

Toward accurate polynomial evaluation in rounded arithmetic

James Demmel*, Ioana Dumitriu[†] and Olga Holtz[‡]

August 17, 2005

Abstract

Given a multivariate real (or complex) polynomial p and a domain \mathcal{D} , we would like to decide whether an algorithm exists to evaluate $p(x)$ accurately for all $x \in \mathcal{D}$ using rounded real (or complex) arithmetic. Here “accurately” means with relative error less than 1, i.e., with some correct leading digits. The answer depends on the model of rounded arithmetic: We assume that for any arithmetic operator $op(a, b)$, for example $a+b$ or $a \cdot b$, its computed value is $op(a, b) \cdot (1+\delta)$, where $|\delta|$ is bounded by some constant ϵ where $0 < \epsilon \ll 1$, but δ is otherwise arbitrary. This model is the traditional one used to analyze the accuracy of floating point algorithms.

Our ultimate goal is to establish a decision procedure that, for any p and \mathcal{D} , either exhibits an accurate algorithm or proves that none exists. In contrast to the case where numbers are stored and manipulated as finite bit strings (e.g., as floating point numbers or rational numbers) we show that some polynomials p are impossible to evaluate accurately. The existence of an accurate algorithm will depend not just on p and \mathcal{D} , but on which arithmetic operators are available (perhaps beyond $+$, $-$, and \cdot), which constants are available to the algorithm (integers, algebraic numbers, ...), and whether branching is permitted in the algorithm. For floating point computation, our model can be used to identify which accurate operators beyond $+$, $-$ and \cdot (e.g., dot products, 3x3 determinants, ...) are necessary to evaluate a particular $p(x)$.

Toward this goal, we present necessary conditions on p for it to be accurately evaluable on open real or complex domains \mathcal{D} . We also give sufficient conditions, and describe progress toward a complete decision procedure. We do present a complete decision procedure for homogeneous polynomials p with integer coefficients, $\mathcal{D} = \mathbb{C}^n$, and using only the arithmetic operations $+$, $-$ and \cdot .

*Mathematics Department and CS Division, University of California, Berkeley, CA 94720. The author acknowledges the support of NSF under grants CCF-0444486, ACI-00090127, CNS-0325873 and of DOE under grant DE-FC02-01ER25478.

[†]Mathematics Department, University of California, Berkeley, CA 94720. The author acknowledges the support of the Miller Institute for Basic Research in Science.

[‡]Mathematics Department, University of California, Berkeley, CA 94720.

Contents

1	Introduction	2
2	Models of Algorithms and Related Work	6
2.1	Exact or Rounded Inputs	7
2.2	Finite Convergence	7
2.3	Basic Arithmetic Operations and Constants	7
2.4	Comparisons and Branching	8
2.5	Nondeterminism	9
2.6	Choice of Domain \mathcal{D}	9
2.7	Summary of Arithmetic and Algorithmic Models	9
2.8	Other Models of Error and Arithmetic	10
3	Evaluating positive polynomials accurately	14
4	Classical arithmetic	15
4.1	Necessity: real and complex	15
4.2	Sufficiency: the complex case	21
4.3	Toward a necessary and sufficient condition in the real case	22
4.3.1	Homogeneity	23
4.3.2	Dominance	24
4.3.3	Pruning	30
4.3.4	Sufficiency of evaluating dominant terms	34
4.3.5	Obstacles to full induction	36
5	“Black-box” arithmetic	37
5.1	Necessity: real and complex	37
5.2	Sufficiency: the complex case	42
6	Accurate Linear Algebra in Rounded Arithmetic	44
6.1	Vandermonde matrices and generalizations	47

1 Introduction

Let $x = (x_1, \dots, x_n)$ be a vector of real (or complex) numbers, let $p(x)$ denote a multivariate polynomial, and let \mathcal{D} be a subset of \mathbb{R}^n (or \mathbb{C}^n). We would ideally like to evaluate $p(x)$ *accurately* for all $x \in \mathcal{D}$, despite any rounding errors in arithmetic operations. The nature of the problem depends on how we measure accuracy, what kinds of rounding errors we consider, the class of polynomials $p(x)$, the domain \mathcal{D} , and what operations and constants our evaluation algorithms may use. Depending on these choices, an accurate algorithm for evaluating $p(x)$ may or may not exist. Our ultimate goal is a decision procedure that will either exhibit an algorithm that evaluates $p(x)$ accurately for all $x \in \mathcal{D}$, or else exhibit a proof that no such algorithm exists.

By *accuracy*, we mean that we compute an approximation $p_{comp}(x)$ to $p(x)$ that has small *relative* error: $|p(x) - p_{comp}(x)| \leq \eta|p(x)|$ for some desired $0 < \eta < 1$. In particular, $\eta < 1$ implies that $p(x) = 0$ if and only if $p_{comp}(x) = 0$. This requirement that p and p_{comp} define the same

variety will be crucial in our development. We justify this definition of accuracy in more detail in Section 2.

Our motivation for this work is two-fold. First, it is common for numerical analysts to seek accurate formulas for particularly important or common expressions. For example, in computational geometry and mesh generation, certain geometric predicates like “Is point x inside, on or outside circle C ?” are expressed as multivariate polynomials $p(\cdot)$ whose signs determine the answer; the correctness of the algorithms depends critically on the correctness of the sign, which is in turn guaranteed by having a relative error less than 1 in the value of $p(\cdot)$ [55]. We would like to automate the process of finding such formulas.

The second motivation is based on recent work of Koev and one of the authors [24] which identified several classes of structured matrices (e.g., Vandermonde, Cauchy, totally positive, certain discretized elliptic partial differential equations, ...) for which algorithms exist to accurately perform some (or all) computations from linear algebra: determinants, inversion, Gaussian elimination, computing singular values, computing eigenvalues, and so on. The proliferation of these classes of structured matrices led us to ask what common algebraic structure these matrix classes possess that made these accurate algorithms possible. This paper gives a partial answer to this question; see section 6.

Now we consider our model of rounded arithmetic. Let $op(\cdot)$ denote a *basic arithmetic operation*, for example $op(x, y) = x + y$ or $op(x, y, z) = x + y \cdot z$. Then we assume that the rounded value of $op(\cdot)$, which we denote $rnd(op(\cdot))$, satisfies

$$rnd(op(\cdot)) = op(\cdot)(1 + \delta) \tag{1}$$

where we call δ the *rounding error*. We assume only that $|\delta|$ is tiny, $|\delta| \leq \epsilon$, where $0 < \epsilon < 1$ and typically $\epsilon \ll 1$; otherwise δ is an arbitrary real (or complex) number. The constant ϵ is called the *machine precision*, by analogy to floating point computation, since this model is the traditional one used to analyze the accuracy of floating point algorithms [34, 63].

To illustrate the obstacles to accurate evaluation that this model poses, consider evaluating $p(x) = x_1 + x_2 + x_3$ in the most straightforward way: add (and round) x_1 and x_2 , and then add (and round) x_3 . If we let δ_1 be the first rounding error and δ_2 be the second rounding error, we get the computed value $p_{comp}(x) = ((x_1 + x_2)(1 + \delta_1) + x_3)(1 + \delta_2)$. To see that this algorithm is not accurate, simply choose $x_1 = x_2 = 1$ and $x_3 = -2$ (so $p(x) = 0$) and $\delta_1 \neq 0$. Then $p_{comp}(1, 1, -2) = 2\delta_1(1 + \delta_2) \neq 0$, so the relative error is infinite. Indeed, it can be shown that there is an open set of x and of (δ_1, δ_2) where the relative error is large, so that this loss of accuracy occurs on a “large” set. We will see that unless $x_1 + x_2 + x_3$ is itself a basic arithmetic operation, or unless the variety $\{x_1 + x_2 + x_3 = 0\}$ is otherwise constructible from varieties derived from basic operations as described in Theorem 5.16, then *no* algorithm exists to evaluate $x_1 + x_2 + x_3$ accurately for all arguments.

In contrast, if we were to assume that the x_i and coefficients of p were given as exact rational numbers (e.g., as floating point numbers), then by performing integer arithmetic with sufficiently large integers it would clearly be a straightforward matter to evaluate any $p(x)$ as an exact rational number. (One could also use floating point arithmetic to accomplish this, see Sections 2.3 and 2.8.) In other words, accurate evaluation is always possible, and the only question is cost. Our model addresses this by identifying which composite operations have to be provided with high precision in order to evaluate $p(x)$ accurately.

We give some examples to illustrate our results. Consider the family of homogeneous polyno-

mials

$$M_{jk}(x) = j \cdot x_3^6 + x_1^2 \cdot x_2^2 \cdot (j \cdot x_1^2 + j \cdot x_2^2 - k \cdot x_3^2)$$

where j and k are positive integers, $\mathcal{D} = \mathbb{R}^n$, and we allow only addition, subtraction and multiplication of two arguments as basic arithmetic operations, along with comparisons and branching.

- When $k/j < 3$, $M_{jk}(x)$ is *positive definite*, i.e., zero only at the origin and positive elsewhere. This will mean that $M_{jk}(x)$ is easy to evaluate accurately using a simple method discussed in Section 3.
- When $k/j > 3$, then we will show that $M_{jk}(x)$ cannot be evaluated accurately by *any* algorithm using only addition, subtraction and multiplication of two arguments. This will follow from a simple necessary condition on the real variety $V_{\mathbb{R}}(M_{jk})$, the set of real x where $M_{jk}(x) = 0$, see Theorem 4.8. This theorem requires the real variety $V_{\mathbb{R}}(p)$ (or the complex variety $V_{\mathbb{C}}(p)$) to lie in a certain explicitly given finite set of varieties called *allowable varieties* in order to be able to evaluate $p(x)$ accurately in real arithmetic (or in complex arithmetic, resp.).
- When $k/j = 3$, i.e., on the boundary between the above two cases, $M_{jk}(x)$ is a multiple of the Motzkin polynomial [53]. Its real variety $V_{\mathbb{R}}(M_{jk}) = \{x : |x_1| = |x_2| = |x_3|\}$ satisfies the necessary condition of Theorem 4.8, and the simplest accurate algorithm to evaluate it that we know is shown (in part) below:

$$\begin{aligned}
& \text{if} && |x_1 - x_3| \leq |x_1 + x_3| \wedge |x_2 - x_3| \leq |x_2 + x_3| && \text{then} \\
p = & && x_3^4 \cdot [4((x_1 - x_3)^2 + (x_2 - x_3)^2 + (x_1 - x_3)(x_2 - x_3))] \\
& && + x_3^3 \cdot [2(2(x_1 - x_3)^3 + 5(x_2 - x_3)(x_1 - x_3)^2 + 5(x_2 - x_3)^2(x_1 - x_3) + \\
& && \quad 2(x_2 - x_3)^3)] \\
& && + x_3^2 \cdot [(x_1 - x_3)^4 + 8(x_2 - x_3)(x_1 - x_3)^3 + 9(x_2 - x_3)^2(x_1 - x_3)^2 + \\
& && \quad 8(x_2 - x_3)^3(x_1 - x_3) + (x_2 - x_3)^4] \\
& && + x_3 \cdot [2(x_2 - x_3)(x_1 - x_3)((x_1 - x_3)^3 + 2(x_2 - x_3)(x_1 - x_3)^2 + \\
& && \quad 2(x_2 - x_3)^2(x_1 - x_3) + (x_2 - x_3)^3)] \\
& && + (x_2 - x_3)^2(x_1 - x_3)^2((x_1 - x_3)^2 + (x_2 - x_3)^2) \\
p = & && j \cdot p \\
& \text{else} && \dots \text{ 7 more analogous cases}
\end{aligned}$$

In general, for a Motzkin polynomial in n real variables ($n = 3$ above), the algorithm has 2^n separate cases. Just n tests and branches are needed to choose the correct case for any input x , so that the cost of running the algorithm is still just a polynomial function of n for any particular x .

In contrast to the real case, when $\mathcal{D} = \mathbb{C}^n$ then Theorem 4.8 will show that $M_{jk}(x)$ is not accurately evaluable using only addition, subtraction and multiplication.

If we still want to evaluate $M_{jk}(x)$ accurately in one of the cases where addition, subtraction and multiplication alone do not suffice, it is natural to ask which composite or “black-box” operations we would need to implement accurately to do so. Section 5 addresses this question.

The necessary condition for accurate evaluability of $p(x)$ in Theorem 4.8 depends only on the variety of $p(x)$. The next example shows that the variety alone is not enough to determine accurate evaluability, at least in the real case. Consider the two irreducible, homogeneous, degree $2d$, real polynomials

$$p_i(x) = (x_1^{2d} + x_2^{2d}) + (x_1^2 + x_2^2)(q_i(x_3, \dots, x_n))^2 \quad \text{for } i = 1, 2 \quad (2)$$

where $q_i(\cdot)$ is a homogeneous polynomial of degree $d - 1$. Both $p_1(x)$ and $p_2(x)$ have the same real variety $V_{\mathbb{R}}(p_1) = V_{\mathbb{R}}(p_2) = \{x : x_1 = x_2 = 0\}$, which is allowable, i.e., satisfies the necessary condition for accurate evaluability in Theorem 4.8. However, near $x_1 = x_2 = 0$, $p_i(x)$ is “dominated” by $(x_1^2 + x_2^2)(q_i(x_3, \dots, x_n))^2$, so accurate evaluability of $p_i(x)$ in turn depends on accurate evaluability of $q_i(x_3, \dots, x_n)$. Since $q_1(\cdot)$ may be accurately evaluable while $q_2(\cdot)$ is not, we see that $V_{\mathbb{R}}(p_i)$ alone cannot determine whether $p_i(x)$ is accurately evaluable. Applying the same principle to $q_i(\cdot)$, we see that any decision procedure must be recursive, expanding $p_i(x)$ near the components of its variety and so on. We show current progress toward a decision procedure in Section 4.3. In particular, Theorem 4.44 shows that, at least for algorithms without branching, being able to compute dominant terms of p (suitably defined) accurately on \mathbb{R}^n is a necessary condition for computing p accurately on \mathbb{R}^n . Furthermore, Theorem 4.46 shows that accurate evaluability of the dominant terms, along with branching, is sufficient to evaluate p accurately.

In contrast to the real case, Theorem 4.21 shows that for the complex case, with $\mathcal{D} = \mathbb{C}^n$, and using only addition, subtraction and multiplication of two arguments, a homogeneous polynomial $p(x)$ with integer coefficients is accurately evaluable if and only if it satisfies the necessary condition of Theorem 4.8. More concretely, $p(x)$ is accurately evaluable for all $x \in \mathbb{C}^n$ if and only if $p(x)$ can be completely factored into a product of factors of the form x_i , $x_i + x_j$ and $x_i - x_j$.

