal{A}$ a non-degenerate circuit implies that $b_i\neq0$ for all i.

We therefore obtain that

$$(c_1e^{a_1\cdot\zeta},\ldots,c_{n+2}e^{a_{n+2}\cdot\zeta})=\alpha(b_1,\ldots,b_{n+2})$$

for some $\alpha\in\mathbb{C}^*$. Dividing coordinate-wise and taking absolute values, we then obtain

$$\left(|c_1/b_1|e^{a_1\cdot\mathbf{Re}(\zeta)},\ldots,|c_{n+2}/b_{n+2}|e^{a_{n+2}\cdot\mathbf{Re}(\zeta)}\right)=(|\alpha|,\ldots,|\alpha|).$$

Taking both sides to the vector power $(b_1,\ldots,b_{n+2})$ we then

clearly obtain
$\left(|c_1/b_1|^{b_1}\cdots|c_{n+2}/b_{n+2}|^{b_{n+2}}\right)\left(e^{(b_1a_1+\cdots b_{n+2}a_{n+2})\cdot\mathbf{Re}(\zeta)}\right)=|\alpha|^{b_1+\cdots+b_n}$.
Since $\hat{A}(b_1,\ldots,b_{n+2})^T=\mathbf{O}$, we thus obtain $\prod_{i=1}^{n+2}\left|\frac{c_i}{b_i}\right|^{b_i}=1$.
Since the last equation is homogeneous in the c_i, its zero set in $\mathbb{P}^{n+1}_{\mathbb{C}}$ actually defines a closed set of $[c_1:\cdots:c_{n+2}]$. So we obtain the containment for $\nabla_{\mathcal{A}}$.

The assertion on the complexity of computing $(b_1,\ldots,b_{n+2})$ then follows immediately from the classic efficient parallel algorithms for linear algebra over $\mathbb{R}$ [Csa76]. ■

Assertion (2): We can proceed by almost exactly the same argument as above, using one simple additional observation: $e^{a_i\cdot\zeta}\in\mathbb{R}_+$ for all i when $\zeta\in\mathbb{R}$. So then, we can replace our use of absolute value by a sign factor, so that all real powers are well-defined. In particular, we immediately obtain the "$\Longleftarrow$" direction of our desired equivalence.

To obtain the "$\Longrightarrow$" direction, note that when
$$Z_{\mathbb{R}}\left(\sum_{i=1}^{n+2}c_ie^{a_i\cdot y}\right)$$
has a degeneracy ζ, we directly obtain $e^{a_i\cdot\zeta}=\text{sign}(b_1c_1)b_i/c_i$ for all i (and the constancy of sign(b_ic_i) in particular). We thus obtain the system of equations
$$\left(e^{(a_2-a_1)\cdot\zeta},\ldots,e^{(a_{n+1}-a_1)\cdot\zeta}\right)=\left(\frac{b_2c_1}{b_1c_2},\ldots,\frac{b_{n+1}c_1}{b_1c_{n+1}}\right),$$
and $a_2-a_1,\ldots,a_{n+1}-a_1$ are linearly independent since $\mathcal{A}$ is a circuit. So, employing Proposition 2.3, we can easily solve the preceding system for ζ by taking the logs of the coordinates of $\left(\frac{b_2c_1}{b_1c_2},\ldots,\frac{b_{n+1}c_1}{b_1c_{n+1}}\right)^{[a_2-a_1,\ldots,a_{n+1}-a_1]^{-1}}$. ■

3.2 Proving Corollary 1.10 and Theorem 1.5

Corollary 1.10 and Assertion (2) of Theorem 1.5: Since our underlying family of putative hard problems shrinks as δ decrease, it clearly suffices to prove the case $\delta<1$. So let assume henceforth that $\delta<1$. Let us also define $\mathbf{QSAT}_{\mathbb{R}}$ to be the problem of deciding whether an input **quartic** polynomial $f\in\bigcup_{n\in\mathbb{N}}\mathbb{R}[x_1,\ldots,x_n]$ has a real root or not. $\mathbf{QSAT}_{\mathbb{R}}$ (referred to as 4-FEAS in [BCSS98]) is one of the fundamental $\mathbf{NP}_{\mathbb{R}}$-complete problems (see Chapter 4 of [BCSS98]).

