Optimizing n-variate $(n+k)$-nomials for small k

Philippe Pébay [1] J. Maurice Rojas [2] David C. Thompson [1]

Sandia National Laboratories, PO Box 969, MS 9159, Livermore, CA 94551, USA. [1]

TAMU 3368, Department of Mathematics, Texas A&M University, College Station, Texas 77843-3368, USA. [2]

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 quadratic in n and the logarithm of a certain condition number. For the special case of n-variate $(n+2)$-nomials with 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 Viro diagrams and $\mathcal{A}$-discriminants to real exponents. We also show that, for any fixed $\delta>0$, deciding whether the supremum of an n-variate $\left(n+n^{\delta}\right)$-nomial exceeds a given number is $\mathbf{NP}_{\mathbb{R}}$-complete.

1 Introduction and Main Results

Maximizing or minimizing polynomial functions is a central problem in science and engineering (see, e.g., [BG-V03,AM10]). Typically, the polynomials

Email addresses: `pppebay@sandia.gov` (Philippe Pébay), `rojas@math.tamu.edu` (J. Maurice Rojas), `dcthomp@sandia.gov` (David C. Thompson).

URL: `www.math.tamu.edu/~rojas` (J. Maurice Rojas).

[1] Supported by the US Dept. of Energy, Office of Defense Programs. Sandia is a multiprogram laboratory operated by Sandia Corp., a Lockheed Martin Company, for the US Dept. of Energy under contract DE-AC04-94AL85000.

[2] Partially supported by NSF grants DMS-0915245, DMS-0349309, DOE ASCR grant DE-SC0002505, and Sandia National Laboratories.

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 non-linear optimization 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 reveal two new families of problems that are $\mathbf{NP}_{\mathbb{R}}$-complete, i.e., the analogue of **NP**-complete for the BSS model over $\mathbb{R}$ [BSS89]. Our framework has both symbolic and numerical aspects in that (a) we deal with real number inputs and (b) our algorithms give either `Yes/No` answers that are always correct, or numerically approximate answers whose precision can be efficiently tuned.

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 $a_j \in \mathbb{R}^n$ 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 m-nomials 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.*[3] *We also call any $f \in \mathcal{F}^*_{n,m}$ an* honest n-variate m-nomial *(or* honestly n-variate*).* $\diamond$

For example, upon substituting $y_1 := x_1^2 x_2 x_3^7 x_4^3$, it is clear that 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, it is 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.

The main computational problems we address are the following.

Definition 1.2 *Let $\mathbb{R}_+$ denote the positive real numbers, and let* **SUP** *denote the problem of deciding, for a given* $(f, \lambda) \in \left(\bigcup_{n \in \mathbb{N}} \mathbb{R}[x^a \mid a \in \mathbb{R}^n] \right) \times (\mathbb{R} \cup \{+\infty\})$, *whether* $\sup_{x \in \mathbb{R}^n_+} f(x) \geq \lambda$ *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 $\mathcal{F}$.*

[3] A d-flat is simply a translated d-dimensional subspace. Note that our definitions of $\mathcal{F}_{n,m}$ and $\mathcal{F}^*_{n,m}$ here permit real coefficients *and* real exponents, unlike [BRS09] where the same notation included a restriction to integer coefficients and exponents.

When $\varepsilon > 0$ we say that $a \in \mathbb{R} \cup \{+\infty\}$ is a **strong** *$(1+\varepsilon)$-factor approximation of $b \in \mathbb{R} \cup \{+\infty\}$ when $a=0$, $a=+\infty$, or $\frac{a}{b} = \left[\frac{1}{1+\varepsilon}, 1+\varepsilon\right]$, according as $b=0$, $b=+\infty$, or $b \in \mathbb{R}^*$. Finally, we let* **FSUP** *(resp.* **FSUP**$(\mathcal{F})$*) denote the obvious functional analogue of* **SUP** *(resp.* **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 $(\bar{x}, \bar{\lambda}) \in (\mathbb{R}^n_+ \cup \{\texttt{``boundary''}\}) \times (\mathbb{R} \cup \{+\infty\})$ where*

1. *$\bar{\lambda}$ is a strong $(1+\varepsilon)$-factor of $\lambda^* := \sup_{x \in \mathbb{R}^n_+} f(x)$.*
2. *$\bar{x} = (\bar{x}_1, \ldots, \bar{x}_n)$ and, for all i, $\bar{x}_i$ is a strong $(1+\varepsilon)$-factor approximation of $x^*_i \in \mathbb{R}^n_+$ with $f(x^*_1, \ldots, x^*_n) = \lambda^* < +\infty$ (when λ^* is finite and attained by f in $\mathbb{R}^n_+$), or $\bar{x} = \texttt{``boundary''}$ (when λ^* is not attained in $\mathbb{R}^n_+$). $\diamond$*

Note that the output to **FSUP** always includes a true declaration of boundedness, or unboundedness, for f over $\mathbb{R}^n_+$.

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 a nonzero constant term. In what follows, our underlying notion of input size is clarified in Definition 2.1 of Section 2.1 below, and illustrated in Example 1.3 immediately following our first main theorem.

Theorem 1 *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* $\mathbf{NP}_{\mathbb{R}}$*-complete. More precisely:*

(0) Both **SUP**$\left(\bigcup_{n\in\mathbb{N}} \mathcal{F}^{**}_{n,n+1}\right)$ *and* **FSUP**$\left(\bigcup_{n\in\mathbb{N}} \mathcal{F}^{**}_{n,n+1}\right)$ *can be solved in time logarithmic in the input size, using a number of processors linear in the input size.*

(1) **SUP**$\left(\bigcup_{n\in\mathbb{N}} \mathcal{F}^{**}_{n,n+2}\right) \in \mathbf{P}_{\mathbb{R}}$. *Moreover, we can solve* **FSUP**$\left(\bigcup_{n\in\mathbb{N}} \mathcal{F}^{**}_{n,n+2}\right)$ *using only a number of arithmetic operations quadratic in the input size and $\log\log\frac{1}{\varepsilon}$.*

(2) For any fixed $\delta > 0$, **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 1.3 *Suppose $\varepsilon > 0$. A very special case of Assertion (1) of Theorem 1 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 by*
$$c_1 + c_2\left(x_1^D \cdots x_n^{D^n}\right) + c_3\left(x_1^{2D} \cdots x_n^{2^n D^n}\right) + \cdots + \left(c_{n+2} x_1^{(n+1)D} \cdots x_n^{(n+1)^n D^n}\right),$$
using a number of arithmetic operations linear in $n^2 \log|nD| + \log\log\frac{1}{\varepsilon}$ (see Definition 2.1 below). 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 $n^n D^n + \log\log\frac{1}{\varepsilon}$ instead, and then only under the assumption that $D \in \mathbb{N}$. Alternative approaches via semidefinite programming (see, e.g.,

[Par03,Las06,DN08,KM09]) also result in complexity bounds superlinear in $n^n D^n$ for our family of examples, and still require $D \in \mathbb{N}$. $\diamond$

For any input f with unbounded supremum, our algorithm evincing Theorem 1 (Algorithm 3.1 of Section 3.2) in fact gives additional information: a curve along which the values of f tend to $+\infty$.