The results described so far from Section 4 consider only addition, subtraction, multiplication and (exact) negation (which we call *classical arithmetic*). Section 5 considers the same questions when accurate *black-box* operations beyond addition, subtraction and multiplication are permitted, such as fused-multiply-add [45], or indeed any enumerable collection of polynomials at all (e.g., dot products, 3x3 determinants, ...). The necessary condition on the variety of p from Theorem 4.8 is generalized to black-boxes in Theorem 5.16, and the sufficient conditions Theorem 4.21 in the complex case are generalized in Theorems 5.27 and 5.28.

The rest of this paper is organized as follows. Section 2 discusses further details of our algorithmic model, explains why it is a useful model of floating point computation, and otherwise justifies the choices we have made in this paper. Section 3 discusses the evaluation of positive polynomials. Section 4 discusses necessary conditions (for real and complex data) and sufficient conditions (for complex data) for accurate evaluability, when using only the classical arithmetic. Section 4.3 describes progress toward devising a decision procedure for accurate evaluability in the real case using classical arithmetic. Section 5 extends Section 4’s necessary conditions to arbitrary black-box arithmetic operations, and gives sufficient conditions in the complex case. Section 6 describes implications for accurate linear algebra on structured matrices.

2 Models of Algorithms and Related Work

Now we state more formally our decision question. We write the output of our algorithm as $p_{comp}(x, \delta)$, where $\delta = (\delta_1, \delta_2, \dots, \delta_k)$ is the vector of rounding errors made during the algorithm.

Definition 2.1. We say that $p_{comp}(x, \delta)$ is an accurate algorithm for the evaluation of $p(x)$ for $x \in \mathcal{D}$ if

$$\begin{aligned} \forall 0 < \eta < 1 & \quad \dots \text{ for any } \eta = \text{desired relative error} \\ \exists 0 < \epsilon < 1 & \quad \dots \text{ there is an } \epsilon = \text{machine precision} \\ \forall x \in \mathcal{D} & \quad \dots \text{ so that for all } x \text{ in the domain} \\ \forall |\delta_i| \leq \epsilon & \quad \dots \text{ and for all rounding errors bounded by } \epsilon \\ |p_{comp}(x, \delta) - p(x)| & \leq \eta \cdot |p(x)| \quad \dots \text{ the relative error is at most } \eta. \end{aligned}$$

Our ultimate goal is a decision procedure (a “compiler”) that takes $p(\cdot)$ and \mathcal{D} as input, and either produces an accurate algorithm p_{comp} (including how to choose the machine precision ϵ given the desired relative error η) or exhibits a proof that none exists.

To be more precise, we must say what our set of possible algorithms includes. The above decision question is apparently not Tarski-decidable [51, 58] despite its appearance, because we see no way to express “there exists an algorithm” in that format.

The basic decisions about algorithms that we make are as follows, with details given in the indicated sections:

Sec. 2.1: We insist that the inputs x are given exactly, rather than approximately.

Sec. 2.2: We insist that the algorithm compute the exact value of $p(x)$ in finitely many steps when all rounding errors $\delta = 0$. In particular, we exclude iterative algorithms which might produce an approximate value of $p(x)$ even when $\delta = 0$.

Sec. 2.3: We describe the basic arithmetic operations we consider, beyond addition, subtraction and multiplication. We also describe the constants available to our algorithms.

Sec. 2.4: We consider algorithms both with and without comparisons and branching, since this choice may change the set of polynomials that we can accurately evaluate.

Sec. 2.5: If the computed value of an operation depends only the values of its operands, i.e., if the same operands x and y of $op(x, y)$ always yield the same δ in $rnd(op(x, y)) = op(x, y) \cdot (1 + \delta)$, then we call our model *deterministic*, else it is *nondeterministic*. We show that comparisons and branching let a nondeterministic machine simulate a deterministic one, and subsequently restrict our investigation to the easier nondeterministic model.

Sec. 2.6: What domains of evaluation \mathcal{D} do we consider? In principle, any semialgebraic set \mathcal{D} is a possibility, but for simplicity we mostly consider open \mathcal{D} , especially $\mathcal{D} = \mathbb{R}^n$ or $\mathcal{D} = \mathbb{C}^n$. We point out issues in extending results to other \mathcal{D} .

Finally, Section 2.7 summarizes the axioms our model satisfies, and Section 2.8 compares our model to other models of arithmetic, and explains the advantages of our model.

2.1 Exact or Rounded Inputs

We must decide whether we assume that the arguments are given exactly [9, 14, 51, 58] or are known only approximately [15, 30, 39, 49]. Not knowing the input x exactly means that at best (i.e., in the absence of any further error) we could only hope to compute the exact value of $p(\hat{x})$ for some $\hat{x} \approx x$, an algorithmic property known as *backward stability* [18, 34]. Since we insist that zero outputs be computed exactly in order to have bounded relative error, this means there is no way to guarantee that $p(\hat{x}) = 0$ when $p(x) = 0$, for nonconstant p . This is true even for simple addition $x_1 + x_2$. So we insist on exact inputs in our model.

2.2 Finite Convergence

Do we consider algorithms that take a bounded amount of time for all inputs $x \in \mathcal{D}$, and return $p_{comp}(x, 0) = p(x)$, i.e., the exact answer when all rounding errors are zero? Or do we consider possibly iterative algorithms that might take arbitrarily long on some inputs to produce an adequately accurate answer? We consider only the former, because (1) it seems natural to use a finite algorithm to evaluate a finite object like a polynomial, (2) we have seen no situations where an iterative algorithm offers any advantage to obtaining guaranteed *relative* accuracy and (3) this lets us write any algorithm as a piecewise polynomial function and so use tools from algebraic geometry.

2.3 Basic Arithmetic Operations and Constants

What are the basic arithmetic operations? For most of the paper we consider addition, subtraction and multiplication of two arguments, since this is necessary and sufficient for polynomial evaluation in the absence of rounding error. Furthermore, we consider negation as a basic operation that is always exact (since this mimics all implementations of rounded arithmetic). Sometimes we will also use (rounded) multiplication by a constant $op(x) = c \cdot x$. We also show how to extend our results to include additional basic arithmetic operations like $op(x, y, z) = x + y \cdot z$. The motivations for considering such additional “black-box” operations are as follows:

1. By considering operations like $x + c$, $x - c$ and $c \cdot x$ for any c in a set C of constants, we may investigate how the choice of C affects accurate evaluability. For example, if C includes the roots of a polynomial like $p(x) = x^2 - 2$, then we can accurately evaluate $p(x)$ with the algorithm $(x - \sqrt{2}) \cdot (x + \sqrt{2})$, but otherwise it may be impossible. We note that having C include all algebraic numbers would in principle let us evaluate any univariate polynomial $p(x)$ accurately by using its factored form $p(x) = c \prod_{i=1}^d (x - r_i)$.

In the complex case, it is natural to consider multiplication by $\sqrt{-1}$ as an exact operation, since it only involves “swapping” and possibly negating the real and imaginary parts. We can accommodate this by introducing operations like $x + \sqrt{-1} \cdot y$ and $x - \sqrt{-1} \cdot y$.

The necessary conditions in Theorem 5.16 and sufficient conditions in Theorems 5.27 and 5.28 do not depend on how one chooses an operation $x - r_i$ from a possibly infinite set, just whether that operation exists in the set. On the other hand, a decision procedure must effectively choose that operation, so our decision procedures will restrict themselves to enumerable (typically finite!) sets of possible operations.¹

¹We could in principle deal with the set of all instructions $x - r$ for r an arbitrary algebraic number, because the algebraic numbers are enumerable.

2. Many computers now supply operations like $x+y \cdot z$ in hardware with the accuracy we demand (the *fused-multiply-add* instruction [45]). It is natural to ask how this operation extends the class of polynomials that we can accurately evaluate.
3. It is natural to build a library (in software or perhaps even hardware) containing several such accurate operations, and ask how much this extends the class of polynomials that can be evaluated accurately. This approach is taken in computational geometry, where the library of accurate operations is chosen to implement certain geometric predicates precisely (e.g., “is point x inside, outside or on circle C ?” written as a polynomial whose sign determines the answer). These precise geometric predicates are critical to performing reliable mesh generation [55].
4. A common technique for extending floating point precision is to simulate and manipulate extra precision numbers by representing a high precision number y as a sum $y = \sum_{i=1}^k y_i$ of numbers satisfying $|y_i| \gg |y_{i+1}|$, the idea being that each y_i represents (nearly) disjoint parts of the binary expansion of y (see [5, 6, 22, 50] and the references therein). This technique can be modeled by the correct choice of black-box operations as we now illustrate. Suppose we include the enumerable set of black-box operations $\sum_{i=1}^n p_i$, where n is any finite number, and each p_i is the product of 1 to d arguments. In other words, we include the accurate evaluation of arbitrary multivariate polynomials in $\mathbb{Z}[x]$ of degree at most d among our black-box operations. Then the following sequence of operations produces as accurate an approximation of any such polynomial

$$\sum_{i=1}^n p_i = y_1 + y_2 + \cdots + y_k$$

as desired:

$$\begin{aligned} y_1 &= \text{rnd}\left(\sum_{i=1}^n p_i\right) = (1 + \delta_1)\left(\sum_{i=1}^n p_i\right) \\ y_2 &= \text{rnd}\left(\sum_{i=1}^n p_i - y_1\right) = (1 + \delta_2)\left(\sum_{i=1}^n p_i - y_1\right) \\ &\dots \\ y_k &= \text{rnd}\left(\sum_{i=1}^n p_i - \sum_{j=1}^{k-1} y_j\right) = (1 + \delta_k)\left(\sum_{i=1}^n p_i - \sum_{j=1}^{k-1} y_j\right) \end{aligned}$$

Induction shows that

$$\sum_{j=1}^k y_j = \left(1 - (-1)^k \left(\prod_{j=1}^k \delta_j\right)\right) \cdot \left(\sum_{i=1}^n p_i\right)$$

so that $y = \sum_{j=1}^k y_j$ approximates the desired quantity with relative error at most ϵ^k . Despite this apparent power, our necessary conditions in Theorem 5.16 and Section 6.1 will still show limits on what can be evaluated accurately.

2.4 Comparisons and Branching

Are we permitted to do comparisons and then branch based on their results? Are comparisons *exact*, i.e., are the computed values of $x > y$, $x = y$ and $x < y$ (true or false) always correct for real x and

y ? (For complex x and y we consider only the comparison $x = y$.) We consider algorithms both without comparisons (in which case $p_{comp}(x, \delta)$ is simply a polynomial), and with exact comparisons and branching (in which case $p_{comp}(x, \delta)$ is a piecewise polynomial, on semialgebraic sets determined by inequalities among other polynomials in x and δ). We conjecture that using comparisons and branching strictly enlarges the set of polynomials that we can evaluate accurately.

We note that by comparing $x - r_i$ to zero for selected constants r_i , we could extract part of the bit representation of x . Since we are limiting ourselves to a finite number of operations, we could at most approximate x this way, and as stated in Section 2.1, this means we could not exploit this to get high relative accuracy near $p(x) = 0$. We note that the model of arithmetic in [15] excludes real \rightarrow integer conversion instructions.

2.5 Nondeterminism

As currently described, our model is *nondeterministic*, e.g., the rounded result of $1 + 1$ is not necessarily identical if it is performed more than once. This is certainly different behavior than the deterministic computers whose behavior we are modeling. However, it turns out that this is not a limitation, because we can always simulate a deterministic machine with a nondeterministic one using comparisons and branching. The idea is simple: The first addition instruction (say) records its input arguments and computed sum in a list. Every subsequent addition instruction compares its arguments to the ones in the list (which it can do exactly), and either just uses the precomputed sum if it finds them, or else does the addition and appends the results to the list. In other words, the existence (or nonexistence) of an accurate algorithm in a model with comparisons and branching does not depend on whether the machine is deterministic. So for simplicity, we will henceforth assume that our machines are nondeterministic.

2.6 Choice of Domain \mathcal{D}

As mentioned in the introduction, it seems natural to consider any semialgebraic set as a possible domain of evaluation for $p(x)$. While some choices, like $\mathcal{D} = \{x : p(x) = 0\}$ make evaluating $p(x)$ trivial, they beg the question of how one would know whether $x \in \mathcal{D}$. Similarly, if \mathcal{D} includes a discrete set of points, then $p(x)$ can be evaluated at these points by looking up the answers in a table. To avoid these pathologies, it may seem adequate restrict \mathcal{D} to be a sufficiently “fat” set, say open. But this still leads to interesting complications; for example the algorithm

$$p_{comp}(x, \delta) = ((x_1 + x_2)(1 + \delta_1) + x_3)(1 + \delta_2)$$

for $p(x) = x_1 + x_2 + x_3$, which is inaccurate on \mathbb{R}^n , is accurate on the open set $\{|x_1 + x_2| > 2|x_3|\}$, whose closure intersects the variety $V_{\mathbb{R}}(p)$ on $\{x_1 + x_2 = 0 \wedge x_3 = 0\}$.

In this paper we will mostly deal with open \mathcal{D} , especially $\mathcal{D} = \mathbb{R}^n$ or $\mathcal{D} = \mathbb{C}^n$, and comment on when our results apply to smaller \mathcal{D} .

2.7 Summary of Arithmetic and Algorithmic Models

We summarize the axioms our arithmetic and algorithms must satisfy. We start with the axioms all arithmetic operations and algorithms satisfy:

Exact Inputs. Our algorithm will be given the input exactly.

Finite Convergence. An accurate algorithm must, when all roundoff errors $\delta = 0$, compute the exact value of the polynomial in a finite number of steps.

Roundoff Model. Except for negation, which is always exact, the rounded value of any arithmetic operation $op(x_1, \dots, x_k)$ satisfies

$$rnd(op(x_1, \dots, x_k)) = op(x_1, \dots, x_k) \cdot (1 + \delta)$$

where δ is arbitrary number satisfying $|\delta| \leq \epsilon$, where ϵ is a nonzero value called the *machine precision*. If the data x is real (or complex), then δ is also real (resp. complex).

Nondeterminism. Every arithmetic operation produces an independent roundoff error δ , even if the arguments to different operations are identical.

Domain \mathcal{D} . Unless otherwise specified, the domain of evaluation \mathcal{D} is assumed to be all of \mathbb{R}^n (or all of \mathbb{C}^n).

We now list the alternative axioms our algorithms may satisfy. In each category, an algorithm must satisfy one set of axioms or the other.