That $\mathbf{SUP}\in\mathbf{NP}_{\mathbb{R}}$ follows immediately from the definition of $\mathbf{NP}_{\mathbb{R}}$. So it suffices to prove that
$$\mathbf{SUP}\left(\bigcup_{\substack{n\in\mathbb{N}\\0<\delta'<\delta}}\mathcal{F}^{**}_{n,n+n^{\delta'}}\cap\mathbb{R}[x_1,\ldots,x_n]\right)$$
is $\mathbf{NP}_{\mathbb{R}}$-hard. We will do this by giving an explicit reduction of $\mathbf{QSAT}_{\mathbb{R}}$ to
$$\mathbf{SUP}\left(\bigcup_{\substack{n\in\mathbb{N}\\0<\delta'<\delta}}\mathcal{F}^{**}_{n,n+n^{\delta'}}\cap\mathbb{R}[x_1,\ldots,x_n]\right),$$
passing through $\mathbf{FEAS}_+\left(\bigcup_{\substack{n\in\mathbb{N}\\0<\delta'<\delta}}\mathcal{F}^{**}_{n,n+n^{\delta'}}\cap\mathbb{R}[x_1,\ldots,x_n]\right)$ along the way.

To do so, let f denote any $\mathbf{QSAT}_{\mathbb{R}}$ instance, involving, say, n variables. Clearly, f has no more than $\binom{n+4}{4}$ monomial terms. Letting $\mathbf{QSAT}_+$ denote the natural variant of $\mathbf{QSAT}_{\mathbb{R}}$ where one instead asks if f has a root in $\mathbb{R}^n_+$, we will first need to show that $\mathbf{QSAT}_+$ is $\mathbf{NP}_{\mathbb{R}}$-hard as an intermediate step. This is easy, via the introduction of slack variables: using $2n$ new variables $\left\{x_i^{\pm}\right\}_{i=1}^n$ and forming the polynomial $f^{\pm}(x^{\pm}):=f\left(x_1^+-x_1^-,\ldots,x_n^+-x_n^-\right)$, it is clear that f has a root in $\mathbb{R}^n$ iff $f^{\pm}$ has a root in $\mathbb{R}^{2n}_+$. Furthermore, we easily see that
$$\text{size}(f^{\pm})=(16+o(1))\text{size}(f).$$
So $\mathbf{QSAT}_+$ is $\mathbf{NP}_{\mathbb{R}}$-hard. We also observe that we may restrict the inputs to quartic polynomials with full-dimensional Newton polytope, since the original proof for the $\mathbf{NP}_{\mathbb{R}}$-hardness of $\mathbf{QSAT}_{\mathbb{R}}$ actually involves polynomials having nonzero constant terms and nonzero x_i^4 terms for all i [BCSS98].

So now let f be any $\mathbf{QSAT}_+$ instance with, say, n variables. Let us also define, for any $M\in\mathbb{N}$, the polynomial $t_M(z):=1+z_1^{M+1}+\cdots+z_M^{M+1}-(M+1)z_1\cdots z_M$. One can then check via the Arithmetic-Geometric Inequality [HLP88] that t_M is nonnegative on $\mathbb{R}^M_+$, with a unique root at $z=(1,\ldots,1)$. Note also that f^2 has no more than $\binom{n+4}{4}^2$ monomial terms. Forming the polynomial $F(x,z):=f(x)^2+t_M(z)$ with $M:=\left\lceil\binom{n+4}{4}^{2/\delta}\right\rceil$, we see that f has a root in $\mathbb{R}^n_+$ iff F has a root in $\mathbb{R}^{n+M}_+$. It is also easily checked that $F\in\mathcal{F}^{**}_{N,N+k}$ with $k\leq N^{\delta'}$, where $N:=n+M$ and $0<\delta'\leq\delta$. In particular,
$$k<\binom{n+4}{4}^2\leq\left\lceil\binom{n+4}{4}^{2/\delta}\right\rceil^{\delta}=M^{\delta}<(n+M)^{\delta}.$$
So we must now have that
$$\mathbf{FEAS}_+\left(\bigcup_{\substack{n\in\mathbb{N}\\0<\delta'<\delta}}\mathcal{F}^{**}_{n,n+n^{\delta}}\cap\mathbb{R}[x_1,\ldots,x_n]\right)$$
is $\mathbf{NP}_{\mathbb{R}}$-hard. (A small digression allows us to succinctly prove that
$$\mathbf{FEAS}_{\mathbb{R}}\left(\bigcup_{\substack{n\in\mathbb{N}\\0<\delta'<\delta}}\mathcal{F}^{**}_{n,n+n^{\delta}}\cap\mathbb{R}[x_1,\ldots,x_n]\right)$$
is $\mathbf{NP}_{\mathbb{R}}$-hard as well: we simply repeat the argument from the last paragraph, but use $\mathbf{QSAT}_{\mathbb{R}}$ in place of $\mathbf{QSAT}_+$, and define $F(x,z):=f(x)^2+t_M(z_1^2,\ldots,z_M^2)$ instead.)