Theorem 1 thus gives 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. Some of the intricacies of extending our techniques to n-variate $(n+k)$-nomials with $k \geq 3$ are detailed in [DRRS07,BHPR10] and Example 1.7 below. Theorem 1 is proved in Section 3.2 below, using an extension of tropical geometric ideas and $\mathcal{A}$-discriminant theory to real exponents (see Theorem 6 of Section 2.3 in particular).

Example 1.4 *Consider the trivariate pentanomial*

$$f(x) := c_1 + c_2 x_1^{999} + c_3 x_1^{73} x_3^{\sqrt{363}} + c_4 x_2^{2009} + c_5 x_1^{74} x_2^{108e} x_3,$$

with $c_1, \ldots, c_4 < 0$ and $c_5 > 0$. Theorem 8 of Section 2.4 below 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 6 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

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

normalized so that $b_5 > 0$. In particular, such a b can be easily computed through 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 positive numbers (and thus λ^*) can be efficiently approximated. Similarly, deciding whether λ^* exceeds a given λ reduces to a simple check of an inequality involving real powers of positive numbers.* $\diamond$

1.1 Origins, and Extensions of Viro Diagrams

Our main technical tool is a combinatorial/numerical characterization of when an $f \in \mathcal{F}^{**}_{n,n+2}$ has an unbounded supremum: Theorem 8 of Section 2.4. Perhaps the best way to understand this result is to recall one of its inspirations: Viro diagrams. First, 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 *coherent* iff it its maximal simplices are exactly the domains of linearity for some function that is convex, continuous, and piecewise linear on the convex hull of[4] $\mathcal{A}$.

Definition 1.5 *(See Proposition 5.2 and Theorem 5.6 of [GKZ94, Ch. 5, pp. 378–393].) Suppose $\mathcal{A} \subset \mathbb{Z}^n$ is finite and the convex hull of $\mathcal{A}$ has positive volume and boundary ∂A. Suppose also that $\mathcal{A}$ is equipped with a coherent triangulation Σ and a function $s : \mathcal{A} \longrightarrow \{\pm\}$ which we will call a* distribution of signs for $\mathcal{A}$. *We then define a piece-wise linear manifold — the* 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 $\mathcal{V}_{\mathcal{A}}(\Sigma, s) := \bigcup_{C \text{ an } n\text{-cell}} L_C \setminus \partial A$. When $\mathcal{A} = \mathrm{Supp}(f)$ and s is the corresponding sequence of coefficient signs, then we also call $\mathcal{V}_\Sigma(f) := \mathcal{V}_{\mathcal{A}}(\Sigma, s)$ the* Viro diagram of f. ⋄

Example 1.6 *Consider $f(x) := 1 - x_1 - x_2 + \frac{6}{5}(x_1^4 x_2 + x_1 x_2^4)$. Then* $\mathrm{Supp}(f) = \{(0,0),(1,0),(0,1),(1,4),(4,1)\}$ *and has convex hull a pentagon. So then there are exactly 5 coherent triangulations, yielding 5 possible Viro diagrams for f (drawn in thicker green lines):*

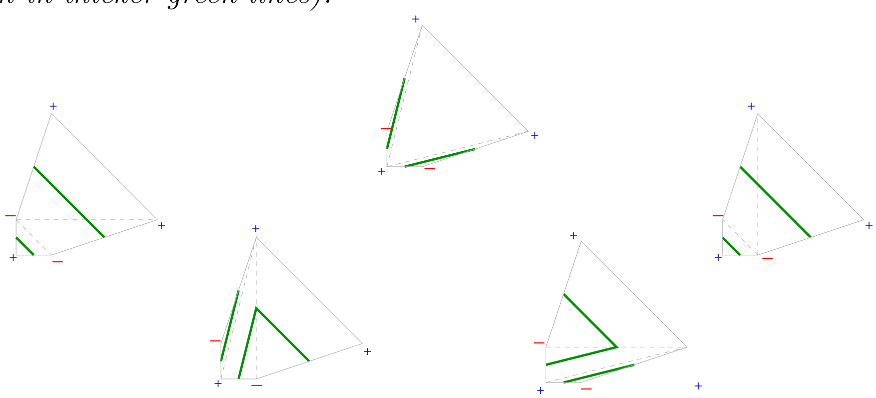

Note that all these diagrams have exactly 2 connected components, with each component isotopic to an open interval. Note also that our f here is a 2-variate $(2+3)$-nomial. ⋄

Viro originally used his diagrams to construct real algebraic varieties with prescribed topological behavior (see, e.g., [Vir84]). In particular, for *certain* f, one can *sometimes* find a triangulation Σ such that $\mathcal{V}_\Sigma(f)$ is isotopic to the set of positive zeroes of f. This construction is sometimes referred to as *patchworking*, but does not always yield all possible topologies for the positive zero set of a real polynomial with support $\mathcal{A}$.

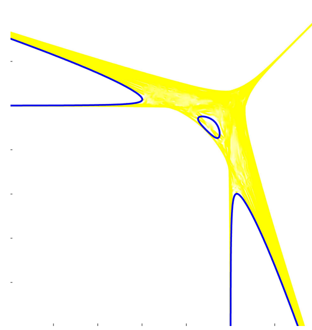

Example 1.7 *Our last example $f(x) := 1 - x_1 - x_2 + \frac{6}{5}(x_1^4 x_2 + x_1 x_2^4)$ is the*

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

simplest example we know where no *Viro diagram has topology matching that of the positive zero set of f: It is easily checked that the positive zero set of our f has exactly* 3 *connected components, one of which is isotopic to a circle. The last illustration above shows the image of the complex roots of f under the map $x \mapsto (\mathrm{Log}|x_1|, \mathrm{Log}|x_2|)$ (with image of the positive roots in blue).* ⋄

On the other hand, when $f \in \mathcal{F}^{**}_{n,n+2}$, it is possible to (a) extend the Viro diagram to real point sets and (b) find a triangulation Σ such that the positive zero set of f is isotopic to $\mathcal{V}_\Sigma(f)$. This is illustrated in 4 additional examples in Section 2.4, and will be pursued further in future work. To the best of our knowledge, Viro diagrams have only been used in the setting $\mathcal{A} \subset \mathbb{Z}^n$.

1.2 Related Work

Let us first recall the following basic inclusions of complexity classes from the BSS model over $\mathbb{R}$: $\mathbf{NC}^1_{\mathbb{R}} \subsetneqq \mathbf{P}_{\mathbb{R}} \subseteq \mathbf{NP}_{\mathbb{R}}$ [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, for any $k \in \mathbb{N} \cup \{0\}$, $\mathbf{NC}^k_{\mathbb{R}}$ is the family of real valued functions (with real inputs) computable by arithmetic circuits [5] with size polynomial in the input size and depth $O\left(\log^k(\text{Input Size})\right)$ (see [BCSS98, Ch. 18] for further discussion). The BSS model over $\mathbb{R}$ has proven quite useful for unifying computational complexity and numerical analysis. For instance, while the BSS model over $\mathbb{R}$ involves field operations with exact arithmetic, many recent results build upon this model to 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 in the Turing model (see, e.g., [Koi99,ABKM09]).