Branching or Not. Some of our algorithms will permit exact comparisons of intermediate quantities ($<$, $=$ and $>$ for real data), and subsequent branching based on the result of the comparison. Other algorithms will not permit branching. In the complex case, we will see that branching does not matter (see Sections 4.2 and 5.2).

Classical or “black-box” operations. Some of our algorithms will use only “classical” arithmetic operations, namely addition, subtraction and multiplication. Others will use a enumerable set of arbitrary polynomial “black-box” operations, like $op(x, y, z) = x + y \cdot z$ or $op(x) = x - \sqrt{2}$, of our choice. In particular, we omit division.

2.8 Other Models of Error and Arithmetic

Our goal in this paper is to model rounded, finite precision computation, i.e., arithmetic with numbers represented in scientific notation, and rounded to their leading k digits, for some fixed k . It is natural to ask about models related to ours.

First, we point out some positive attributes of our model:

1. The model $rnd(op(a, b)) = op(a, b)(1 + \delta)$ has been the most widely used model for floating point error analysis [34] since the early papers of von Neumann [62], Turing [60] and Wilkinson [63].
2. The extension to include black-boxes includes widely used floating point techniques for extending the precision.
3. Though the model is for real (or complex) arithmetic, it can be efficiently simulated on a conventional Turing machine by using a simple variation of floating point numbers $m \cdot 2^e$, stored as the pair of integers (m, e) , where m is of fixed length, and $|e|$ grows as necessary. In particular, any sequence of n addition, subtraction, multiplication or division (by nonzero) operations can increase the largest exponent e by at most $O(n)$ bits, and so can be done in time polynomial in the input size. See [23] for further discussion. This is in contrast to repeated squaring in the BSS model [7] which can lead to exponential time simulations.

Models of arithmetic may be categorized according to several criteria (the references below are not exhaustive, but illustrative):

- Are numbers (and any errors) represented discretely (e.g., as bit strings such as floating point numbers) [23, 34, 63], or as a (real or complex) continuum [10, 13]?
- Is arithmetic exact [10, 8] or rounded [14, 15, 34, 63]? If it is rounded, is the error bounded in a relative sense [34], absolute sense [10], or something else [42, 17, 16] [34, Sec. 2.9]?
- In which of these metrics is the final error assessed?
- Is the input data exact [10] or considered “rounded” from its true value [12, 30, 39, 49, 52] (and if rounded, again how is the error bounded)?
- Do we want a “worst case” error analysis [34, 63], or by modeling rounding errors as random variables, a statistical analysis [61, 37, 56] [34, Sec. 2.8]? Does a condition number appear explicitly in the complexity of the problem [15]?

First we consider floating point arithmetic itself, i.e., where real numbers are represented by a pair of integers (m, n) representing the real number $m \cdot r^n$, where r is a fixed number called the *radix* (typically $r = 2$ or $r = 10$). Either by using one of many techniques in the literature for using an array (x_1, \dots, x_s) of floating point numbers to represent $x = \sum_{i=1}^s x_i$ to very high accuracy and to perform arithmetic on such high precision numbers (e.g., [5, 6, 50]), or by converting $m \cdot r^n$ to an exact rational number and performing exact rational arithmetic, one can clearly evaluate *any* polynomial $p(x)$ without error, and the only question is cost. In light of this, our results on classical vs black-box arithmetic can be interpreted as saying when such high precision techniques are necessary, and which black-box operations must be implemented this way, in order to evaluate p accurately.

Second, consider our goal of guaranteed high *relative* accuracy. One might propose that *absolute* accuracy is a more tractable goal, i.e., guaranteeing $|p_{comp}(x, \delta) - p(x)| \leq \eta$ instead of $|p_{comp}(x, \delta) - p(x)| \leq \eta|p(x)|$. However, we claim that as long as our basic arithmetic operations are defined to have bounded relative error ϵ , then trying to attain relative error in p_{comp} is the most natural goal.

Indeed, we claim that tiny absolute accuracy is impossible to attain for *any* nonconstant polynomial $p(x)$ when $\mathcal{D} = \mathbb{R}^n$ or $\mathcal{D} = \mathbb{C}^n$. For example, consider $p(x) = x_1 + x_2$, for which the obvious algorithm is $p_{comp}(x, \delta) = (x_1 + x_2)(1 + \delta)$. Thus the absolute error $|p_{comp}(x, \delta) - p(x)| = |x_1 + x_2|\delta \leq |x_1 + x_2|\epsilon$. This absolute error is at most η precisely when $|x_1 + x_2| \leq \eta/\epsilon$, i.e., for x in a diagonal strip in the (x_1, x_2) plane. For $p(x) = x_1 \cdot x_2$ we analogously get accuracy only for x in a region bounded by hyperbolas. In other words, even for the simplest possible polynomials that take one operation to evaluate, they cannot be evaluated to high absolute accuracy on most of $\mathcal{D} = \mathbb{R}^n$ or \mathbb{C}^n . The natural error model to consider when trying to attain low absolute error in $p(x)$ is to have low absolute error in the basic arithmetic operations, and this is indeed the approach taken in [15] (though as stated before, repeated squaring can lead to an exponential growth in the number of bits a real number represents [7]).

One could also consider more complicated error models, for example *mixed absolute/relative error*: $|p_{comp}(x, \delta) - p(x)| \leq \eta \cdot \max(|p(x)|, 1)$. Similar models have been used to model underflow error in floating point arithmetic [16]. A small mixed error implies that either the relative error or the absolute error must be small, and so may be easier to attain than either small absolute

error or small relative error alone. But we argue that, at least for the class of homogeneous polynomials evaluated on homogeneous \mathcal{D} , the question of whether $p(x)$ is accurately evaluable yields the same answer whether we mean accuracy in the relative sense or mixed sense. To see why, note that $x \in \mathcal{D}$ if and only if $\alpha x \in \mathcal{D}$ for any scalar α , since \mathcal{D} is homogeneous, and that $p(\alpha x) = \alpha^d p(x)$, where $d = \text{degree}(p)$. Thus for any nonzero $p(x)$, scaling x to αx will make $\eta \cdot \max(|p(\alpha x)|, 1) = \eta |p(\alpha x)|$ once α is large enough, i.e., relative error η must be attained. By results in Section 4.3, this will mean that $p_{\text{comp}}(x, \delta)$ must also be homogeneous in x of the same degree, i.e., $p_{\text{comp}}(\alpha x, \delta) = \alpha^d p_{\text{comp}}(x, \delta)$. Thus for any $x \in \mathcal{D}$ at which we can evaluate $p(x)$ with high mixed accuracy, we can choose α large enough so that

$$\alpha^d |p_{\text{comp}}(x, \delta) - p(x)| = |p_{\text{comp}}(\alpha x, \delta) - p(\alpha x)| \leq \eta \cdot \max(|p(\alpha x)|, 1) = \eta \cdot |p(\alpha x)| = \alpha^d \cdot \eta \cdot |p(x)|$$

implying that $p(\alpha x)$ can be evaluated with high relative accuracy for all α . In summary, changing our goal from relative accuracy to mixed relative/absolute accuracy will not change any of our results, for the case of homogeneous p and homogeneous \mathcal{D} .

Yet another model is to assume that the input x is given only approximately, instead of exactly as we assume. This corresponds to the approach taken in [30, 39, 49], in which one can imagine reading as many leading bits as desired of each input x_i can be read from an infinite tape, after which one tries to compute the answer using a conventional Turing machine model. This gives yet different results, since, for example, the difference $x_1 - x_2$ cannot be computed with small relative error in a bounded amount of time, since x_1 and x_2 may agree in arbitrarily many leading digits. Absolute error is more appropriate for this model.

It is worth commenting on why high accuracy of the sort we want is desirable in light of inevitable uncertainties in the inputs. Indeed, many numerical algorithms are successfully analyzed using *backward error analysis* [34, 18], where the computed results are shown to be the exact result for a slightly perturbed value of the input. This is the case, for example, for polynomial evaluation using Horner's rule where one shows that one gets the exact value of a polynomial at x but with slightly perturbed coefficients. Why is this not always accurate enough?

We already mentioned mesh generation [55], where the inputs are approximately known physical coordinates of some physical object to be triangulated, but where geometric predicates about the vertices defining the triangulation must be answered consistently; this means the signs of certain polynomials must be computed exactly, which is in turn guaranteed by guaranteeing any relative accuracy $\eta < 1$.

More generally, in many physical simulations, the parameters describing the physical system to be simulated are often known to only a few digits, if that many. Nonetheless, intermediate computations must be performed to much higher accuracy than the input data is known, for example to make sure the computed system conserves energy (which it should to high accuracy for the results to be meaningful, even if the initial conditions are uncertain).

Another example where high accuracy is important are the trigonometric functions: When x is very large and slightly uncertain, the value of $\sin x$ may be completely uncertain. Still, we want the computed trigonometric functions to (nearly) satisfy identities like $\sin^2 x + \cos^2 x = 1$ and $\sin 2x = 2 \sin x \cos x$ so that we can reason about program correctness. Many other examples of this sort can be found in articles posted at [36].

In the spirit of backward error analysis, one could consider the polynomial p fixed, but settle for accurately computing $p(\hat{x})$ where \hat{x} differs from x by only a small relative change in each component x_i . This is not guaranteed by Horner's rule, which is equivalent to changing the polynomial p slightly

but not x . Would it be easier to compute $p(\hat{x})$ accurately than $p(x)$ itself? This is the case for some polynomials, like $x_1 + x_2 + x_3$ or $c_1x_2^2x_3^3 + c_2x_1^2x_3^3 + c_3x_1x_2^4$, where there is a unique x_i that we can associate with each monomial to “absorb” the rounding error from Horner’s rule. In particular, with Horner’s rule, the number of monomials in $p(x)$ may at most be equal to the number of x_i . In analogy to this paper, one could ask for a decision procedure to identify polynomials that permit accurate evaluation of $p(\hat{x})$ using any algorithm. This is a possible topic for future work.

Another possibility is to consider error probabilistically [34, Sec. 2.8]. This has been implemented in a practical system [61], where a program is automatically executed repeatedly with slightly different rounding errors made at each step in order to assess the distribution of the final error. This approach is criticized in [37] for improperly modeling the discrete, non-random behavior of roundoff, and for possibly invalidating (near) identities like $\sin 2x = 2 \sin x \cos x$ upon which correctness may depend.

In *smoothed analysis* [56], one considers complexity (or for us, relative error) by averaging over a Gaussian distribution around each input. For us, input could mean either the argument x of a fixed polynomial p , or the polynomial itself, or both. First consider the case of a fixed polynomial p with a randomly perturbed x . This case is analogous to the previous paragraph, because the inputs can be thought of as slightly perturbed before starting the algorithm. Indeed, one could imagine rounding the inputs slightly to nearby rational or floating point numbers, and then computing exactly. But in this case, it is easy to see that, at least for codimension 1 varieties of p , the “smoothed” relative error is finite or infinite precisely when the worst case relative error is finite or infinite. So smoothing does not change our basic analysis.² Now suppose one smooths over the polynomial p , i.e., over its coefficients. If we smooth using a Gaussian distribution, then as we will see, the genericity of “bad” p will make the smoothed relative error infinite for all polynomials. Changing the distribution from Gaussian to one with a small support would only distinguish between positive definite polynomials, the easy case discussed in section 3, and polynomials that are not positive definite.

In *interval arithmetic* [46, 47, 2] one represents each number by a floating point interval guaranteed to contain it. To do this one rounds interval endpoints “outward” to ensure that, for example, the sum $c = a + b$ of two intervals yields an interval c guaranteed to contain the sum of any two numbers in a and b . It is intuitive that if an interval algorithm existed to evaluate $p(x)$ for $x \in \mathcal{D}$ that always computed an interval whose width was small compared to the number of smallest magnitude in the interval, and if the algorithm obeyed the rules in Section 2.7, then it would satisfy our accuracy requirements. Conversely, one might conjecture that an algorithm accurate by our criteria would straightforwardly provide an accurate interval algorithm, where one would simply replace all arithmetic operation by interval operations. The issue of interpreting comparisons and branches using possibly overlapping intervals makes this question interesting, and a possible subject for future work.

Finally, many authors use condition numbers in their analysis of the complexity of solving certain problems. This is classical in numerical analysis [34]; more recent references are [12, 15, 13]. In this approach, one is willing to do more and more work to get an adequate answer as the condition number grows, perhaps without bound. Such a conditioning question appears in our approach, if we ask how small the machine precision ϵ must be as a function of the desired relative error η , as well as p , \mathcal{D} , and allowed operations. Computing this condition number (outside the easy case described in Section 3) is an open question.

²The logarithm of the relative error, like the logarithm of many condition numbers, does however have a finite average.

3 Evaluating positive polynomials accurately

Here we address the simpler case where the polynomial $p(x)$ to be evaluated has no zeros in the domain of evaluation \mathcal{D} . It turns out that we need more than this to guarantee accurate evaluability: we will require that $|p(x)|$ be bounded both above and below in an appropriate manner on \mathcal{D} .

We let $\bar{\mathcal{D}}$ denote the closure of \mathcal{D} .

Theorem 3.1. *Let $p_{comp}(x, \delta)$ be any algorithm for $p(x)$ satisfying $p_{comp}(x, 0) = p(x)$, i.e. it computes the right value in the absence of rounding error. Let $p_{min} \equiv \inf_{x \in \bar{\mathcal{D}}} |p(x)|$. Suppose $\bar{\mathcal{D}}$ is compact and $p_{min} > 0$. Then $p_{comp}(x, \delta)$ is an accurate algorithm for $p(x)$ on \mathcal{D} .*

Proof. Since the relative error on \mathcal{D} is $|p_{comp}(x, \delta) - p(x)|/|p(x)| \leq |p_{comp}(x, \delta) - p(x)|/p_{min}$, it suffices to show that the numerator approaches 0 uniformly as $\delta \rightarrow 0$. This follows by writing the value of $p_{comp}(x, \delta)$ along any branch of the algorithm as $p_{comp}(x, \delta) = p(x) + \sum_{\alpha > 0} p_\alpha(x) \delta^\alpha$, where $\alpha > 0$ is a multiindex with at least one component exceeding 0. By compactness of $\bar{\mathcal{D}}$, $|\sum_{\alpha > 0} p_\alpha(x) \delta^\alpha| \leq C \sum_{\alpha > 0} |\delta|^\alpha$ for some constant C , which goes to 0 uniformly as the upper bound ϵ on each $|\delta_i|$ goes to zero. \square

Next we consider domains \mathcal{D} whose closure is not compact. To see that merely requiring $p_{min} > 0$ is not enough, consider evaluating $p(x) = 1 + (x_1 + x_2 + x_3)^2$ on \mathbb{R}^3 . Intuitively, $p(x)$ can only be accurate if its “dominant term” $(x_1 + x_2 + x_3)^2$ is accurate, once it is large enough, and this is not possible using only addition, subtraction and multiplication. (These observations will be formalized in Sections 4.3 and 4, respectively.)