To conclude, note that $F(x,z)$ is nonnegative on $\mathbb{R}^n_+$. So by checking whether $-F$ has supremum ≥ 0 in $\mathbb{R}^n_+$, we can decide if F has a root in $\mathbb{R}^n_+$. In other words,
$$\mathbf{SUP}\left(\bigcup_{\substack{n\in\mathbb{N}\\0<\delta'<\delta}}\mathcal{F}^{**}_{n,n+n^{\delta'}}\cap\mathbb{R}[x_1,\ldots,x_n]\right)$$
must be $\mathbf{NP}_{\mathbb{R}}$-hard has well. So we are done. ■

Assertion (0) of Theorem 1.5: Letting (f,ε) denote any instance of $\mathbf{FSUP}\left(\bigcup_{n\in\mathbb{N}}\mathcal{F}^{**}_{n,n+1}\right)$, first note that via Corollary 2.4 we can assume that
$$f(x)=c_1+x_1+\cdots+x_\ell-x_{\ell+1}-\cdots-x_n,$$
after a computation in $\mathbf{NC}^1_{\mathbb{R}}$. Clearly then, f has an unbounded supremum iff $\ell\geq 1$. Also, if $\ell=0$, then the supremum of f is exactly c_1. So $\mathbf{FSUP}\left(\bigcup_{n\in\mathbb{N}}\mathcal{F}^{**}_{n,n+1}\right)\in\mathbf{NC}^1_{\mathbb{R}}$. That $\mathbf{SUP}\left(\bigcup_{n\in\mathbb{N}}\mathcal{F}^{**}_{n,n+1}\right)\in\mathbf{NC}^1_{\mathbb{R}}$ is obvious as well: after checking the signs of the c_i, we make merely decide the sign of $c_1-\lambda$. ■

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 $d=\dim\text{Supp}(f)$ 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): We will first derive the **HPTAS** result. Let us assume $f\in\mathcal{F}^{**}_{n,n+2}$ and observe the following algorithm:

ALGORITHM 3.2.

Input: *A coefficient vector* $c := (c_1, \ldots, c_{n+2})$*, a (possibly degenerate) circuit* $\mathcal{A} = \{a_1, \ldots, a_{n+2}\}$ *of cardinality* $n+2$*, and a precision parameter* $\varepsilon > 0$.

Output: *A pair* $(\bar{x}, \bar{\lambda}) \in (\mathbb{R}_+ \cup \{0, +\infty\})^n \times (\mathbb{R} \cup \{+\infty\})$ *with* $\bar{x} = (\bar{x}_1, \ldots, \bar{x}_n)$ *(resp.* $\bar{\lambda}$*) an* **HPTAS** *for* x^* *(resp.* λ^**) where* $f(x) := \sum_{i=1}^{n+2} c_i x^{a_i}$ *and* $\lambda^* := \sup_{x \in \mathbb{R}_+^n} f = \lim_{x \to x^*} f(x)$ *for some* $x^* = (x_1^*, \ldots, x_n^*) \in (\mathbb{R}_+ \cup \{0, +\infty\})^n$.