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 $\mathbf{NP}$-complete: deciding the existence of 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 3 below), now join this short list.

While sparsity has been profitably explored in the context of interpolation (see,

---

[5] 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 2.4 below.

e.g., [KY07,GLL09]) and factorization over number fields [Len99,KK06,AKS07], it has been mostly ignored in numerical analysis (for non-linear polynomials) and the study of the BSS model over $\mathbb{C}$ and $\mathbb{R}$. One important exception is work of Gabrielov and Vorobjov [GV04] that yields singly exponential complexity bounds for computing certain types of cell decompositions for a class of sub-analytic sets far more general than those we consider here. Nevertheless, there appear to have been 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 $m \geq 3$.

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].

Corollary 2

(0) We can find strong $(1+\varepsilon)$-factor approximations for the real roots of any trinomial in $\mathcal{F}_{1,3} \cap \mathbb{R}[x_1]$ in time polynomial in the input size and $\log\log\frac{1}{\varepsilon}$.

(1) $\mathbf{SUP}(\mathcal{F}^{**}_{1,4} \cap \mathbb{R}[x_1]) \in \mathbf{P}_{\mathbb{R}}$. *Moreover,* $\mathbf{FSUP}(\mathcal{F}^{**}_{1,4} \cap \mathbb{R}[x_1])$ *can be solved in time polynomial in the input size and* $\log\log\frac{1}{\varepsilon}$.

Assertion (1) appears to be new. Our corollary is proved in Section 3.3.

As for earlier complexity lower bounds for **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). An important precursor to our work here is thus the paper [BRS09], which deals with decision problems (i.e., `Yes/No` answers) and bit complexity (as opposed to arithmetic complexity). In fact, we can extend some of the complexity lower bounds from [BRS09] as follows. (See Section 3.2 for the proof.)

Definition 1.8 *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* $\mathcal{F}$. $\diamond$

Corollary 3 *For any* $\delta > 0$, *both* $\mathbf{FEAS}_{\mathbb{R}}\left(\bigcup_{\substack{n\in\mathbb{N}\\ 0<\delta'<\delta}} \mathcal{F}^{*}_{n,n+n^{\delta'}} \cap \mathbb{R}[x_1, \ldots, x_n]\right)$ *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* $\mathbf{NP}_{\mathbb{R}}$*-complete.*

2 Background

2.1 Input Size

Unlike integer exponents, real exponents can come arbitrarily close to each other. We will thus need to incorporate geometric information on the spread or proximity of the exponents of f when discussing the hardness of optimizing f. So we use the following notation 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 $m \geq n+1$, let us define $\hat{\mathcal{A}}$ to be the $(n+1) \times m$ matrix whose $j^{\underline{th}}$ column is the transpose of $\{1\} \times a_j$, and β_J the determinant of the submatrix of $\hat{\mathcal{A}}$ consisting of those columns of $\hat{\mathcal{A}}$ (written in increasing order of their index) 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*

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

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

For all computational problems in this paper, save for **SUP***, the size of an input f is defined to be $\log \mathcal{C}(f)$. For* **SUP***, the size of an instance (f, λ) is defined to be $\log\left(\max^*\left(3, |\lambda|, \frac{1}{|\lambda|}\right)\right) + \log \mathcal{C}(f)$.* ⋄

For $f \in \mathbb{Z}[x_1, \ldots, x_n]$ it is easy to show that $\log \mathcal{C}(f) = O\left((n+k)^{\min\{n+1,k-1\}} S(f)\right)$ 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 thus much faster than those with complexity polynomial in n and $\deg(f)$. The papers [Len99,KK06,AKS07,KY07,GLL09,BRS09] provide other important examples of algorithms with complexity polynomial in $S(f)$.

2.2 Tricks with Exponents

By substituting monomials in new variables it is easy to reduce n-variate $(n+k)$-nomials to a simpler canonical form.

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 $y^M := (y_1^{m_{1,1}} \cdots y_n^{m_{n,1}}, \ldots, y_1^{m_{1,n}} \cdots y_n^{m_{n,n}})$.* ⋄

Proposition 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* $m_U(x) := x^U$ *is an analytic automorphism of* $\mathbb{R}_+^n$, *and preserves smooth points and singular points of positive zero sets of analytic functions. In particular, for any* $f \in \mathcal{F}_{n,n+1}^{**}$ *we can compute* $c \in \mathbb{R}^*$ *and* $\ell \in \{0, \ldots, n\}$ *within* $\mathbf{NC}_{\mathbb{R}}^1$ *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)* $\bar{f}\left(\mathbb{R}_+^n\right) = f\left(\mathbb{R}_+^n\right)$.

The last assertion can be seen easily upon substituting $x = y^{A^{-1}}$ (where A is the matrix with columns the exponents of f) and rescaling the norms of the variables of f.

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.

Recall that the *affine hull* of a point set $\mathcal{A} = \{a_1, \ldots, a_m\} \subset \mathbb{R}^n$ is simply the set of all linear combinations of the form $\lambda_1 a_1 + \cdots + \lambda_m a_m$ where $\lambda_1, \ldots, \lambda_m \in \mathbb{R}$ satisfy $\lambda_1 + \cdots + \lambda_m = 1$. Such linear combinations are called *affine* linear combinations. In particular, $\mathcal{A}$ is said to be *affinely dependent* iff there is an affine linear combination satisfying $\lambda_1 a_1 + \cdots + \lambda_m a_m = \mathbf{O}$ (the zero vector).

Definition 2.3 *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* $\nabla_{\mathcal{A}} \subset \mathbb{P}_{\mathbb{C}}^{m-1}$ *— the* generalized $\mathcal{A}$-discriminant variety *— to be the closure of the set of all* $[c_1 : \cdots : c_m] \in \mathbb{P}_{\mathbb{C}}^{m-1}$ *such that* $g(y) = \sum_{i=1}^m c_i e^{a_i \cdot y}$ *has a degenerate root in* $\mathbb{C}^n$. *In particular, we call* g *an* n-variate exponential m-sum. ◇

We use the appelation "generalized" because $\mathcal{A}$-discriminants were originally developed by Gelfand, Kapranov, and Zelevinsky with $\mathcal{A} \subset \mathbb{Z}^n$ [GKZ94]. The more general setting $\mathcal{A} \subset \mathbb{C}^n$ is pursued further in [CR10].