Instead, we consider a homogeneous polynomial $p(x)$ evaluated on a homogeneous \mathcal{D} , i.e. one where $x \in \mathcal{D}$ implies $\gamma x \in \mathcal{D}$ for any scalar γ . Even though such \mathcal{D} are unbounded, homogeneity of p will let us consider just the behavior of $p(x)$ on \mathcal{D} intersected with the unit ball S^{n-1} in \mathbb{R}^n (or S^{2n-1} in \mathbb{C}^n). On this intersection we can use the same compactness argument as above:

Theorem 3.2. *Let $p(x)$ be a homogeneous polynomial, let \mathcal{D} be a homogeneous domain, and let S denote the unit ball in \mathbb{R}^n (or \mathbb{C}^n). Let*

$$p_{min,homo} \equiv \inf_{x \in \bar{\mathcal{D}} \cap S} |p(x)|$$

Then $p(x)$ can be evaluated accurately if $p_{min,homo} > 0$.

Proof. We describe an algorithm $p_{comp}(x, \delta)$ for evaluating $p(x)$. There are many such algorithms, but we only describe a simple one. (Indeed, we will see that the set of all accurate algorithms for this situation can be characterized completely by Definition 4.26 and Lemma 4.27.) Write $p(x) = \sum_{\alpha} c_{\alpha} x^{\alpha}$, where α is a multiindex $(\alpha_1, \dots, \alpha_n)$, $x^{\alpha} \equiv x_1^{\alpha_1} \cdots x_n^{\alpha_n}$, and $c_{\alpha} \neq 0$ is a scalar. Homogeneity implies $|\alpha| = \sum_i \alpha_i$ is constant. Then the algorithm simply

1. computes each x^{α} term by repeated multiplication by x_i s,
2. computes each $c_{\alpha} x^{\alpha}$ either by multiplication by c_{α} or by repeated addition if c_{α} is an integer, and
3. sums the $c_{\alpha} x^{\alpha}$ terms.

Since each multiplication, addition and subtraction contributes a $(1 + \delta_i)$ term, it is easy to see that

$$p_{comp}(x, \delta) = \sum_{\alpha} c_{\alpha} x^{\alpha} \Delta_{\alpha}$$

where each Δ_{α} is the product of at most some number f of factors of the form $1 + \delta_i$.

Now let $\|x\|_2 = (\sum_i |x_i|^2)^{1/2}$, so $\hat{x} = x/\|x\|_2$ is in the unit ball S . Then the relative error may be bounded by

$$\begin{aligned} \left| \frac{p_{comp}(x, \delta) - p(x)}{p(x)} \right| &= \left| \frac{\sum_{\alpha} c_{\alpha} x^{\alpha} \Delta_{\alpha} - \sum_{\alpha} c_{\alpha} x^{\alpha}}{\sum_{\alpha} c_{\alpha} x^{\alpha}} \right| \\ &= \left| \frac{\sum_{\alpha} c_{\alpha} \hat{x}^{\alpha} (\Delta_{\alpha} - 1)}{\sum_{\alpha} c_{\alpha} \hat{x}^{\alpha}} \right| \\ &\leq \frac{\sum_{\alpha} |c_{\alpha}| \cdot |\Delta_{\alpha} - 1|}{p_{min}} \\ &\leq \frac{\sum_{\alpha} |c_{\alpha}| \cdot ((1 + \epsilon)^f - 1)}{p_{min}} \end{aligned}$$

which goes to zero uniformly in ϵ . □

4 Classical arithmetic

In this section we consider the simple or classical arithmetic over the real or complex fields, with the three basic operations $\{+, -, \cdot\}$, to which we add negation. The model of arithmetic is governed by the laws in Section 2.7. We remind the reader that this arithmetic model *does not allow* the use of constants.

In Section 4.1 we find a necessary condition for accurate evaluability over either field, and in Section 4.2 we prove that this condition is also sufficient for the complex case.

Throughout this section, we will make use of the following definition of allowability.

Definition 4.1. *Let p be a polynomial over \mathbb{R}^n or \mathbb{C}^n , with variety $V(p) := \{x : p(x) = 0\}$. We call $V(p)$ allowable if it can be represented as a union of intersections of sets of the form*

$$1. \quad Z_i = \{x : x_i = 0\}, \tag{3}$$

$$2. \quad S_{ij} = \{x : x_i + x_j = 0\}, \tag{4}$$

$$3. \quad D_{ij} = \{x : x_i - x_j = 0\}. \tag{5}$$

If $V(p)$ is not allowable, we call it unallowable.

Remark 4.2. *For a polynomial p , having an allowable variety $V(p)$ is obviously a Tarski-decidable property (following [58]), since the number of unions of intersections of hyperplanes (3)-(5) is finite.*

4.1 Necessity: real and complex

All the statements and proofs in this section work equally well for both the real and the complex case, and thus we may treat them together. At the end of the section we use the necessity condition to obtain a partial result relating to domains.

Definition 4.3. From now we will refer to the space of variables as $\mathcal{S} \in \{\mathbb{R}^n, \mathbb{C}^n\}$.

To state and prove the main result of this section, we need to introduce some additional notions and notation.

Definition 4.4. Given a polynomial p over \mathcal{S} with unallowable variety $V(p)$, consider all sets W that are finite intersections of allowable hyperplanes defined by (3), (4), (5), and subtract from $V(p)$ those W for which $W \subset V(p)$. We call the remaining subset of the variety points in general position and denote it by $G(p)$.

Remark 4.5. If $V(p)$ is not allowable, then from definition 4.4 it follows that $G(p) \neq \emptyset$. One may also think of points in $G(p)$ as “unallowable” or “problematic”, because, as we will see, we necessarily get large relative errors in their vicinity.

Definition 4.6. Given $x \in \mathcal{S}$, define the set $\text{Allow}(x)$ as the intersection of all allowable planes going through x :

$$\text{Allow}(x) := (\cap_{x \in Z_i} Z_i) \cap (\cap_{x \in S_{ij}} S_{ij}) \cap (\cap_{x \in D_{ij}} D_{ij}),$$

with the understanding that

$$\text{Allow}(x) := \mathcal{S} \quad \text{whenever} \quad x \notin Z_i, S_{ij}, D_{ij} \quad \text{for all} \quad i, j.$$

Note that $\text{Allow}(x)$ is a linear subspace of \mathcal{S} .

We will be interested in the sets $\text{Allow}(x)$ primarily when $x \in G(p)$. For such cases we make the following observation.

Remark 4.7. For each $x \in G(p)$, the set $\text{Allow}(x)$ is not a subset of $V(p)$:

$$\text{Allow}(x) \not\subseteq V(p),$$

which follows directly from the definition of $G(p)$.

We can now state the main result of this section, which is a necessity condition for the evaluability of polynomials over domains.

Theorem 4.8. Let p be a polynomial over a domain $\mathcal{D} \in \mathcal{S}$. Let $G(p)$ be the set of points in general position on the variety $V(p)$. If there exists $x \in \mathcal{D} \cap G(p)$ such that $\text{Allow}(x) \cap \text{Int}(\mathcal{D}) \neq \emptyset$, then p is not accurately evaluable on \mathcal{D} .

To prove Theorem 4.8, we need to recall the notion of Zariski topology (see, e.g., [35]).

Definition 4.9. A subset $Y \subseteq \mathbb{R}^n$ (or \mathbb{C}^n) is called a Zariski closed set if there a subset T of the polynomial ring $\mathbb{R}[x_1, \dots, x_n]$ (or $\mathbb{C}[x_1, \dots, x_n]$) such that Y is the variety of T : $Y = V(T)$. A complement of a Zariski closed set is said to be Zariski open. The class of Zariski open sets defines the Zariski topology on \mathcal{S} .

In this paper, we consider the Zariski topology not on \mathcal{S} , but on a hypercube centered at the origin in δ -space (the space in which the vector of error variables δ lies). This topology is defined in exactly the same fashion.

Note that a Zariski closed set has measure zero unless it is defined by the zero polynomial only; then the set is the whole space. In the coming proof we will deal with nonempty Zariski open sets,

which are all of full measure. Finally, it is worth noting that the Zariski sets we will work with are algorithm-dependent.

Finally, we represent any algorithm as in [11, 1] by a directed acyclic graph (DAG) with input nodes, branching nodes, and output nodes. For simplicity in dealing with negation (given that negation is an *exact* operation), we define a special type of edge which indicates that the value carried along the edge is negated. We call these special edges *dotted*, to distinguish them from the regular *solid* ones.

Every computational node has two inputs (which may both come from a single other computational node); depending on the source of these inputs we have computational nodes with inputs from two distinct nodes and computational nodes with inputs from the same node. The latter type correspond either to

1. doubling $((x, x) \overset{+}{\mapsto} 2x)$,
2. doubling and negating $((-x, -x) \overset{+}{\mapsto} -2x)$,
3. computing zero exactly $((-x, x) \overset{+}{\mapsto} 0, (-x, -x) \overset{-}{\mapsto} 0, \text{ or } (x, x) \overset{-}{\mapsto} 0)$,
4. squaring $((x, x) \overset{\cdot}{\mapsto} x^2 \text{ or } (-x, -x) \overset{\cdot}{\mapsto} x^2)$,
5. squaring and negating $((-x, x) \overset{\cdot}{\mapsto} -x^2)$.

All nodes are labeled by $(op(\cdot), \delta_i)$ with $op(\cdot)$ representing the operation that takes place at that node. It means that at each node, the algorithm takes in two inputs, executes the operation, and multiplies the result by $(1 + \delta_i)$.

Finally, for each branch, there is a single destination node, with one input and no output, whose input value is the result of the algorithm.

Throughout the rest of this section, unless specified, we consider only non-branching algorithms.

Definition 4.10. *For a given $x \in \mathcal{S}$, we say that a computational node N is of non-trivial type if its output is a nonzero polynomial in the variables δ when the algorithm is run on the given x and with symbolic δ s.*

Definition 4.11. *For a fixed x , let N be any non-trivial computational node in an algorithm. We denote by $L(N)$ (resp., $R(N)$) the set of computational nodes in the left (resp., right) subgraphs of N . If both inputs come from the same node, i.e. $L(N)$ and $R(N)$ overlap, we will only talk about $L(N)$.*

Definition 4.12. *For a given $\epsilon > 0$, we denote by H_ϵ the hypercube of edge length 2ϵ centered at the origin, in δ -space.*

We will need the following Proposition.

Proposition 4.13. *Given any algorithm, any $\epsilon > 0$, and a point $x \in G(p)$, there exists a Zariski open set Δ in H_ϵ such that no non-trivial computational node has a zero output on the input x for all $\delta \in \Delta$.*

Proof. The proof follows from the definition of the non-trivial computational node.

Since every non-trivial computational node outputs a non-trivial polynomial in δ , it follows that each non-trivial computational node is nonzero on a Zariski open set (corresponding to the output

polynomial in δ) in H_ϵ . Intersecting this finite number of Zariski open sets we obtain a Zariski open set which we denote by Δ ; for any $\delta \in \Delta$ the output of any non-trivial computational node is nonzero. \square

We can now state and prove the following crucial lemma.

Lemma 4.14. *For a given algorithm, any $x \in G(p)$, and $\epsilon > 0$, exactly one of the following holds:*

1. *there exists a Zariski open set $\Delta \subseteq H_\epsilon$ such that the value $p_{comp}(x)$ computed by the algorithm is not zero when the algorithm is run with source input x and $\delta \in \Delta$;*
2. *$p_{comp}(y, \delta) = 0$ for all $y \in \text{Allow}(x)$ and all δ in H_ϵ .*

Proof of Lemma 4.14. We recall that the algorithm can be represented as a DAG, as described in the paragraphs preceding Definition 4.10.

Fix a point $x \in G(p)$. Once x is fixed, the result of each computation is a polynomial expression in the δ s. Consider the Zariski open set Δ whose existence is guaranteed by Proposition 4.13. There are now two possibilities: either the output node is of non-trivial type, in which case $p_{comp}(x, \delta) \neq 0$ for all $\delta \in \Delta$, or the output node is not of non-trivial type, in which case $p_{comp}(x, \delta_0) = 0$ for some $\delta_0 \in \Delta$.

In the latter case the output of the computation is zero; we trace back this zero to its origin, by marking in descending order all computational nodes that produced a zero (and thus we get a set of paths in the DAG, all of whose nodes produced exact zeros). Note that we are not interested in *all* nodes that produced a 0; only those which are on paths of zeros to the output node.

We will examine the last occurrences of zeros on paths of marked vertices, i.e. the zeros that are farthest from the output on such paths.

Lemma 4.15. *The last zero on such a path must be either*

1. *a source;*
2. *the output of a node where $(-x, x) \stackrel{\perp}{\mapsto} 0$, $(-x, -x) \stackrel{\leftarrow}{\mapsto} 0$, or $(x, x) \stackrel{\rightarrow}{\mapsto} 0$ are performed;*
3. *the output of an addition or subtraction node with two nonzero source inputs.*

Proof of Lemma 4.15. Note that a nonzero non-source output will be a non-constant polynomial in the δ specific to that node.

Clearly the last zero output cannot happen at a multiplication node; we have thus to show that the last occurrence of a zero output cannot happen at an addition or subtraction node which has two nonzero inputs from different nodes, at least one of which is a non-source. We prove the last statement by reductio ad absurdum.

Assume we could have a zero output at a node N with two nonzero inputs, at least one of which is not a source. Let $R(N)$ and $L(N)$ be as in Definition 4.11. Let $\delta(L(N))$ and $\delta(R(N))$ be the sets of errors δ_i corresponding to the left, respectively the right subtrees of N .

By assumption, $\delta(R(N)) \cup \delta(L(N)) \neq \emptyset$ (since at least one of the two input nodes is a non-source). Let δ_l (δ_r) denote the δ associated to the left (right) input node of N . Then we claim that either $\delta_l \notin \delta(R(N))$ or $\delta_r \notin \delta(L(N))$. (There is also the possibility that one of the two input nodes is a source and does not have a δ , but in that case the argument in the next paragraph becomes trivial.)