Description:

1. *If* $c_i > 0$ *for some* i *with* $a_i \neq \mathbf{O}$ *a vertex of* $\mathrm{Conv}\mathcal{A}$ *then output*
   "`f tends to` $+\infty$ `along a curve of the form` $\{ct^{a_i}\}_{t \to +\infty}$*"*
   and `STOP`.

2. *Let* $b := (b_1, \ldots, b_{n+2})$ *where* b_j *is* $(-1)^j$ *times the determinant of the matrix with columns* $1 \times a_1, \ldots, \widehat{1 \times a_j}, \ldots, a_{n+2}$ *(*$\widehat{(\cdot)}$ *denoting omission). If* b *or* $-b$ *has a unique negative coordinate* $b_{j'}$*, and* $c_{j'}$ *is the unique negative coordinate of* c*, then do the following:*

   (a) *Replace* b *by* $-\mathrm{sign}(b_{j'})b$ *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* $a_i \neq \mathbf{O}$ *for all* $i \in \{1, \ldots, j'\}$ *and*
   $$\prod_{i=1}^{j'} \left(\mathrm{sign}(b_{j'}) \frac{c_i}{b_i}\right)^{\mathrm{sign}(b_{j'}) b_i} < 1$$
   then output
   "`f tends to` $+\infty$ `along a curve of the form` $\{ct^{a_{j'}}\}_{t \to +\infty}$*"*
   and `STOP`.

   (c) *If* $a_j = \mathbf{O}$ *for some* $j \in \{1, \ldots, j'\}$ *then output*
   "`f(z) tends to a supremum of` $\bar{\lambda}$ `as` z `tends to the point` $\bar{x}$ `on the` $(j'-2)$`-dimensional sub-orbit corresponding to` $\{a_1, \ldots, a_{j'}\}$.*"*, *where* $x \in \mathbb{R}_+^{j'-2}$ *is the unique solution to the binomial system*
   $$\left(x^{a_2 - a_1}, \ldots, x^{a_{j'-1} - a_1}\right) = \left(\frac{b_2 c_1}{b_1 c_2}, \ldots, \frac{b_{n+1} c_1}{b_1 c_{n+1}}\right),$$
   $\bar{x}$ *is a* $(1+\varepsilon)$*-factor approximation*[8] *of* x*,* λ *is the unique solution of*
   $$\left(\mathrm{sign}(b_{j'}) \frac{c_j - \lambda}{b_j}\right)^{\mathrm{sign}(b_{j'}) b_j} \times \prod_{i \in \{1, \ldots, j'\} \setminus \{j\}} \left(\mathrm{sign}(b_{j'}) \frac{c_i}{b_i}\right)^{\mathrm{sign}(b_{j'}) b_i} = 1$$
   with $(c_j - \lambda) b_j b_{j'} > 0$*, and* $\bar{\lambda}$ *is a* $(1+\varepsilon)$*-factor approximation*[8] *of* λ*; and* `STOP`.

3. *Output*
   "`f approaches a supremum of` c_j `as all` x^{a_i} `with` a_i `incident to` a_j `approach 0.`*"*,
   where $a_j = \mathbf{O}$*, and* `STOP`.

Our proof then reduces to proving correctness and a suitable complexity bound for Algorithm 3.2. In particular, correctness follows immediately from Theorem 2.10. So we now focus on a complexity analysis.

Steps 1 and 3 can clearly be done within $\mathbf{NC}^1_{\mathbb{R}}$, so let us focus on Step 2.

For Step 2, the dominant complexity comes from Part (b). (Part (a) can clearly be done in $\mathbf{NC}^1_{\mathbb{R}}$, and Part (c) can clearly be done in $\mathbf{NC}^2_{\mathbb{R}}$ via Csanky's method [Csa76].) The latter can be done by taking the logarithm of each term, thus reducing to checking the sign of a linear combination of logarithms of positive real numbers. So the arithmetic complexity of our algorithm is $O\left(\log \mathcal{C}(f) + \log\log \frac{1}{\varepsilon}\right)$ and we thus obtain our **HPTAS** result.

The proof that $\mathbf{SUP}\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}^{**}_{n,n+2}\right) \in \mathbf{P}_{\mathbb{R}}$ is almost completely identical. ■

Note that 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, 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 this is ultimately the evaluation of $n+2$ determinants.