Remark 5 *Note that by taking logs and exponentials, optimizing* n*-variate exponential* m*-sums over the real numbers is essentially the same as optimizing* n*-variate* m*-nomials over the positive numbers. To simplify our development, we will henceforth deal with exponential sums.* ◇

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

Definition 2.4 *We call* $\mathcal{A} \subset \mathbb{R}^n$ *a* (non-degenerate) circuit [6] *iff* $\mathcal{A}$ *is affinely*

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

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.* $\diamond$

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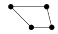 and 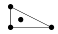 are circuits, but 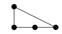 is a degenerate circuit. In general, for any degenerate circuit $\mathcal{A}$, the subset $\mathcal{B}$ named above is always unique.

Definition 2.5 *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 $\mathcal{A}$.* $\diamond$

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 6 *Suppose $\mathcal{A} = \{a_1, \ldots, a_{n+2}\} \subset \mathbb{R}^n$ is a non-degenerate circuit and, following the notation of Definition 2.1, let $b := (b_1, \ldots, b_{n+2})$ where $b_i := (-1)^i \beta_{\{1,\ldots,n+2\}\setminus\{i\}}$. 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 $[c_1 : \cdots : c_{n+2}] \in \mathbb{P}^{n+1}_{\mathbb{R}}$ satisfying the two conditions

(i) $\mathrm{sign}(c_1 b_1) = \cdots = \mathrm{sign}(c_{n+2} b_{n+2})$

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

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 $e^{a_i \cdot \zeta} = |b_i/c_i|$ for all i, and thus the real zero set of g has at most one degenerate point.

Theorem 6 is proved in Section 3 below. As a warm-up, it is worth noting that the range of certain exponential sums supported on circuits can be described quite explicitly. Let $\mathrm{Conv}\mathcal{A}$ denote the convex hull of $\mathcal{A}$.

Lemma 2.6 *Suppose $\mathcal{A} = \{a_1, \ldots, a_j\} \subset \mathbb{R}^n$ is a non-degenerate circuit with a_j in the relative interior of $\mathrm{Conv}\{a_1, \ldots, a_{j-1}\}$. Suppose also that $g(y) := \sum_{i=1}^{j} c_i e^{a_i \cdot y}$ with $c_1, \ldots, c_{j-1} < 0$ and $c_j > 0$. Finally, let $b = (b_1, \ldots, b_j) \in \mathbb{R}^j \setminus \{\mathbf{O}\}$ be any vector such that $b_1 a_1 + \cdots + b_j a_j = \mathbf{O}$ and $b_1 + \cdots + b_j = 0$. Then:*

(1) If $a_j = \mathbf{O}$ then $\sup_{y \in \mathbb{R}^n} g(y) = c_j - |b_j| \left| \prod_{i=1}^{j-1} (c_i/b_i)^{b_i} \right|^{|1/b_j|}$, and this value is attained exactly on an $(n+2-j)$-flat perpendicular to the affine hull of $\mathcal{A}$.

(2) *If* $a_r=\mathbf{O}$ *for some* $r<j$ *then* $\sup_{y\in\mathbb{R}^n} g(y)=c_r+|b_r|\left|\frac{\prod_{i\not\in\{j,r\}}(c_i/b_i)^{|b_i|}}{(c_j/b_j)^{|b_j|}}\right|^{-|1/b_r|}$, *and this value is attained exactly on an* $(n+2-j)$*-flat perpendicular to the affine hull of* $\mathcal{A}$.

(3) *If* $j<n+2$ *and the affine hull of* $\mathcal{A}$ *does not contain* $\mathbf{O}$ *then*
(a) $\left|\prod_{i=1}^{j-1}(c_i/b_i)^{|b_i|}\right|<(c_j/b_j)^{|b_j|}\Longrightarrow \sup_{y\in\mathbb{R}^n} g(y)=+\infty$, *and*
(b) $\left|\prod_{i=1}^{j-1}(c_i/b_i)^{|b_i|}\right|\geq(c_j/b_j)^{|b_j|}\Longrightarrow \sup_{y\in\mathbb{R}^n} g(y)\leq 0$.

Proof: First note that, by assumption, a_j is a convex (and positive) linear combination of $a_1,\ldots,a_{j-1}$. In other words, the b_i are all nonzero, and the signs of $b_1,\ldots,b_{j-1}$ are all identical and opposite to that of b_j. Since Assertions (1)–(3) are clearly invariant under the sign flip $b\mapsto -b$, let us henceforth assume that $b_1,\ldots,b_{j-1}>0>b_j$ to simplify matters.

Now recall the *Weighted Arithmetic-Geometric Inequality (AGI)* [HLP88, Sec. 2.5, pp. 16–18]: For any nonnegative real numbers $w_1,u_1,\ldots,w_k,u_k\geq 0$ with $w_1+\cdots+w_k>0$, we have

$$\frac{w_1u_1+\cdots+w_ku_k}{w_1+\cdots+w_k}\geq\left(u_1^{w_1}\cdots u_k^{w_k}\right)^{1/(w_1+\cdots+w_k)},$$

with equality iff all u_i with $w_i>0$ are equal. Substituting $k:=j-1$, and $w_i=b_i$ and $u_i=\frac{-c_ie^{a_i\cdot y}}{b_i}$ for all $i\in\{1,\ldots,j-1\}$, we then easily obtain

$$(\star)\qquad -\left(c_1e^{a_1\cdot y}+\cdots+c_{j-1}e^{a_{j-1}\cdot y}\right)\geq -b_je^{a_j\cdot y}\left|\prod_{i=1}^{j-1}(c_i/b_i)^{b_i}\right|^{-1/b_j}.$$

(Note also that the last inequality is invariant under scaling of the vector b since $b_1+\cdots+b_j=0$.) By the weighted AGI, equality holds in $(\star)$ iff $\frac{c_1e^{a_1\cdot y}}{b_1}=\cdots=\frac{c_{j-1}e^{a_{j-1}\cdot y}}{b_{j-1}}$. In particular, the latter equalities have an $(n+2-j)$-flat as their solution set: this follows immediately upon taking log and using the fact that $a_2-a_1,\ldots,a_{j-1}-a_1$ are linearly independent. Furthermore, this flat is clearly perpendicular to the affine hull of $\mathcal{A}$, by the orthogonality of left nullspaces and column spaces.

Assertion (1) then follows immediately from Inequality $(\star)$, and its conditions for equality, upon setting $a_j=\mathbf{O}$.

The proof of Assertion (2) follows a similar approach, but via Theorem 6 instead of the weighted AGI: we first observe that, by our hypotheses, g has supremum λ^* iff $g(y)-\lambda^*$ has a degenerate real zero set. (Indeed, via a slight variant of our application of the weighted AGI above, it is easy to see that g is bounded from above. That g is unbounded from below is easy to see by evaluating $e^{a_i\cdot y}$ along a suitable ray — a trick we'll expand on in the next paragraph.) In particular, the second assertion of Theorem 6 tells us that we must have $\left|\frac{c_r-\lambda^*}{b_r}\right|\prod_{i\neq r}\left|\frac{c_i}{b_i}\right|^{b_i}=1$ and $c_r-\lambda^*<0$ (since $c_jb_j<0$). Solving for λ^* we then immediately obtain the stated formula for the supremum. The statement on where the supremum is attained follows almost identically as in Assertion (1), except that the exponential equalities come from the second assertion of Theorem 6 and involve the vectors $\{a_1,\ldots,a_j\}\setminus a_r$ instead.