Indeed, since each δ is specific to a node, if δ_l were in $\delta(R(N))$, there would be a path from the left input node to the right input node. Similarly, if δ_r were in $\delta(L(N))$, then there would be a path from the right input node of N to the left input node of N . So if both events were to happen at the same time, there would be a cycle in the DAG. This cannot happen, hence either $\delta_l \notin \delta(R(N))$ or $\delta_r \notin \delta(L(N))$.

Assume w.l.o.g. $\delta_l \notin \delta(R(N))$. Then the left input of N is a non-trivial polynomial in δ_l , while the right input does not depend on δ_l at all. Hence their sum or difference is still a non-trivial polynomial in δ_l . Contradiction. \square

Now that Lemma 4.15 has been proven, we can state the crucial fact of the proof of Lemma 4.14: *all last occurrences of a zero appear at nodes which either correspond to allowable constraints (i.e., zero sources, or sums and differences of sources), or are addition/subtraction nodes with both inputs from the same node, which always, on any source inputs, produce a zero.*

Take now any point $y \in \text{Allow}(x)$; then y produces the same chains of consecutive zeros constructed (marked) in Lemma 4.14 as x does, with errors given by $\delta_0 \in \Delta$. Indeed, any node on such a chain that has a zero output at x when the error variables are δ_0 can trace this zero back to an allowable constraint (which is satisfied by both x and y) or to an addition/subtraction node with both inputs from the same node; hence the node will also have a zero output at y with errors δ_0 . In particular, if $p_{\text{comp}}(x, \delta_0) = 0$ for $\delta_0 \in \Delta$, then $p_{\text{comp}}(y, \delta_0) = 0$. Moreover, changing δ_0 can only introduce additional zeros, but cannot eliminate zeros on the zero paths that we traced for x (by the choice of Δ). Therefore, $p_{\text{comp}}(y, \delta) = 0$ for all $y \in \text{Allow}(x)$ and $\delta \in H_\epsilon$. This completes the proof of Lemma 4.14. \square

Corollary 4.16. *Lemma 4.14 implies that, for any $\epsilon > 0$ and any $x \in G(p)$, exactly one of the following holds: the relative error of computation, $|p_{\text{comp}} - p|/|p|$, is either infinity at x for all δ in a Zariski open set or 1 at all points $y \in (\text{Allow}(x) \setminus V(p))$ and all $\delta \in H_\epsilon$.*

We now consider algorithms with or without branches.

Theorem 4.17. *Given a (branching or non-branching) algorithm with output function $p_{\text{comp}}(\cdot)$, $x \in G(p)$, and $\epsilon > 0$, then one of the following is true:*

1. *there exists a set Δ_1 of positive measure in H_ϵ such that $p_{\text{comp}}(x, \delta)$ is nonzero whenever the algorithm is run with errors $\delta \in \Delta_1$, or*
2. *there exists a set Δ_2 of positive measure in H_ϵ such that for every $\delta \in \Delta_2$, there exists a neighborhood $N_\delta(x)$ of x such that for every $y \in N_\delta(x) \cap (\text{Allow}(x) \setminus V(p))$, $p_{\text{comp}}(y, \delta) = 0$ when the algorithm is run with errors δ .*

Remark 4.18. *This implies that, on a set of positive measure in H_ϵ , the relative accuracy of any given algorithm is either ∞ or 1.*

Proof. With $p_{\text{comp}}(\cdot)$ the output function and x a fixed point in general position, we keep the δ s symbolic. Depending on the results of the comparisons, the algorithm splits into a finite number of non-branching algorithms, which all start in the same way (with the input nodes) and then differ in accordance with a finite set of polynomial constraints on the δ s and x s.

Some of these branches will be chosen by sets of δ s of measure zero; at least one of the branches will have to be chosen by a set of δ s of positive measure whose interior is nonempty (all constraints being polynomials). Call that branch B , and let the set of δ s that choose it be called Δ_B .

By Proposition 4.13, there exists a Zariski open set $\Delta \in H_\epsilon$ such that, for all $\delta \in \Delta$, no non-trivial node in the subgraph representing our branch B has a zero output. In particular, this includes all quantities computed for comparisons that define B . Let $\Delta_2 := \text{Int}(\Delta_B \cap \Delta)$, where Int denotes the interior of a set. By the choice of Δ_B and Δ , the obtained set Δ_2 is non-empty.

Suppose the algorithm is run with errors $\delta_0 \in \Delta_2$ and $p_{\text{comp}}(x, \delta_0) \neq 0$. Then, by continuity, there must be a neighborhood Δ_1 in the set Δ_2 on which the computation will still be directed to branch B and $p_{\text{comp}}(x, \cdot)$ will still be nonzero, so we are in Case 1.

Assume now that we are not in Case 1, i.e. there is no $\delta \in \Delta_2$ such that $p_{\text{comp}}(x, \delta) \neq 0$. In this case we show by contradiction that $p_{\text{comp}}(y, \delta) = 0$ for all $y \in \text{Allow}(x)$ if y is sufficiently close to x (since $\text{Allow}(x)$ is a linear subspace containing x , there exist points in $\text{Allow}(x)$ which are arbitrarily close to x), thus, that Case 2 must be fulfilled.

If this claim is not true, then there is no neighborhood $N_\delta(x)$ of x such that when $y \in N_\delta(x) \cap \text{Allow}(x)$, the algorithm is directed to branch B on δ . In that case, there must be a sequence $\{y_n\} \in (\text{Allow}(x) \setminus V(p))$ such that $y_n \rightarrow x$ and y_n is always directed elsewhere for this choice of δ . The reason for this is that $\text{Allow}(x)$ is a linear subspace which is *not* contained in $V(p)$; hence no neighborhood of x in $\text{Allow}(x)$ can be contained in $V(p)$, and then such a sequence y_n must exist.

Since there is a finite number of branches, we might as well assume that all y_n will be directed to the same branch B' for this δ , and that they split off at the same branching node (pigeonhole principle).

Now consider the branching node where the splitting occurs, and let $r(z, \delta)$ be the quantity to be compared to 0 at that node. Since we always go to B' with y_n but to B with y , it follows that we *necessarily must have* $r(y_n, \delta) \neq 0$ whereas $r(x, \delta) = 0$. On the other hand, until that splitting point the algorithm followed the same path with y_n and with x , computing with the same errors δ . Applying then case 2 of Lemma 4.14 (which can be read to state that any algorithm computing r , and obtaining $r(x, \delta) = 0$, will also obtain $r(y_n, \delta) = 0$), we get a contradiction.

This completes the proof of Theorem 4.17. \square

Corollary 4.19. *Let p be a polynomial over \mathcal{S} with unallowable variety $V(p)$. Choose any algorithm with output function $p_{\text{comp}}(\cdot)$, any point $x \in G(p)$, $\epsilon > 0$, and $\eta < 1$. Then there exists a set Δ_x of positive measure arbitrarily close to x and a set Δ of positive measure in H_ϵ , such that $|p_{\text{comp}} - p|/|p|$ is strictly larger than η when computed at a point $y \in \Delta_x$ using any vector of relative errors $\delta \in \Delta$.*

Proof. On symbolic input x and with symbolic δ , the algorithm will have m branches B_1, \dots, B_m that correspond to constraints yielding (semi-algebraic) sets of positive measure S_1, \dots, S_m in (x, δ) -space. Choose $x \in G(p)$, and let $(x, 0)$ be a point in (x, δ) -space.

1. If $(x, 0)$ is in $\text{Int}(S_i)$ (the interior of some region S_i), then by Lemma 4.14 and Corollary 4.16 there exists either
 - (a) a δ_0 in δ -space sufficiently small such that (x, δ_0) is in $\text{Int}(S_i)$ and $p_{\text{comp}}(x, \delta_0) \neq 0$. The relative error at (x, δ_0) is in this case ∞ , and (by continuity) there must be a small ball around (x, δ_0) which is still in $\text{Int}(S_i)$, on which the minimum relative error is arbitrarily large, certainly larger than 1;
 - (b) a δ_0 in δ -space sufficiently small and a $y \in \text{Allow}(x) \setminus V(p)$ sufficiently close to x such that (y, δ_0) is in $\text{Int}(S_i)$ and $p_{\text{comp}}(y, \delta_0) = 0$. In this case the relative error at (y, δ_0) is 1, and (by continuity) there must be a small ball around (y, δ_0) which is still in $\text{Int}(S_i)$, on which the relative error is strictly larger than our $\eta < 1$.

2. Otherwise, $(x, 0)$ must be on the boundary of some of the regions S_i ; assume w.l.o.g. that it is on the boundary of the regions S_1, \dots, S_l . In this case, we choose a small hyperdisk $B_{\tilde{\epsilon}}((x, 0))$ in the linear subspace (x, \cdot) such that $B_{\tilde{\epsilon}}((x, 0))$ intersects the closures of S_1, \dots, S_l (and no other S_i s). We can do this because the sets S_i are all semi-algebraic.

- (a) If there exists a δ_0 in δ -space such that $(x, \delta_0) \in B_{\tilde{\epsilon}}((x, 0))$ and $(x, \delta_0) \in \text{Int}(S_i)$ for some $i \in \{1, \dots, l\}$, then by the same argument as in case 1 we obtain a small ball included in $\text{Int}(S_i)$ on which the relative error is greater than η ;
- (b) Otherwise, if there exists a δ_0 such that $(x, \delta_0) \in B_{\tilde{\epsilon}}((x, 0))$ is on the boundary of some region S_i for which the local algorithm corresponding to it would yield $p_{\text{comp}}(x, \delta_0) \neq 0$, then (by continuity) there exists a small ball around (x, δ_0) such that the intersection of that small ball with S_i is of positive measure, and the relative error on that small ball as computed by the algorithm corresponding to S_i is greater than 1;
- (c) Finally, otherwise, choose some point $(x, \delta_1) \in B_{\tilde{\epsilon}}((x, 0))$, so that (x, δ_1) is on the boundary of a subset of regions $S \subset \{S_1, \dots, S_l\}$. We must have that $p_{\text{comp}}(x, \delta_1) = 0$ when computed using any of the algorithms that correspond to any $S_i \in S$.

Let now $B(x)$ be a small ball around x in x -space, and consider $\tilde{B}(x) \equiv B(x) \cap (\text{Allow}(x) \setminus V(p))$.

There exists some $y \in \tilde{B}(x)$, close enough to x , such that (y, δ_1) is either in the interior or on the boundary of some $S_k \in S$.

By Lemma 4.14, since we must necessarily have $p_{\text{comp}}(y, \delta_1) = 0$ as computed by the algorithm corresponding to S_k , it follows (by continuity) that there is a small ball around (y, δ_1) on which the relative error, when computed using the algorithm corresponding to S_k , is greater than η . The intersection of that small ball with S_k must have positive measure.

From the above analysis, it follows that there is always a set of positive measure, arbitrarily close to $(x, 0)$, on which the algorithm will produce a relative error larger than η . \square

Proof of Theorem 4.8. Follows immediately from Theorem 4.17 and Corollary 4.19. \square

Remark 4.20. Consider the polynomial $p(x, y) = (1 - xy)^2 + x^2$, whose variety is at infinity. We believe that Theorem 4.8 can be extended to show that polynomials like $p(x, y)$ cannot be evaluated accurately on \mathbb{R} ; this is future work.

4.2 Sufficiency: the complex case

Suppose we now restrict input values to be complex numbers and use the same algorithm types and the notion of accurate evaluability from the previous sections. By Theorem 4.8, for a polynomial p of n complex variables to be accurately evaluable over \mathbb{C}^n it is necessary that its variety $V(p) := \{z \in \mathbb{C}^n : p(z) = 0\}$ be allowable.

The goal of this section is that this condition is also sufficient, as stated in the following theorem.

Theorem 4.21. Let $p : \mathbb{C}^n \rightarrow \mathbb{C}$ be a polynomial with integer coefficients and zero constant term. Then p is accurately evaluable on $\mathcal{D} = \mathbb{C}^n$ if and only if the variety $V(p)$ is allowable.

To prove this we first investigate what allowable complex varieties can look like. We start by recalling a basic fact about complex polynomial varieties, which can for example be deduced from Theorem 3.7.4 in [59, page 53]. Let V denote any complex variety. To say that $\dim_{\mathbb{C}}(V) = k$ means that, for each $z \in V$ and each $\delta > 0$, there exists $w \in V \cap B(z, \delta)$ such that w has a V -neighborhood that is homeomorphic to a real $2k$ -dimensional ball.

Theorem 4.22. *Let p be a non-constant polynomial over \mathbb{C}^n . Then*

$$\dim_{\mathbb{C}}(V(p)) = n - 1.$$

Corollary 4.23. *Let $p : \mathbb{C}^n \rightarrow \mathbb{C}$ be a nonconstant polynomial whose variety $V(p)$ is allowable. Then $V(p)$ is a union of allowable hyperplanes.*

Proof. Suppose $V(p) = \cup_j S_j$, where each S_j is an intersection of the sets in Definition 4.1 and, for some j_0 , S_{j_0} is not a hyperplane but an irredundant intersection of hyperplanes. Let $z \in S_{j_0} \setminus \cup_{j \neq j_0} S_j$. Then, for some $\delta > 0$, $B(z, \delta) \cap V(p) \subset S_{j_0}$. Since $\dim_{\mathbb{C}}(S_{j_0}) < n - 1$, no point in $B(z, \delta) \cap V(p)$ has a $V(p)$ -neighborhood that is homeomorphic to a real $2(n - 1)$ -dimensional ball. Contradiction. \square

Corollary 4.24. *If $p : \mathbb{C}^n \rightarrow \mathbb{C}$ is a polynomial whose variety $V(p)$ is allowable, then it is a product $p = c \prod_j p_j$, where each p_j is a power of x_i , $(x_i - x_j)$, or $(x_i + x_j)$.*

Proof. By Corollary 4.23, the variety $V(p)$ is a union of allowable hyperplanes. Choose a hyperplane H in that union. If $H = Z_{j_0}$ for some J_0 , expand p into a Taylor series in x_{j_0} . If $H = D_{i_0 j_0}$ or $H = S_{i_0 j_0}$ for some i_0, j_0 , expand p into a Taylor series in $(x_{i_0} - x_{j_0})$ (or $(x_{i_0} + x_{j_0})$). In either case, in this expansion, the zeroth coefficient of p must be the zero polynomial in x_j , $j \neq j_0$ (or $j \notin \{i_0, j_0\}$). Hence there is a k such that $p(x) = x_{j_0}^k \tilde{p}(x)$ in the first case, or $p(x) = (x_{i_0} \pm x_{j_0})^k \tilde{p}(x)$ in the second/third one. In any case, we choose k maximal, so that the variety $V(\tilde{p})$ is the closure of the set $V(p) \setminus Z_{j_0}$ in the first case, or $V(p) \setminus D_{i_0 j_0}$ ($V(p) \setminus S_{i_0, j_0}$) in the second/third case. Then proceed by factoring \tilde{p} in the same fashion. \square

Proof of Theorem 4.21. By Corollary 4.24, $p = c \prod_j p_j$, with each p_j a power of x_k or $(x_k \pm x_l)$. It also follows that c must be an integer since all coefficients of p are integers.