3.3 The Proof of Corollary 1.8

Assertion (0): Since the roots of f in $\mathbb{R}_+$ are unchanged under multiplication by monomials, we can clearly assume $f \in \mathcal{F}^{**}_{1,3} \cap \mathbb{R}[x_1]$. Moreover, via the classical Cauchy bounds on the size of roots of polynomials, it is easy to show that the log of any root of f is $O(\log \mathcal{C}(f))$. We can then invoke Theorem 1 of [RY05] to obtain our desired **HPTAS** as follows: If $D := \deg(f)$, [RY05, Theorem 1] tells us that we can count exactly the number of positive roots of f using $O(\log^2 D)$ arithmetic operations, and ε-approximate all the roots of f in $(0, R)$ within $O\left((\log D) \log\left(D \log \frac{R}{\varepsilon}\right)\right)$ arithmetic operations. Since we can take $\log R = O(\log \mathcal{C}(f))$ via our root bound observed above, we are done. ■

Assertion (1): Writing any $f \in \mathcal{F}^{**}_{1,4} \cap \mathbb{R}[x_1]$ as $f(x) = c_1 + c_2 x^{a_2} + c_3 x^{a_3} + c_4 x^{a_4}$ with $a_2 < a_3 < a_4$, note that f has unbounded supremum on $\mathbb{R}_+$ iff $c_4 > 0$ So let us assume $c_4 < 0$.

Clearly then, the supremum of f is attained either at a critical point in $\mathbb{R}_+$ or at 0. But then, any positive critical point is a positive root of a trinomial, and by Assertion (0), such critical points must admit an **HPTAS**. Similarly, since f is a tetranomial (and thus evaluable within $O(\log \deg(f))$ arithmetic operations), we can efficiently approximate (as well as efficiently check inequalities involving) $\sup_{x \in \mathbb{R}_+} f$. So we are done. ■

Acknowledgements

We thank Peter Bürgisser, Felipe Cucker, Johan Hastad, and Gregorio Malajovich for discussions on complexity classes over $\mathbb{R}$. The second author would also like to thank MSRI and the Wenner Gren Foundation for their support during the completion of this paper. In particular, special thanks go to Mikael Passare and Boris Shapiro of Stockholm University for their hospitality.

4. REFERENCES

[ABKM09] Allender, E.; Bürgisser, P.; Kjeldgaard-Pedersen, J.; and Miltersen, P., *"On the Complexity of Numerical Analysis,"* SIAM J. Comput. 38(5): pp. 1987-2006 (2009).

[ACGKM-SP99] Ausiello, G.; Crescenzi, P.; Gambosi, G.; Kann, V.; Marchetti-Spaccamela, A. and Protasi, M., *"Complexity and Approximation: Combinatorial optimization problems and their approximability properties,"* Springer-Verlag, 1999.

[AKS07] Avendaño, Martin; Krick, Teresa; and Sombra, Martin, *"Factoring bivariate sparse (lacunary) polynomials,"* J. Complexity 23 (2007), no. 2, pp. 193–216.

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

[8] We compute $\bar{x}$ and $\bar{\lambda}$ via Proposition 2.3 and the Brent-Salamin Theorem.

[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.

[BGKNT07] Bayati, Mohsen; Gamarnik, David; Katz, Dimitriy; Nair, Chandra; and Tetali, Prasad, *"Simple deterministic approximation algorithms for counting matchings,"* Proceedings of the 39th annual ACM symposium on Theory of computing (STOC 2007, San Diego, California, USA), pp. 122–127, ACM, New York, 2007.

[BRS09] Bihan, Frederic; Rojas, J. Maurice; Stella, Case E., *"Faster Real Feasibility via Circuit Discriminants,"* proceedings of International Symposium on Symbolic and Algebraic Computation (ISSAC 2009, July 28–31, Seoul, Korea), to appear.

[BSS89] Blum, Lenore; Shub, Mike; and Smale, Steve, *"On a theory of computation and complexity over the real numbers: NP-completeness, recursive fnctions and universal machines„* Bulletin of the Amer. Math. Soc. 21, pp. 1–46.