Assertion (3) then follows easily from Assertion (1) and a geometric construction: By applying Assertion (1) to $g(y)/e^{a_j \cdot y}$ we obtain that

$$|c_j/b_j|^{-b_j} > \left|\prod_{i=1}^{j-1}(c_i/b_i)^{b_i}\right| \Longrightarrow g(y_0) = g_0 \text{ for some } y_0 \in \mathbb{R}^n \text{ and } g_0 > 0.$$

Now let v be any nonzero vector perpendicular to the affine hull of $\mathcal{A}$. (Such a vector must exist since $j < n+2 \Longrightarrow \dim \mathrm{Conv}\mathcal{A} < n$.) Clearly $a_i \cdot v$ is nonzero and constant (say, equal to μ) for all $i \in \{1, \ldots, j\}$. So, replacing v by $-v$ if necessary, we may assume $\mu > 0$. So then

$$g(y_0 + tv) = \left(\sum_{i=1}^{j} c_i e^{a_i \cdot (y_0 + tv)}\right) = \left(\sum_{i=1}^{j} c_i e^{a_i \cdot y_0} (e^{a_i \cdot v})^t\right) = e^{\mu t}\left(\sum_{i=1}^{j} c_i e^{a_i \cdot y_0}\right) = e^{\mu t} g_0,$$

which is an unbounded increasing function of t. So $\sup_{y \in \mathbb{R}^n} g(y) := +\infty$.

Assuming $|c_j/b_j|^{-b_j} \leq \left|\prod_{i=1}^{j-1}(c_i/b_i)^{b_i}\right|$ instead, we clearly have $g(y)/e^{a_j \cdot y}$ is non-positive for all $y \in \mathbb{R}^n$, using Assertion (1) once again. So we are done. ■

Remark 7 *If* **O** *lies in the affine hull of $\mathcal{A}$ but not in the convex hull of $\mathcal{A}$ then Assertion (3) still holds: one simply sets v to be the point in $\mathcal{A}$ of maximal norm and uses almost the proof as above. This strengthening may be of future use in the optimization of certain n-variate exponential $(n+3)$-sums.* $\diamond$

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.* ■

In particular, for any $a > 0$ and $b \in \mathbb{R}^*$, it is easy to show via the identity $a^b = e^{b \log a}$ that a $(1+\varepsilon)$-factor approximation of a^b can be computed using just $O\left(|\log a| + |\log b| + \log\log \frac{1}{\varepsilon}\right)$ 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, Sign Checks, and Generalized Viro Diagrams

Optimizing an $f \in \mathcal{F}^{**}_{n,n+1}$ will ultimately reduce to checking simple inequalities involving just the coefficients of f. The supremum 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 supremum can be a transcendental function of the coefficients.

To formalize the harder case, we will continue working in the realm 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 $\mathcal{F}^{**}_{n,m}$.

Theorem 8 *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 g so that $\mathcal{B} := \{a_1, \ldots, a_j\}$ is the unique non-degenerate sub-circuit of $\mathcal{A}$ and, if* $\mathrm{Conv}\mathcal{B}$ *is a simplex, a_j is in the relative interior of* $\mathrm{Conv}\mathcal{B}$*. Also let $b := (b_1, \ldots, b_{n+2})$ be the vector defined in Theorem 6. Then $\sup\limits_{y \in \mathbb{R}^n} g(y) = +\infty \iff$ one of the following 2 conditions holds:*

(1) $c_s > 0$ for some vertex a_s of $\mathrm{Conv}\mathcal{A}$ *not equal to* $\mathbf{O}$.
(2) $\mathrm{Conv}\mathcal{B}$ *is a simplex, $c_j > 0$, $c_i < 0$ for all $i < j$, $\mathbf{O} \notin \mathcal{B}$, and $\left|\prod_{i=1}^{j-1} (c_i/b_i)^{|b_i|}\right| < |c_j/b_j|^{|b_j|}$.*

Finally, if $\lambda^* := \sup\limits_{y \in \mathbb{R}^n} g(y) < +\infty$, let r be such that $a_r = \mathbf{O}$. Then either

(3) $\lambda^* = c_r$ and $c_i < 0$ for all $i \neq r$, or
(4) $\mathrm{Conv}\mathcal{B}$ *is a simplex, $\mathbf{O} \in \mathcal{B}$, $c_j > 0$, $c_i < 0$ for all $i < j$, and λ^* is either*

$$c_j - |b_j| \left|\prod_{i=1}^{j-1} (c_i/b_i)^{|b_i|}\right|^{|1/b_j|} \quad \text{or} \quad c_r + |b_r| \left|\frac{\prod_{i \notin \{j,r\}} (c_i/b_i)^{|b_i|}}{(c_j/b_j)^{|b_j|}}\right|^{-|1/b_r|},$$

according as $a_j = \mathbf{O}$ or not.

Note that Conditions (1) and (2) can not hold simultaneously, by virtue of their respective restrictions on the signs of coefficients c_i (and since a_j can not be a vertex if $\mathrm{Conv}\mathcal{B}$ is a simplex). Similarly, Conditions (3) and (4) can not hold simultaneously. While the 4 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 $\mathbb{R}^n$.

More precisely, for any $\mathcal{A} \subset \mathbb{R}^n$, Σ a triangulation of $\mathcal{A}$, s a distribution of signs for $\mathcal{A}$, and f any real n-variate m-nomial with support $\mathcal{A}$, we can mimic Definition 1.5 of Section 1.1 to define the *generalized* Viro diagrams $\mathcal{V}_{\mathcal{A}}(\Sigma, s)$ and $\mathcal{V}_{\Sigma}(f)$. Some representative examples for $n = 3$ are illustrated below: The right-hand illustrations show the graphs of explicit honest tetranomials $f(x_1, x_2)$, while the left-hand illustrations show the corresponding generalized Viro diagrams of $f(x_1, x_2) - x_3$. (A blue (resp. red) vertex correspond to a positive (resp. negative) monomial coefficient.)