Since each of the factors is accurately evaluable, and we can get any integer constant c in front of p by repeated addition (followed, if need be, by negation), which are again accurate operations, the algorithm that forms their product and then adds/negates to obtain c evaluates p accurately. \square

Remark 4.25. *From Theorem 4.21, it follows that only homogeneous polynomials are accurately evaluable over \mathbb{C}^n .*

4.3 Toward a necessary and sufficient condition in the real case

In this section we show that accurate evaluability of a polynomial over \mathbb{R}^n is ultimately related to accurate evaluability of its “dominant terms”. This latter notion is formally defined later in this section. Informally, it describes the terms of the polynomial that dominate the remaining terms in a particular semialgebraic set close to a particular component of its variety; thus it depends on how we “approach” the variety of a polynomial.

For reasons outlined in Section 3, we consider here only homogeneous polynomials. Furthermore, most of this section is devoted to non-branching algorithms, but we do need branching for

our statements at the end of the section. The reader will be alerted to any change in our basic assumptions.

Here is a short walk through this section:

- In Section **4.3.1. Homogeneity**, we discuss an expansion of the relative error $|p_{comp}(x, \delta) - p(x)|/|p(x)|$ as a function of x and δ , and prove a result about accurate evaluability of homogeneous polynomials that will be used in Section **4.3.3. Pruning**.
- In Section **4.3.2. Dominance**, we introduce the notion of dominance and present different ways of looking at an irreducible component of the variety $V(p)$ using various simple linear changes of variables. These changes of variables allow us to identify all the dominant terms of the polynomial, together with the “slices” of space where they dominate.
- In Section **4.3.3. Pruning**, we explain how to “prune” an algorithm to manufacture an algorithm that evaluates one of its dominant terms, and prove a necessary condition for the accurate evaluation of a homogeneous polynomial by a non-branching algorithm. Roughly speaking, this condition says that accurate evaluation of the dominant terms we identified in Section **4.3.2. Dominance**, is necessary.
- In Section **4.3.4. Sufficiency of evaluating dominant terms**, we identify a special collection of dominant terms, together with the slices of space where they dominate. If accurately evaluable by (branching or non-branching) algorithms, these dominant terms allow us to construct a branching algorithm for the evaluation of the polynomial over the entire space (Theorem 4.46). These are just some of the terms present in the statement of Theorem 4.44.

4.3.1 Homogeneity

We begin by establishing some basic facts about non-branching algorithms that evaluate homogeneous polynomials.

Definition 4.26. *We call an algorithm $p_{comp}(x, \delta)$ with error set δ for computing $p(x)$ homogeneous of degree d if*

1. *the final output is of degree d in x ;*
2. *no output of a computational node exceeds degree d in x ;*
3. *the output of every computational node is homogeneous in x .*

Lemma 4.27. *If $p(x)$ is a homogeneous polynomial of degree d and if a non-branching algorithm evaluates $p(x)$ accurately by computing $p_{comp}(x, \delta)$, the algorithm must itself be homogeneous of degree d .*

Proof. First note that the output of the algorithm must be of degree at least d in x , since $p_{comp}(x, \delta) = p(x)$ when $\delta = 0$. Let us now write the overall relative error as

$$rel_{err}(x, \delta) = \frac{p_{comp}(x, \delta) - p(x)}{p(x)} = \sum_{\alpha} \frac{p_{\alpha}(x)}{p(x)} \delta^{\alpha}$$

where α is a multi-index. If $p_{comp}(x, \delta)$ is accurate then $p_{\alpha}(x)/p(x)$ must be a bounded rational function on the domain $(\mathbb{R}^n, \text{ or in the homogeneous case the sphere } S^{(n-1)})$. This implies, in

particular, that the output cannot be of degree higher than d in x . So, Condition 1 of Definition 4.26 must be satisfied.

Now suppose Condition 3 of Definition 4.26 is violated. We would like to show that the final output is also inhomogeneous. We can assume without loss of generality that the algorithm does not contain nodes that do operations like $x - x$ or $0 \cdot x$ (these can be “pre-pruned” and replaced with a 0 source). There exists a highest node $(op(\cdot), \delta_i)$ whose output is not homogeneous. If it is the output node, we are done. Otherwise, look at the next node $(op(\cdot), \delta_j)$ on the path toward the output node. The output of $op(\cdot, \delta_j)$ is homogeneous. On the other hand, the output of $(op(\cdot), \delta_i)$ (which is one of the two inputs to $(op(\cdot), \delta_j)$) must be inhomogeneous in x and must contain a term $\delta_i r(x)$ with $r(x)$ an inhomogeneous polynomial in x .

If $op(\cdot, \delta_i)$ is the only input to $op(\cdot, \delta_j)$, then inhomogeneity will be present in both outputs, since neither doubling nor squaring can cancel it; contradiction. Otherwise there is another input to $op(\cdot, \delta_j)$ (call it $op(\cdot, \delta_k)$). The output of $op(\cdot, \delta_k)$ must therefore also be inhomogeneous to cancel the inhomogeneous $r(x)$. Since the DAG is acyclic, δ_i is not present in the output of $op(\cdot, \delta_k)$ or δ_k is not present in the output $op(\cdot, \delta_i)$. Without loss of generality, assume the former case. Then the term $\delta_i r(x)$ will create inhomogeneity in the output of $(op(\cdot), \delta_j)$, and hence $(op(\cdot), \delta_i)$ is *not* a highest node with inhomogeneous output, contradiction. Hence $p_{comp}(x, \delta)$ is not homogeneous in x , thus one of the $p_\alpha(x)$ ’s has to contain terms in x of higher or smaller degree than d .

Similarly, if Condition 2 of Definition 4.26 were violated, then for some δ s the final output would be a polynomial of higher degree in x , and that would also mean some $p_\alpha(x)$ would be of higher degree in x .

In either of these cases, if some $p_\alpha(x)$ contained terms of smaller degree than d , by scaling the variables appropriately and letting some of them go to 0, we would deduce that $p_\alpha(x)/p(x)$ could not be bounded. If some $p_\alpha(x)$ contained terms of higher degree than d , by scaling the variables appropriately and letting some of them go to ∞ , we would once again obtain that $p_\alpha(x)/p(x)$ could not be bounded. \square

This proof shows that an algorithm evaluates a homogeneous polynomial p accurately on \mathbb{R}^n if and only if each fraction p_α/p is bounded on \mathbb{R}^n . It also shows each p_α has to be homogeneous of the same degree as p . Therefore, each fraction p_α/p is bounded on \mathbb{R}^n if and only if it is bounded on the unit sphere $S^{(n-1)}$. We record this as a corollary.

Corollary 4.28. *A non-branching homogeneous algorithm is accurate on \mathbb{R}^n if and only if it is accurate on $S^{(n-1)}$.*

4.3.2 Dominance

Now we begin our description of “dominant terms” of a polynomial. Given a polynomial p with an allowable variety $V(p)$, let us fix an irreducible component of $V(p)$. Any such component is described by linear allowable constraints, which, after reordering variables, can be grouped into l groups as

$$x_1 = \cdots = x_{k_1} = 0, \quad x_{k_1+1} = \cdots = \pm x_{k_2}, \quad \dots, \quad x_{k_{l-1}+1} = \cdots = \pm x_{k_l}.$$

To consider terms of p that “dominate” in a neighborhood of that component, we will change variables to map any component of a variety to a set of the form

$$\tilde{x}_1 = \cdots = \tilde{x}_{k_1} = 0, \quad \tilde{x}_{k_1+2} = \cdots = \tilde{x}_{k_2} = 0, \quad \dots, \quad \tilde{x}_{k_{l-1}+2} = \cdots = \tilde{x}_{k_l} = 0. \quad (6)$$

The supresets \tilde{S} for this example are the set S itself together with the set $\{x : x_1 + x_2 = 0, x_3 = -x_4 = x_5\}$ and the set $\{x : x_1 - x_2 = 0, x_3 = -x_4 = x_5\}$.

Note that we can write the vector of new variables \tilde{x} as Cx where C is a matrix, so can label the change of variables by the matrix C .

Now let us consider components of the variety $V(p)$. We have seen that any given component of $V(p)$ can be put into the form $x_1 = x_2 = \dots = x_k = 0$ using a standard change of variables, provided $V(p)$ is allowable. (To avoid cumbersome notation, we renumber all the variables set to zero as x_1 through x_k for our discussion that follows. We will return to the original description to introduce the notion of pruning.)

Write the polynomial $p(x)$ in the form

$$p(x) = \sum_{\lambda \in \Lambda} c_\lambda x_{[1:k]}^\lambda q_\lambda(x_{[k+1:n]}), \quad (8)$$

where, almost following MATLAB notation, we write $x_{[1:k]} := (x_1, \dots, x_k)$, $x_{[k+1:n]} := (x_{k+1}, \dots, x_n)$. Also, we let Λ be the set of all multi-indices $\lambda := (\lambda_1, \dots, \lambda_k)$ occurring in the monomials of $p(x)$.

To determine all dominant terms associated with the component $x_1 = x_2 = \dots = x_k = 0$, consider the Newton polytope P of the polynomial p with respect to the variables x_1 through x_k only, i.e., the convex hull of the exponent vectors $\lambda \in \Lambda$ (see, e.g., [44, p. 71]). Next, consider the normal fan $N(P)$ of P (see [64, pp. 192–193]) consisting of the cones of all row vectors η whose dot products with $x \in P$ are maximal for x on a fixed face of P . That means that for every nonempty face F of P we take

$$N_F := \{\eta = (n_1, \dots, n_k) \in (\mathbb{R}^k)^* : F \subseteq \{x \in P : \eta x := \sum_{j=1}^k n_j x_j = \max_{y \in P} \eta y\}\}$$

and

$$N(P) := \{N_F : F \text{ is a face of } P\}.$$

Finally, consider the intersection of the negative of the normal fan $-N(P)$ and the nonnegative quadrant \mathbb{R}_+^k . This splits the first quadrant \mathbb{R}_+^k into several regions S_{Λ_j} according to which subsets Λ_j of exponents λ “dominate” close to the considered component of the variety $V(p)$, in the following sense:

Definition 4.31. *Let Λ_j be a subset of Λ that determines a face of the Newton polytope P of p such that the negative of its normal cone $-N(P)$ intersects $(\mathbb{R}^k)_+^*$ nontrivially (not only at the origin). Define $S_{\Lambda_j} \in (\mathbb{R}^k)_+^*$ to be the set of all nonnegative row vectors η such that*

$$\eta \lambda_1 = \eta \lambda_2 < \eta \lambda, \quad \forall \lambda_1, \lambda_2 \in \Lambda_j, \quad \text{and } \lambda \in \Lambda \setminus \Lambda_j.$$

Note that if x_1 through x_k are small, then the exponential change of variables $x_j \mapsto -\log |x_j|$ gives rise to a correspondence between the nonnegative part of $-N(P)$ and the space of original variables $x_{[1:k]}$. We map back the sets S_{Λ_j} into a neighborhood of 0 in \mathbb{R}^k by lifting.³

Definition 4.32. *Let $F_{\Lambda_j} \subseteq [-1, 1]^k$ be the set of all points $x_{[1:k]} \in \mathbb{R}^k$ such that*

$$\eta := (-\log |x_1|, \dots, -\log |x_k|) \in S_{\Lambda_j}.$$

³This is reminiscent of the concept of an amoeba introduced in [32].

Remark 4.33. For any j , the closure of F_{Λ_j} contains the origin in \mathbb{R}^k .

Remark 4.34. Given a point $x_{[1:k]} \in F_{\Lambda_j}$, and given $\eta = (n_1, n_2, \dots, n_k) \in S_{\Lambda_j}$, for any $t \in (0, 1)$, the vector $(x_1 t^{n_1}, \dots, x_k t^{n_k})$ is in F_{Λ_j} . Indeed, if $(-\log |x_1|, \dots, -\log |x_k|) \in S_{\Lambda_j}$, then so is $(-\log |x_1|, \dots, -\log |x_k|) - \log |t| \eta$, since all equalities and inequalities that define S_{Λ_j} will be preserved, the latter because $\log |t| < 0$.

Example 4.35. Consider the following polynomial

$$p(x_1, x_2, x_3) = x_2^8 x_3^{12} + x_1^2 x_2^2 x_3^{16} + x_1^8 x_3^{12} + x_1^6 x_2^{14} + x_1^{10} x_2^6 x_3^4.$$

We show below the Newton polytope P of p with respect to the variables x_1, x_2 , its normal fan $N(P)$, the intersection $-N(P) \cap \mathbb{R}_+^2$, the regions S_{Λ_j} , and the regions F_{Λ_j} .




Figure 1. A Newton polytope P .




Figure 2. Its normal fan $N(P)$.




Figure 3. The intersection $-N(P) \cap \mathbb{R}_+^k$.




Figure 4. The regions S_{Λ_j} .




Figure 5. The regions F_{Λ_j} .

Definition 4.36. We define the dominant term of $p(x)$ corresponding to the component $x_1 = \dots = x_k = 0$ and the region F_{Λ_j} by

$$p_{dom_j}(x) := \sum_{\lambda \in \Lambda_j} c_\lambda x_{[1:k]}^\lambda q_\lambda(x_{[k+1:n]}).$$

The following observations about dominant terms are immediate.