[BCSS98] Blum, Lenore; Cucker, Felipe; Shub, Mike; and Smale, Steve, *Complexity and Real Computation,* Springer-Verlag, 1998.

[Bre76] Brent, Richard P., *"Fast Multiple-Precision Evaluation of Elementary Functions,"* Journal of the Association for Computing Machinery, vol. 23, No. 2, April 1976, pp. 242–251.

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

[CCP03] Cheung, Dennis; Cucker, Felipe; and Peña, Javier, *"Unifying Condition Numbers for Linear Programming,"* Math. of Oper. Res., 28, pp. 609–624, 2003.

[CR09] Crofts, Evelyn L. and Rojas, J. Maurice, *"Complex Vertices for $\mathcal{A}$-discriminants and Mixed Volumes,"* in progress.

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

[CS99] Cucker, Felipe and Smale, Steve, *"Complexity estimates depending on condition and round-off error,"* J. ACM 46 (1999), no. 1, pp. 113–184.

[DN08] Demmel, James and Nie, Jiawang, *"Sparse SOS relaxations for minimizing functions that are summations of small polynomials,"* SIAM Journal On Optimization, Vol. 19, No. 4, pp. 1534–1558 (2008).

[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).

[Ful93] Fulton, William, *Introduction to Toric Varieties*, Annals of Mathematics Studies, no. 131, Princeton University Press, Princeton, New Jersey, 1993.

[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.

[GLL09] Giesbrecht, Mark; Labahn, George; Lee, Wen-Shin, *"Symbolic-numeric sparse interpolation of multivariate polynomials,"* Journal of Symbolic Computation, vol. 44, no. 8, pp. 943–959, 2009.

[HLP88] Hardy, G. H.; Littlewood, J. E.; Pólya, G., *Inequalities,* reprint of the 1952 edition, Cambridge Mathematical Library, Cambridge University Press, Cambridge, 1988.

[KK06] Kaltofen, Erich and Koiran, Pascal, *"Finding small degree factors of multivariate supersparse (lacunary) polynomials over algebraic number fields,"* in proceedings of ISSAC 2006 (Internat. Symp. Symbolic Algebraic Comput.), pp. 162–168.

[KY07] Kaltofen, Erich and Yang, Zhengfeng, *"On exact and approximate interpolation of sparse rational functions,"* in proceedings of ISSAC 2007 (Internat. Symp. Symbolic Algebraic Comput.), pp. 203–210.

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

[Koi99] Koiran, Pascal, *"The Real Dimension Problem is* $\mathrm{NP}_{\mathbb{R}}$*-complete,"* J. Complexity 15 (1999), no. 2, pp. 227–238.

[KM09] Kojima, M. and Muramatsu, M., *"A Note on Sparse SOS and SDP Relaxations for Polynomial Optimization Problems over Symmetric Cones,"* Computational Optimization and Applications Vol. 42 (1), pp. 31–41 (2009).

[Las06] Lassere, Jean B., *"Convergent SDP-Relaxations in Polynomial Optimization with Sparsity,"* SIAM J. Optim., Vol. 17, No. 3, pp. 822–843.

[Len99] Lenstra (Jr.), Hendrik W., *"Finding Small Degree Factors of Lacunary Polynomials,"* Number Theory in Progress, Vol. 1 (Zakopane-Kóscielisko, 1997), pp. 267–276, de Gruyter, Berlin, 1999.

[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.

[MP98] Mourrain, Bernard and Pan, Victor Y. *"Asymptotic Acceleration of Solving Multivariate Polynomial Systems of Equations,"* Proc. ACM STOC 1998.

[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.

[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.

[S08] Safey El Din, Mohab, *"Computing the Global Optimum of a Multivariate Polynomial over the Reals,"* proceedings of ISSAC 2008 (Hagenberg, Austria), pp. 71–78, ACM, New York, 2008.

[Sal76] Salamin, E., *"Computation of π using arithmetic-geometric mean,"* Math. Comput., 30 (1976), pp. 565–570

[VY93] Vavasis, Steve and Ye, Yinyu, *"A primal dual accelerated interior point method whose running time depends only on A,"* Math. Progr. 74 (1996), pp. 79–120.