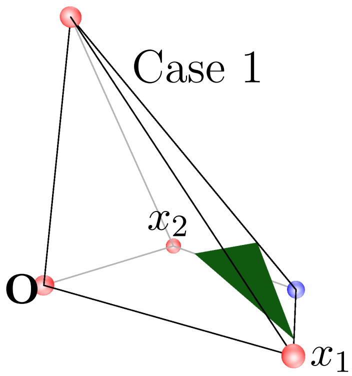
Case 1
x_2
O
x_1

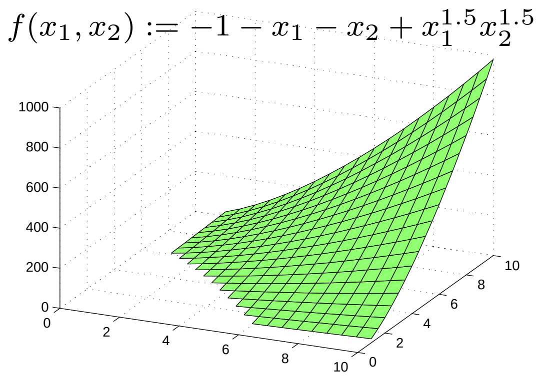
$f(x_1, x_2) := -1 - x_1 - x_2 + x_1^{1.5} x_2^{1.5}$

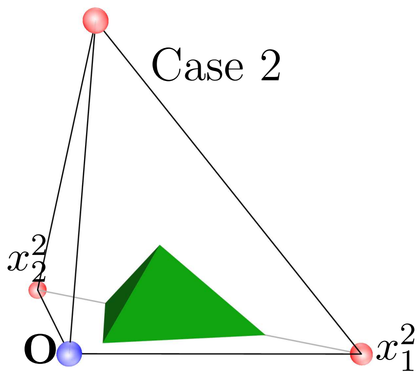
Case 2
x_2^2
O
x_1^2

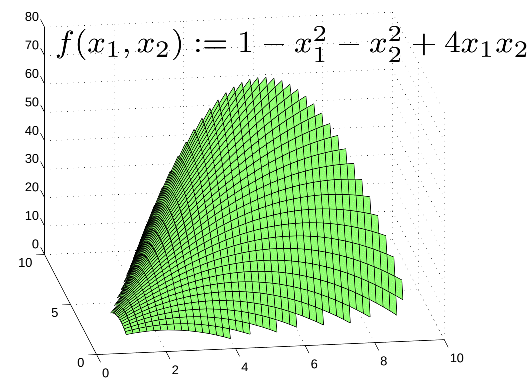
$f(x_1, x_2) := 1 - x_1^2 - x_2^2 + 4x_1 x_2$

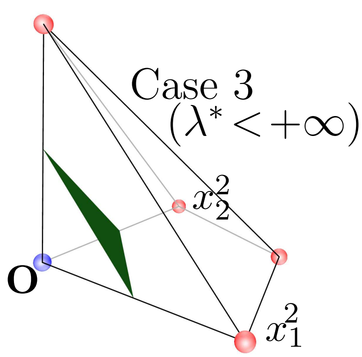
Case 3
$(\lambda^* < +\infty)$
x_2^2
O
x_1^2

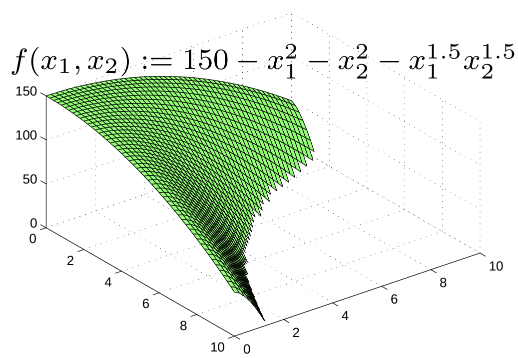
$f(x_1, x_2) := 150 - x_1^2 - x_2^2 - x_1^{1.5} x_2^{1.5}$

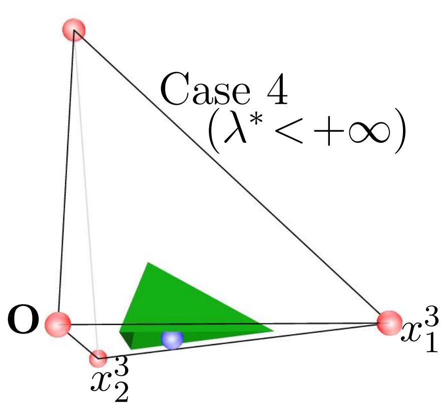
Case 4
$(\lambda^* < +\infty)$
O
x_1^3
x_2^3

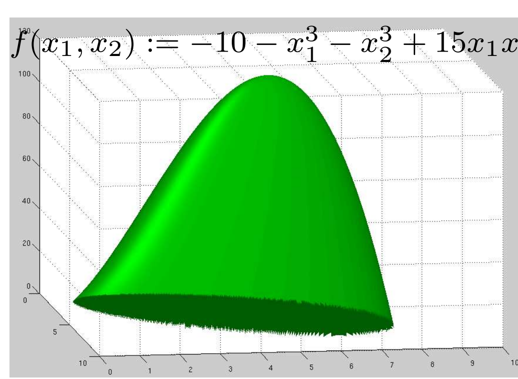
$f(x_1, x_2) := -10 - x_1^3 - x_2^3 + 15x_1 x_2$

Note that there is a hidden blue vertex in the left-hand illustration for Case 2. Also, to simplify the illustrations, we have not drawn the underlying triangulations. (The triangulations underlying Cases 1–3 involve exactly 2 tetrahedra each, while the triangulation for Case 4 involves exactly 3 tetrahedra meeting along the line segment from the top red vertex to the sole blue vertex.) Letting $g(y_1, y_2) := f(e^{y_1}, e^{y_2})$, the first four illustrations thus show how there can exist rays along which g increases without bound. Similarly, the last 4 illustrations respectively show cases where g either approaches a finite supremum as some $y_i \longrightarrow -\infty$ or g has a unique maximum in the real plane.

Proof of Theorem 8: Let $P := \text{Conv}\mathcal{A}$. We have 2 cases to consider.

($\sup_{y\in\mathbb{R}^n} g(y) = +\infty$): We must prove that Condition (1) or Condition (2) holds iff $\sup_{y\in\mathbb{R}^n} g(y) = +\infty$. Let us start with the "only if" direction.

First, if Condition (1) is true, then let v be any outer normal vector to the vertex a_s. The quantity $a_i \cdot v$, for $a_i \in \mathcal{A}$, then clearly has a unique positive maximum μ. So $g(tv) = c_s e^{\mu t} + o(e^{\mu t})$ is an unbounded increasing function of t, and thus $\sup_{y\in\mathbb{R}^n} g(y) = +\infty$.

On the other hand, if Condition (2) is true, then $\text{Conv}\mathcal{B}$ is a face of P not incident to $\mathbf{O}$, and $\mathbf{O}$ avoids the affine hull of $\mathcal{B}$ as well. So let us instead take v to be any outer normal vector to this face. By construction, we have that the quantity $a_i \cdot v$ is maximized (positively) exactly for $a_i \in \mathcal{B}$. Call this positive maximum μ once again. By Assertion (1) of Lemma 2.6, we can clearly find a $y_0 \in \mathbb{R}^n$ such that $\sum_{i=1}^{j} c_i e^{a_i \cdot y_0} = g_0$ for some $g_0 > 0$. Moreover, $g(y_0 + tv) = g_0 e^{\mu t} + o(e^{\mu t})$ and is thus an unbounded increasing function of t. So $\sup_{y\in\mathbb{R}^n} g(y) = +\infty$.