Lemma 4.37. Let $\eta = (n_1, \dots, n_k) \in S_{\Lambda_j}$ and let $d_j := \sum_{\lambda_i \in \Lambda_j} \lambda_i n_i$. Let x^0 be fixed and let

$$x(t) := (x_1(t), \dots, x_n(t)), \quad x_j(t) := \begin{cases} t^{n_j} x_j^0 & j = 1, \dots, k, \\ x_j^0, & j = k+1, \dots, n. \end{cases}$$

Then $p_{dom_j}(x(t))$ has degree d_j in t and is the lowest degree term of $p(x(t))$ in t , that is

$$p(x(t)) = p_{dom_j}(x(t)) + o(t^{d_j}) \quad \text{as } t \rightarrow 0, \quad \deg_t p_{dom_j}(x(t)) = d_j.$$

Proof. Follows directly from the definition of a dominant term. □

Corollary 4.38. Under the assumptions of Lemma 4.37 suppose that $p_{dom_j}(x^0) \neq 0$. Then

$$\lim_{t \rightarrow 0} \frac{p_{dom_j}(x(t))}{p(x(t))} = 1.$$

Thus p_{dom_j} is the leading term along each curve traced by $x(t)$ as t tends to zero from above. An important question now is whether the dominant term p_{dom_j} indeed dominates the remaining terms of p in the region F_{Λ_j} in the sense that $p_{dom_j}(x)/p(x)$ is close to 1 sufficiently close to the component $x_1 = \cdots = x_k = 0$ of the variety $V(p)$. This requires, at a minimum, that the variety $V(p_{dom_j})$ does not have a component strictly larger than the set $x_1 = \cdots = x_k = 0$. Note that most dominant terms of a polynomial actually fail this requirement. Indeed, most dominant terms of p are monomials, which correspond to regions F_{Λ_j} indexed by singletons Λ_j , hence to the vertices of the Newton polytope of p . The dominant terms corresponding to larger sets Λ_j are more useful, since they pick up terms relevant not only in the region F_{Λ_j} but also in its neighborhood. In Example 4.35 above the dominant terms for $F_{\{(2,2),(8,0)\}}$ and $F_{\{(2,2),(0,8)\}}$, corresponding to the edges of the Newton polygon, are the useful ones. This points to the fact that we should be ultimately interested only in dominant terms corresponding to the facets, i.e., the highest-dimensional faces, of the Newton polytope of p . Note that the convex hull of Λ_j is a facet of the Newton polytope N if and only if the set S_{Λ_j} is a one-dimensional ray.

The next lemma will be instrumental for our results in Section 4.3.4. It shows that each dominant term p_{dom_j} such that the convex hull of Λ_j is a facet of the Newton polytope of p and whose variety $V(p_{dom_j})$ does not have a component strictly larger than the set $x_1 = \cdots = x_k = 0$ indeed dominates the remaining terms in p in a certain “slice” \tilde{F}_{Λ_j} around F_{Λ_j} .

Lemma 4.39. *Let p_{dom_j} be the dominant term of a homogeneous polynomial p corresponding to the component $x_1 = \cdots = x_k = 0$ of the variety $V(p)$ and to the set Λ_j whose convex hull is a facet of the Newton polytope N .*

Let \tilde{S}_{Λ_j} be any closed pointed cone in $(\mathbb{R}^k)_+$ with vertex at 0 that does not intersect other one-dimensional rays S_{Λ_l} , $l \neq j$, and contains $S_{\Lambda_j} \setminus \{0\}$ in its interior. Let \tilde{F}_{Λ_j} be the closure of the set

$$\{x_{[1:k]} \in [-1, 1]^k : (-\log|x_1|, \dots, -\log|x_k|) \in \tilde{S}_{\Lambda_j}\}. \quad (9)$$

Suppose the variety $V(p_{dom_j})$ of p_{dom_j} is allowable and intersects \tilde{F}_{Λ_j} only at 0. Let $\|\cdot\|$ be any norm. Then, for any $\delta = \delta(j) > 0$, there exists $\varepsilon = \varepsilon(j) > 0$ such that

$$\left| \frac{p_{dom_j}(x_{[1:k]}, x_{[k+1:n]})}{p(x_{[1:k]}, x_{[k+1:n]})} - 1 \right| < \delta \quad \text{whenever} \quad \frac{\|x_{[1:k]}\|}{\|x_{[k+1:n]}\|} \leq \varepsilon \quad \text{and} \quad x_{[1:k]} \in \tilde{F}_{\Lambda_j}. \quad (10)$$

Proof. We prove the lemma in the case \tilde{F}_{Λ_j} does not intersect nontrivially any of the coordinate planes (the proof extends to the other case via limiting arguments). Let $\|x_{[k+1:n]}\| = 1$ and let x_1 through x_k be ± 1 . If $\eta = (n_1, \dots, n_k) \in \tilde{S}_{\Lambda_j}$, then, directly from the definition of the set \tilde{F}_{Λ_j} , the curve $(t^{n_1}x_1, \dots, t^{n_k}x_k)$, $t \in (0, 1]$, lies in \tilde{F}_{Λ_j} (and every point in \tilde{F}_{Λ_j} lies on such a curve). Denote $(t^{n_1}x_1, \dots, t^{n_k}x_k, x_{[k+1:n]})$ by $x(t)$ and let t decrease from 1 to 0, keeping the x_m , $m = 1, \dots, n$, fixed. By the assumption of the Lemma, $p_{dom_j}(x(t))$ does not vanish for sufficiently small $t > 0$. Moreover, by Lemma 4.37, p_{dom_j} is the leading term of p in F_{Λ_j} . Since the cone \tilde{S}_{Λ_j} around S_{Λ_j} does not intersect any other one-dimensional rays S_{Λ_l} , $l \neq j$, all the monomials present in any term that dominates in $\tilde{F}_{\Lambda_j} \setminus F_{\Lambda_j}$ are already present in p_{dom_j} . Thus p_{dom_j} contains all terms that dominate in \tilde{F}_{Λ_j} . Therefore, there exists $\varepsilon(x) > 0$ such that $|p_{dom_j}(x(t))/p(x(t)) - 1| < \delta$ whenever $t < \varepsilon(x)$. The function $f : x \rightarrow \varepsilon(x)$ is lower semicontinuous. Since the set $S := \{x : x_m = \pm 1, m = 1, \dots, k, \|x_{[k+1:n]}\| = 1\}$ is compact, the minimum $\varepsilon := \min f(S)$ is necessarily positive and satisfies (10). \square

The above discussion of dominance was based on the transformation of a given irreducible component of the variety to the form $x_1 = \dots = x_k = 0$. We must reiterate that the identification of dominant terms becomes possible only after a suitable change of variables C is used to put a given irreducible component into the standard form $x_1 = \dots = x_k = 0$ and then the sets Λ_j are determined. Note however that the polynomial p_{dom_j} is given in terms of the original variables, i.e., as a sum of monomials in the original variables x_q and sums/differences $x_q \pm x_r$. We will therefore use the more precise notation $p_{dom_j, C}$ in the sequel.

Without loss of generality we can assume that any standard change of variables has the form

$$\begin{aligned} x &= (x_{[1:k_1]}, x_{[k_1+1:k_2]}, \dots, x_{[k_{l-1}+1:k_l]}) \mapsto \tilde{x} = (\tilde{x}_{[1:k_1]}, \tilde{x}_{[k_1+1:k_2]}, \dots, \tilde{x}_{[k_{l-1}+1:k_l]}), \text{ where} \\ \tilde{x}_{k_{m+1}} &:= x_{k_{m+1}}, \quad \tilde{x}_{k_{m+2}} := x_{k_{m+2}} - \sigma_{k_{m+2}} x_{k_{m+1}}, \quad \dots, \quad \tilde{x}_{k_{m+1}} := x_{k_{m+1}} - \sigma_{k_{m+1}} x_{k_{m+1}}, \\ k_0 &:= 0, \quad \sigma_r = \pm 1 \text{ for all pertinent } r \end{aligned} \quad (11)$$

Note also that we can think of the vectors $\eta \in S_{\Lambda_j}$ as being indexed by integers 1 through k_l , i.e., $\eta = (n_1, \dots, n_{k_l})$. Moreover, to define pruning in the next subsection we will assume that

$$n_{k_{m+1}} \leq n_r \text{ for all } r = k_m + 2, \dots, k_{m+1} \text{ and for all } m = 0, \dots, l-1. \quad (12)$$

Remark 4.40. *This condition is trivially satisfied if $n_{k_{m+1}} = 0$, as is the case for any group $x_{k_{m+1}} = \sigma_{k_{m+2}} x_{k_{m+2}} = \dots = \sigma_{k_{m+1}} x_{k_{m+1}}$ of original conditions that define the given irreducible component of $V(p)$, since $x_{k_{m+1}}$ does not have to be close to 0 in the neighborhood of that component of $V(p)$. If, however, the same group of equalities was created from the original conditions $x_{k_{m+1}} = x_{k_{m+2}} = \dots = x_{k_{m+1}} = 0$ due to the particular change of variables C , the condition (12) is no longer forced upon us. Yet (12) can be assumed without loss of generality. Indeed, if, say, $n_{k_{m+2}} < n_{k_{m+1}}$, then we can always switch to another standard change of variables by taking $x_{k_{m+2}}$ to be the representative of the group $x_{k_{m+1}}, \dots, x_{k_{m+1}}$ and taking the sums/differences with $x_{k_{m+2}}$ as the other new variables. Also note that (12) is satisfied either by all or by no vectors in S_{Λ_j} . In other words, (12) is a property of the entire set S_{Λ_j} . So, with a slight abuse of terminology we will say that a set S_{Λ_j} satisfies or fails (12).*

Finally note that the curves $(x(t))$ corresponding to the change of variables (11) are described as follows:

$$\begin{aligned} x(t) &:= (x_{[1:k_1]}(t), x_{[k_1+1:k_2]}(t), \dots, x_{[k_{l-1}+1:k_l]}(t), x_{[k_l+1:n]}), \text{ where} \\ x_{[k_{m+1}:k_{m+1}]}(t) &:= \\ &\quad (t^{n_{k_m}} x_{k_{m+1}}, t^{n_{k_m}+2} x_{k_{m+2}} + \sigma_{k_{m+2}} x_{k_{m+1}}, \dots, t^{n_{k_{m+1}}} x_{k_{m+1}} + \sigma_{k_{m+1}} x_{k_{m+1}}) \\ \text{where } k_0 &:= 1, \quad m = 0, \dots, l. \end{aligned} \quad (13)$$

This description will be instrumental in our discussion of pruning, which follows immediately.

4.3.3 Pruning

Now we discuss how to convert an accurate algorithm that evaluates a polynomial p into an accurate algorithm that evaluates a selected dominant term $p_{dom_j, C}$. This process, which we will refer to as *pruning*, will consist of deleting some vertices and edges and redirecting certain other edges in the DAG that represents the algorithm.

Definition 4.41 (Pruning). Given a non-branching algorithm represented by a DAG for computing $p_{\text{comp}}(x, \delta)$, a standard change of variables C of the form (11) and a subset $\Lambda_j \in \Lambda$ satisfying (12), we choose any $\eta \in S_{\Lambda_j}$, we input (formally) the expression (13), and then perform the following process.

We can perform one of two actions: redirection or deletion. By redirection (of sources) we mean replacing an edge from a source node corresponding to a variable x_j to a computational node i by an edge from the representative x_{rep} of x_j followed by exact negation if $\sigma_j = -1$. This corresponds to replacing x_j by the product $\sigma_j x_{\text{rep}}$. To define deletion, consider a node i with distinct input nodes j and k . Then deletion of node i from node j means deleting the out-edge to node i from node j , changing the origin of all out-edges from node i to input node k , and deleting node i .

Starting at the sources, we process each node as follows, provided that both its inputs have already been processed (this can be done because of acyclicity). Let the node being processed be node i , i.e., $(\text{op}(\cdot), \delta_i)$, and assume it has input nodes k and l . Both inputs being polynomials in t , we determine the lowest degree terms in t present in either of them and denote these degrees by $\deg(k)$ and $\deg(l)$.

if $\text{op}(\cdot) = \cdot$ **and one or both inputs are sources**, *then*

redirect each source input.

if $\text{op}(\cdot) = \pm$, *then*

if $\deg(k) \neq \deg(l)$, **say** $\deg(k) > \deg(l)$, *delete input node i from node k .*

else *If nodes k and l are sources and the operation $\text{op}(\cdot)$ leads to cancellation of their lowest degree terms in t , examine their second-lowest degree terms. If those degrees coincide or if one second-lowest term is missing, we change nothing. If one is bigger than the other, we do not change the source containing the lower degree term in t , but redirect the other source.*

If only one of nodes k and l is a source or if both inputs are sources but there is no cancellation of lowest degree terms, redirect each source.

We then delete inductively all nodes which no longer are on any path to the output.

We call this process pruning, and denote the output of the pruned algorithm by $p_{\text{dom}_j, C, \text{comp}}(x, \delta)$.

Remark 4.42. *Note that the outcome of pruning does not depend on the choice of $\eta \in S_{\Lambda_j}$. Since each region S_{Λ_j} is determined by linear homogeneous equalities and inequalities with integer coefficients, the vector η can always be chosen to have all integer entries.*

Example 4.43. *Figure 6 shows an example of pruning an algorithm that evaluates the polynomial*

$$x_1^2 x_2^2 + (x_2 - x_3)^4 + (x_3 - x_4)^2 x_5^2$$

using the substitution

$$(tx_1, x_2, tx_3 + x_2, tx_4 + x_2, x_5)$$

near the component

$$x_1 = 0, \quad x_2 = x_3 = x_4.$$




Figure 6. Pruning an algorithm for $p(x) = x_1^2 x_2^2 + (x_2 - x_3)^4 + (x_3 - x_4)^2 x_5^2$.

The result of pruning is an algorithm that evaluates the dominant term

$$x_1^2 x_2^2 + (x_3 - x_4)^2 x_5^2.$$

One of two branches leading to the node A is pruned due to the fact that it computes a quantity of order $O(t^4)$ whereas the other branch produces a quantity of order $O(t^2)$.

The output of the original algorithm is given by

$$\begin{aligned} & ((x_1^2(1 + \delta_1)x_2^2(1 + \delta_2)(1 + \delta_3) + (x_2 - x_3)^4(1 + \delta_4)^4(1 + \delta_5)^2(1 + \delta_6)) (1 + \delta_7) \\ & + (x_3 - x_4)^2(1 + \delta_8)^2(1 + \delta_9)x_5^2(1 + \delta_{10})(1 + \delta_{11})) (1 + \delta_{12}). \end{aligned}$$

The output of the pruned algorithm is

$$(x_1^2(1 + \delta_1)x_2^2(1 + \delta_2)(1 + \delta_3))(1 + \delta_7) + (x_3 - x_4)^2(1 + \delta_8)^2(1 + \delta_9)x_5^2(1 + \delta_{10})(1 + \delta_{11})) (1 + \delta_{12}).$$

Let us prove that this process will indeed produce an algorithm that accurately evaluates the corresponding dominant term.