To prove the "if" direction it suffices to show that the failure of both Conditions (1) and (2) implies that g is bounded from above. Toward this end, observe that the failure of both Conditions (1) and (2) implies that c_i is negative for all vertices a_i of P (save possibly $\mathbf{O}$), and thus at least one of the following conditions holds: (a) $\text{Conv}\mathcal{B}$ is not a simplex, (b) $c_j < 0$, (c) $\mathbf{O} \in \mathcal{B}$, (d) $\left|\prod_{i=1}^{j-1} (c_i/b_i)^{|b_i|}\right| \geq |c_j/b_j|^{|b_j|}$. We will conclude by showing that the truth of any of these conditions implies that g is bounded from above.

If (a) holds then every point of $\mathcal{B}$ is a vertex of $\mathcal{A}$, which in turn implies that the only potentially positive coefficient of g is its constant term. So then g would be bounded from above. Similarly, (b) also implies that g is bounded from above. So let us assume henceforth that $\text{Conv}\mathcal{B}$ is a simplex and $c_j > 0$ (i.e., the failure of both (a) and (b)) to see the role of conditions (c) and (d).

If (c) fails, then $\mathbf{O}$ can not lie in the affine hull of $\mathcal{B}$, since $\mathbf{O} \in \mathcal{A}$. So then, the hypotheses of Assertion (3)(b) of Lemma 2.6 are satisfied for $\sum_{i=1}^{j} c_i e^{a_i \cdot y}$.

In other words, (d) implies that $\sum_{i=1}^{j} c_i e^{a_i \cdot y}$ is non-positive for all y, and thus g is bounded from above again.

On the other hand, should (c) hold true, then the hypotheses of Assertion (2) of Lemma 2.6 are satisfied for $\sum_{i=1}^{j} c_i e^{a_i \cdot y}$. In particular, $\sum_{i=1}^{j} c_i e^{a_i \cdot y}$ is bounded from above, and thus g is bounded from above yet again.

We have thus proved the "if" direction, by showing that the failure of both Conditions (1) and (2) implies that g is bounded from above. ■

($\sup_{y \in \mathbb{R}^n} g(y) < +\infty$): From our last proof, we know that g is bounded from above iff Conditions (1) and (2) both fail, and this in turn implies the truth of at least one of the conditions (a), (b), (c), or (d) mentioned above.

From our previous analysis, it is clear that Conditions (a) or (b) imply the supremum formula in Condition (3).

Similarly, our previous analysis reveals that Condition (c) implies that the formulae of Condition (4) hold, thanks to Assertions (1) and (2) of Lemma 2.6.

To conclude, the truth of Condition (d) provides a minor subtlety: Should Condition (c) fail, then, as we saw earlier, $\sum_{i=1}^{j} c_i e^{a_i \cdot y}$ has a negative supremum, attained at some value $y_0 \in \mathbb{R}^n$. Letting v be an inner normal to the face $\mathrm{Conv}\mathcal{B}$ of P it is then easily checked that $g(y_0 + tv) \longrightarrow c_r$ as $t \longrightarrow +\infty$. So the formula from Condition (3) holds.

Having covered all the necessary cases, we are done. ■

3 The Proofs of Our Main Results: Theorems 6 and 1, and Corollaries 3 and 2

3.1 The Proof of Theorem 6

In what follows, we let $Z_K(g)$ denote the zero set in K^n of g, for any field K and any function $g : K^n \longrightarrow K$.

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

$$\hat{\mathcal{A}}\begin{bmatrix} c_1 e^{a_1\cdot\zeta} \\ \vdots \\ c_{n+2}e^{a_{n+2}\cdot\zeta}\end{bmatrix} = \begin{bmatrix} 0 \\ \vdots \\ 0\end{bmatrix}.$$

In which case, $(c_1e^{a_1\cdot\zeta},\ldots,c_{n+2}e^{a_{n+2}\cdot\zeta})^T$ 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 $b_i\neq 0$ 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

$$\left(|c_1/b_1|^{b_1}\cdots|c_{n+2}/b_{n+2}|^{b_{n+2}}\right)\left(e^{(b_1a_1+\cdots+b_{n+2}a_{n+2})\cdot\mathbf{Re}(\zeta)}\right) = |\alpha|^{b_1+\cdots+b_{n+2}}.$$

Since $\hat{\mathcal{A}}(b_1,\ldots,b_{n+2})^T = \mathbf{O}$, we thus obtain $\prod\limits_{i=1}^{n+2}\left|\frac{c_i}{b_i}\right|^{b_i} = 1$. 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 Csanky's 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 $\zeta\in\mathbb{R}^n$. 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 $Z_{\mathbb{R}}\left(\sum_{i=1}^{n+2} c_i e^{a_i\cdot y}\right)$ has a degeneracy ζ, we directly obtain $e^{a_i\cdot\zeta} = \mathrm{sign}(b_1c_1)b_i/c_i$ for all i (and the constancy of $\mathrm{sign}(b_ic_i)$ 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 4, we can easily solve the preceding system for ζ by taking the logs of the coordinates of $\left(\frac{b_2c_1}{b_1c_2},\ldots,\frac{b_{n+1}c_1}{b_1c_{n+1}}\right)^{[a_2-a_1,\ldots,a_{n+1}-a_1]^{-1}}$. ■

Corollary 3 and Assertion (2) of Theorem 1: Since our underlying family of putative hard problems shrinks as δ decreases, it clearly suffices to prove the case $\delta < 1$. So let us 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 Chapter 4 of [BCSS98]) is one of the fundamental $\mathbf{NP}_{\mathbb{R}}$-complete problems.

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 $\mathrm{size}(f^{\pm}) = (16+o(1))\mathrm{size}(f)$. 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: 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 Proposition 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 merely need to decide the sign of $c_1 - \lambda$. ■

Remark 9 *Note that checking whether a given $f \in \mathcal{F}_{n,n+1}$ lies in $\mathcal{F}^*_{n,n+1}$ can be done within $\mathbf{NC}^2_\mathbb{R}$: one simply finds $d := \dim \mathrm{Supp}(f)$ in $\mathbf{NC}^2_\mathbb{R}$ by computing the rank of the matrix whose columns are $a_2 - a_1, \ldots, a_{n+1} - a_1$ (via the parallel algorithm of Csanky [Csa76]), and then checks whether $d = n$.* ◇

Assertion (1) of Theorem 1: We will first show how to effectively solve **FSUP**. Observe the following algorithm:

Algorithm 3.1
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 $(\bar{x}, \bar{\lambda}) \in \mathbb{R}^n_+ \times (\mathbb{R} \cup \{+\infty\})$ where*

1. *$\bar\lambda$ is a strong $(1+\varepsilon)$-factor of $\lambda^* := \sup_{x\in\mathbb{R}^n_+} f(x)$.*
2. *$\bar x = (\bar x_1, \ldots, \bar x_n)$ and, for all i, $\bar x_i$ is a strong $(1+\varepsilon)$-factor approximation of $x^*_i \in \mathbb{R}^n_+$ with $f(x^*) = \lambda^* < +\infty$ (when λ^* is attained by f in $\mathbb{R}^n_+$);* **or**, *should $\sup_{x\in\mathbb{R}^n_+} f(x)$ not be attained in $\mathbb{R}^n_+$, $\bar\lambda$ as in Statement (1) and a monomial curve along which $\sup_{x\in\mathbb{R}^n_+} f(x)$ is attained.*

Description:

(1) If $c_i > 0$ for some i with $a_i \neq \mathbf{O}$ a vertex of $\mathrm{Conv}\mathcal{A}$ then let v be any outer normal vector to a_i and output

"f `tends to` $+\infty$ `along a curve of the form` $\{ct^v\}_{t\rightarrow+\infty}$"

and `STOP`.

(2) Let $b := (b_1, \ldots, b_{n+2})$ be the vector defined in Theorem 6. If b or $-b$ has a unique negative coordinate b_j, and c_j is the unique positive 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 $i < j$]. Also let v be any outer normal vector to the face $\{a_1, \ldots, a_j\}$ of $\mathrm{Conv}\mathcal{A}$.

(b) If $\mathbf{O} \not\in \{a_1, \ldots, a_j\}$ and $\left|\prod_{i=1}^{j-1}(c_i/b_i)\right|^{|b_i|} < (c_j/b_j)^{|b_j|}$ then output
"$f \longrightarrow +\infty$ `along a curve of the form` $\{ct^v\}_{t\rightarrow+\infty}$" *and* `STOP`.

(c) If $\mathbf{O} \in \{a_1, \ldots, a_j\}$ then compute, via the Brent-Salamin Theorem, a strong $(1+\varepsilon)$-factor approximation $\bar\lambda$ of

$$c_j - |b_j| \left|\prod_{i=1}^{j-1}(c_i/b_i)^{|b_i|}\right|^{|1/b_j|} \quad \text{or} \quad c_r + |b_r| \left|\frac{\prod_{i\not\in\{j,r\}}(c_i/b_i)^{|b_i|}}{(c_j/b_j)^{|b_j|}}\right|^{-|1/b_r|},$$

according as $a_j = \mathbf{O}$ or $a_r = \mathbf{O}$ for some $r \neq j$. If $j < n+2$ then output

"f `tends to a supremum of` $\bar\lambda$ `along a curve of the form` $\{ct^v\}_{t\rightarrow+\infty}$."

and `STOP`.

(d) Compute, via Proposition 4 and the Brent-Salamin Theorem, $\bar x \in \mathbb{R}^n_+$ having coordinates that are respective strong $(1+\varepsilon)$-factor approximations of the unique solution to the binomial system

$$(x^{a_2-a_1}, \ldots, x^{a_{n+1}-a_1}) = \left(\frac{b_2c_1}{b_1c_2}, \ldots, \frac{b_{n+1}c_1}{b_1c_{n+1}}\right).$$

Then output "f `attains a supremum of` $\bar\lambda$ `at` $\bar x$." *and* `STOP`.

(3) Let v be any outer normal to the vertex $a_r = \mathbf{O}$ of $\mathrm{Conv}\mathcal{A}$. Then output

"f `approaches a supremum of` c_r `along a curve of the form` $\{ct^v\}_{t\rightarrow+\infty}$."

and `STOP`.

Our proof then reduces to proving correctness, and a suitable complexity bound, for Algorithm 3.1. In particular, correctness follows immediately from Theorem 8. So we now focus on a complexity analysis.

First note that $\mathrm{Conv}\mathcal{A}$ and j can be computed via $n+2$ determinants, simply by solving linear systems to determine which point of $\mathcal{A}$ lies in the relative interior of the others. Also, b can be computed in $\mathbf{NC}^2_{\mathbb{R}}$ thanks to Theorem 6. So Steps 1 and 3 (and the computation of any face normals) can clearly be done within $\mathbf{NC}^2_{\mathbb{R}}$. Moreover, the number of processors needed is $O(n^4)$ [Csa76].

For Step 2, the dominant complexity comes from Parts (b)–(d). These steps can be done by taking logarithms, checking the sign of linear combination of logarithms of positive real numbers, and approximating linear combination of logarithms of positive real numbers. By the Brent-Salamin Theorem (applied $O(n)$ times to approximate λ^* and n times to compute each coordinate of $\bar{x}$), the arithmetic complexity of our algorithm is $O\left(n\left(\log \mathcal{C}(f) + \log\log \frac{1}{\varepsilon}\right)\right)$, and we thus obtain our efficient solution of **FSUP**.

That $\mathbf{SUP}\left(\bigcup_{n\in\mathbb{N}} \mathcal{F}^{**}_{n,n+2}\right) \in \mathbf{P}_{\mathbb{R}}$ now follows directly: we merely need to compare λ against the formulae of our algorithm above. Since we are deciding inequalities, we can actually attain correct answers simply by using sufficient precision, and this can be attained within $\mathbf{P}_{\mathbb{R}}$ thanks to our formulae, our definition of $\mathcal{C}(f)$, and the Brent-Salamin Theorem. ■

Note that just as in Remark 9, checking whether a given $f \in \mathcal{F}_{n,n+2}$ lies in $\mathcal{F}^*_{n,n+2}$ can be done within $\mathbf{NC}^2_{\mathbb{R}}$ by computing $d = \dim \mathrm{ConvSupp}(f)$ efficiently.

3.3 The Proof of Corollary 2

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 strong $(1+\varepsilon)$-factor approximations 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 $f(x) = c_1 + c_2x^{a_2} + c_3x^{a_3} + c_4x^{a_4}$ with $0 < a_2 < a_3 < a_4$, 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 admit efficient strong $(1+\varepsilon)$-factor approximations. 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) $\sup_{x\in\mathbb{R}_+} f(x)$. So we are done. ■

Acknowledgements

We thank Eric Allender, Peter Bürgisser, Felipe Cucker, Johan Håstad, 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 generous hospitality while this paper was completed. Finally, we thank the referees for their valuable suggestions which helped greatly improve this paper.

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).

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

[AM10] Avendaño, Martin and Mortari, Daniele, *"A closed-form solution to the minimum $\Delta V^2_{\mathrm{tot}}$ Lambert's problem,"* Celest. Mech. Dyn. Astr. (2010), 106:25–37.

[BHPR10] Bastani, Osbert; Hillar, Chris; Popov, Dimitar; and Rojas, J. Maurice, *"Randomization, Sums of Squares, and Faster Real Root Counting for Tetranomials and Beyond,"* submitted for publication.

[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.

[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.

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

[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.

[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.

[CR10] 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).

[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

[Vir84] Viro, Oleg Ya., *"Gluing of plane real algebraic curves and constructions of curves of degrees* 6 *and* 7*,"* Topology (Leningrad, 1982), pp. 187–200, Lecture Notes in Math., 1060, Springer, Berlin, 1984.