Theorem 4.44. *Suppose a non-branching algorithm evaluates a polynomial p accurately on \mathbb{R}^n by computing $p_{comp}(x, \delta)$. Suppose C is a standard change of variables (11) associated with an irreducible component of $V(p)$. Let $p_{dom_j, C}$ be one of the corresponding dominant terms of p and let S_{Λ_j} satisfy (12). Then the pruned algorithm defined in Definition 4.41 with output $p_{dom_j, C, comp}(x, \delta)$ evaluates $p_{dom_j, C}$ accurately on \mathbb{R}^n . In other words, being able to compute all such $p_{dom_j, C}$ for all components of the variety $V(p)$ and all standard changes of variables C accurately is a condition necessary to compute p accurately.*

Proof. Directly from the definition of pruning it can be seen that the output of each computational node is a homogeneous polynomial in t . This can be checked inductively starting from computational nodes operating on two sources, using the pruning rules. Moreover, the pruning rules are equivalent to taking the lowest degree terms in t (as well as setting some δ s to zero). This can be checked inductively as well, once we rule out the situation when a \pm node in the original algorithm leads to exact cancellation of lowest degree terms of the inputs, and at least one of the inputs is not a source. Indeed, in that case one of the inputs contains a factor $(1 + \delta)$ and that δ by acyclicity is not present in the other input. Therefore no exact cancellation of lowest degree terms can occur.

Thus, the final output $p_{dom_j, C, comp}(x, \delta)$ of the pruned algorithm takes the lowest degree terms in t of the final output of the original algorithm, so $p_{dom_j, C, comp}(x, \delta)$ is homogeneous in t (of degree $d_j = \sum_{\lambda_i \in \Lambda_j} \lambda_i n_i$). We write

$$\frac{p_{comp}(x, \delta) - p(x)}{p(x)} = \sum_{\alpha} \frac{p_{\alpha}(x)}{p(x)} \delta^{\alpha}, \quad \frac{p_{dom_j, C, comp}(x) - p_{dom_j, C}(x)}{p_{dom_j, C}(x)} = \sum_{\alpha} \frac{p_{\alpha, dom_j, C}(x)}{p_{dom_j, C}(x)} \delta^{\alpha}.$$

Note that, in eliminating nodes and redirecting the out-edges in the pruning process, we do the equivalent of setting some of the δ s to 0, and hence the set of monomials δ^{α} present in the second sum is a subset of the set of monomials δ^{α} present in the first sum.

Indeed, first of all, we can focus on the effect of deleting some nodes, since redirection does not affect δ s at all, because only sources can be redirected. So, if δ^{α} appears in the second sum, there is a path which yields the corresponding multi-index, on which some term in $p_{\alpha, dom_j, C}(x)$ is computed. But since this path survived the pruning, there is a corresponding path in the original DAG which perhaps has a few more nodes that have been deleted in the pruning process and a few source nodes that were redirected. In the computation, the effect of deleting a node was to set that δ equal to 0 (and make some terms of higher degree disappear). So the surviving term was also present in the computation of $p_{comp}(x, \delta)$, with the same multi-index: just choose the 1 in the $(1 + \delta)$ each time when you hit a node that will be deleted (i.e., whose δ will be set to 0).

Now note that $p_{\alpha, dom_j, C}(x)$ is the leading term of $p_{\alpha}(x)$, i.e., the term of smallest degree d_j in t . This happens since each term of degree d_j in $p_{\alpha}(x)$ must survive on the same path in the DAG, with the same choices of 1 in $(1 + \delta)$ each time we hit a deleted node.

We can now prove that $p_{dom_j, C, comp}(x)$ is accurate. To do that, it is enough to show that each $p_{\alpha, dom_j, C}(x)/p_{dom_j, C}(x)$ is bounded, provided that there is some constant M such that $|p_{\alpha}(x)/p(x)| \leq M$ for all x .

Choose a point $x = (x_1, \dots, x_n)$ not on the variety of $p_{dom_j, C}$ and consider the curve traced by the associated point $x(t)$ from (13) as t tends to 0. Since both $p_{\alpha, dom_j, C}(x(t))$ and $p_{dom_j, C}(x(t))$

are homogeneous of degree d_j in t , we have

$$\left| \frac{p_{\alpha, \text{dom}_j, C}(x)}{p_{\text{dom}_j, C}(x)} \right| = \left| \frac{p_{\alpha, \text{dom}_j, C}(x(t))}{p_{\text{dom}_j, C}(x(t))} \right| = \lim_{t \rightarrow 0} \left| \frac{p_{\alpha, \text{dom}_j, C}(x(t))}{p_{\text{dom}_j, C}(x(t))} \right|.$$

Since both $p_{\alpha, \text{dom}_j, C}(x(t))$ and $p_{\text{dom}_j, C}(x(t))$ are the dominant terms in $p_{\alpha}(x)$, respectively $p(x)$, along the curve $\{x(t) : t \rightarrow 0\}$, we conclude that

$$\left| \frac{p_{\alpha, \text{dom}_j, C}(x)}{p_{\text{dom}_j, C}(x)} \right| = \lim_{t \rightarrow 0} \left| \frac{p_{\alpha, \text{dom}_j, C}(x(t))}{p_{\text{dom}_j, C}(x(t))} \right| = \lim_{t \rightarrow 0} \left| \frac{p_{\alpha}(x(t))}{p(x(t))} \right| \leq M.$$

Invoking the density of the Zariski open set $\{x : p_{\text{dom}_j, C}(x) \neq 0\}$, we are done. \square

4.3.4 Sufficiency of evaluating dominant terms

Our next goal is to prove a converse of a sort to Theorem 4.44. Strictly speaking, the results that follow do not provide a true converse, since branching is needed to construct an algorithm that evaluates a polynomial p accurately from algorithms that evaluate its dominant terms accurately.

For the rest of this section, we use two blanket assumptions, viz., that our polynomial p is homogeneous and irreducible. The latter assumption effectively reduces the problem to that of accurate evaluation of a nonnegative polynomial, due to the following lemma.

Lemma 4.45. *If a polynomial p is irreducible and has an allowable variety $V(p)$, then it is either a constant multiple of a linear form that defines an allowable hyperplane or it does not change its sign in \mathbb{R}^n .*

Proof. Suppose that p changes its sign. Then the sets $\{x : p(x) > 0\}$ and $\{x : p(x) < 0\}$ are both open, hence the part of the variety $V(p)$ whose neighborhood contains points from both sets must have dimension $n - 1$. The only allowable sets of dimension $n - 1$ are allowable hyperplanes. Therefore, the linear polynomial that defines an allowable hyperplane must divide p . As p is assumed to be irreducible, p must be a constant multiple of the linear polynomial.

Thus, unless p is a constant multiple of a linear factor of the type x_i or $x_i \pm x_j$, it must satisfy $p(x) \geq 0$ for all $x \in \mathbb{R}^n$ or $p(x) \leq 0$ for all $x \in \mathbb{R}^n$. \square

From now on we therefore restrict ourselves to the nontrivial case when a (homogeneous and irreducible) polynomial p is nonnegative everywhere in \mathbb{R}^n .

Theorem 4.46. *Let p be a homogeneous nonnegative polynomial whose variety $V(p)$ is allowable. Suppose that all dominant terms $p_{\text{dom}_j, C}$ for all components of the variety $V(p)$, all standard changes of variables C and all subsets Λ_j satisfying (12) are accurately evaluable. Then there exists a branching algorithm that evaluates p accurately over \mathbb{R}^n .*

Proof. We first show how to evaluate p accurately in a neighborhood of each irreducible component of its variety $V(p)$. We next evaluate p accurately off these neighborhoods of $V(p)$. The final algorithm will involve branching depending on which region the input belongs to, and the subsequent execution of the corresponding subroutine. We fix the relative accuracy η that we want to achieve.

Consider a particular irreducible component V_0 of the variety $V(p)$. Using any standard change of variables C , say, a basic change of variables associated with V_0 , we map V_0 to a set of the form $\tilde{x}_1 = \cdots = \tilde{x}_k = 0$. Our goal is to create an ε -neighborhood of V_0 where we can evaluate

p accurately. It will be built up from semialgebraic ε -neighborhoods. We begin with any set S containing a neighborhood V_0 , say, we let S coincide with $[-1, 1]^k \times \mathbb{R}^{n-k}$ after the change of variables C . We partition the cube $[-1, 1]^k$ into sets \tilde{F}_{Λ_j} of type (9) as follows: We consider, as in Section 4.3.2 above, the Newton polytope P of p , form the intersection of the negative of its normal fan $-N(P)$ with the nonnegative quadrant \mathbb{R}_+^k (in the new coordinate system), and determine the sets S_{Λ_j} . If condition (12) fails for some of the sets S_{Λ_j} , we transform them using a suitable standard change of variables as described in Remark 4.40 so as to meet condition (12). For the rest of the argument, we can assume that all sets S_{Λ_j} satisfy (12).

To form conic neighborhoods of one-dimensional rays of $-N(P) \cap \mathbb{R}_+^k$ (which are normal to facets of the Newton polytope), we intersect $-N(P) \cap \mathbb{R}_+^k$ with, say, the hyperplane $\tilde{x}_1 + \cdots + \tilde{x}_k = 1$. Perform the Voronoi tessellation (see, e.g., [64]) of the simplex $\tilde{x}_1 + \cdots + \tilde{x}_k = 1, \tilde{x}_j \geq 0, j = 1, \dots, k$ relative to the intersection points of $-N(P) \cap \mathbb{R}_+^k$ with the hyperplane $\tilde{x}_1 + \cdots + \tilde{x}_k = 1$. Connecting each Voronoi cell of the tessellation to the origin $\tilde{x}_1 = \cdots = \tilde{x}_k = 0$ by straight rays, we obtain cones \tilde{S}_{Λ_j} and the corresponding sets \tilde{F}_{Λ_j} of type (9). Note that the Voronoi cells and therefore the cones \tilde{S}_{Λ_j} are determined by rational inequalities since the tessellation centers have rational coordinates. Hence the sets \tilde{F}_{Λ_j} are semialgebraic, and moreover are determined by polynomial inequalities with integer coefficients.

Consider a particular ‘‘slice’’ \tilde{F}_{Λ_j} and the dominant term $p_{dom_j, C}$. By Theorem 4.44, the dominant term $p_{dom_j, C}$ must be accurately evaluable everywhere. Hence, in particular, its variety $V(p_{dom_j, C})$ must be allowable. Since the polynomial p vanishes on V_0 , the dominant term $p_{dom_j, C}$ must vanish on V_0 as well. So, there are two possibilities: $V(p_{dom_j, C}) \cap \tilde{F}_{\Lambda_j}$ either coincides with V_0 or is strictly larger.

In the first case we apply Lemma 4.39 to show that we can evaluate p accurately in \tilde{F}_{Λ_j} sufficiently close to V_0 , as follows: Since the polynomial $p_{dom_j, C}$ is accurately evaluable everywhere, for any number $\eta_j > 0$ there exists an algorithm with output $p_{dom_j, C, comp}$ such that $|p_{dom_j, C, comp} - p_{dom_j, C}|/|p_{dom_j, C}| < \eta_j$ everywhere. Next, by Lemma 4.39, for any $\delta_j > 0$ there exists $\varepsilon_j > 0$ such that (10) holds. Choose η_j and δ_j so that

$$\eta_j(1 + \delta_j) + \delta_j < \eta.$$

Then we have

$$\begin{aligned} \left| \frac{p_{dom_j, C, comp} - p}{p} \right| &\leq \frac{|p_{dom_j, C, comp} - p_{dom_j, C}| + |p_{dom_j, C} - p|}{|p|} \\ &= \frac{|p_{dom_j, C, comp} - p_{dom_j, C}|}{|p_{dom_j, C}|} \cdot \frac{|p_{dom_j, C}|}{|p|} + \frac{|p_{dom_j, C} - p|}{|p|} \leq \eta_j(1 + \delta_j) + \delta_j < \eta \end{aligned}$$

in the ε_j -neighborhood of V_0 within the set \tilde{F}_{Λ_j} . Therefore, p can be evaluated by computing $p_{dom_j, C, comp}$ to accuracy η in the ε_j -neighborhood of V_0 within \tilde{F}_{Λ_j} .

In the second case $V(p_{dom_j, C})$ has an irreducible component, say, W_0 , that is strictly larger than V_0 and intersects \tilde{F}_{Λ_j} nontrivially. Since $V(p_{dom_j, C})$ is allowable, it follows from Definition 4.29 that there exists a standard change of variables C_1 associated with V_0 that maps W_0 to a set $\tilde{x}_1 = \cdots = \tilde{x}_l = 0, l < k$. Use that change of variables and consider the new Newton polytope and dominant terms. Since the polynomial p is positive in $S \setminus V_0$, there are terms in p that do not contain variables \tilde{x}_1 through \tilde{x}_l . Therefore each new p_{dom_r, C_1} picks up some of those nonvanishing terms. Hence the (allowable) varieties $V(p_{dom_r, C_1})$ have irreducible components strictly smaller than W_0

(but still containing V_0). So, we can subdivide the set \tilde{F}_{Λ_j} further using the sets \tilde{F}_{Λ_r} coming from the change of variables C_1 and the resulting dominant terms will vanish on a set strictly smaller than W_0 . In this fashion, we refine our subdivision repeatedly until we obtain a subdivision of the original set S into semialgebraic pieces (S_j) such that the associated dominant terms p_{dom_j} vanish in S_j only on V_0 . Applying Lemma 4.39 in each such situation, we conclude that p can be evaluated accurately sufficiently close to V_0 within each piece S_j .

For each V_0 , we therefore can find a collection (S_j) of semialgebraic sets, all determined by polynomial inequalities with integer coefficients, and the corresponding numbers ε_j , so that the polynomial p can be evaluated with accuracy η in each ε_j -neighborhood of V_0 within the piece S_j . Note that we can assume that each ε_j is a reciprocal of an integer, so that testing whether a particular point x is within ε_j of V_0 within S_j can be done by branching based on polynomial inequalities with integer coefficients.

The final algorithm will be organized as follows. Given an input x , determine by branching whether x is in S_j and within the corresponding ε_j of a component V_0 . If that is the case, evaluate $p(x)$ using the algorithm that is accurate in S_j in that neighborhood of V_0 . For x not in any of the neighborhoods, evaluate p by Horner's rule. Since the polynomial p is strictly positive off the neighborhoods of the components of its variety, the reasoning of Section 3 applies, showing that the Horner's rule algorithm is accurate. If x is on the boundary of a set S_j , any applicable algorithm will do, since the inequalities we use are not strict. Thus the resulting algorithm for evaluating p will have accuracy η as required. \square

4.3.5 Obstacles to full induction

The reasoning above suggests that there could be an inductive decision procedure that would allow us to determine whether or not a given polynomial is accurately evaluable by reducing the problem for the original polynomial p to the same problem for its dominant terms, then their dominant terms, and so forth, going all the way to monomials or other polynomials that are easy to analyze. However, this idea would only work if the dominant terms were somehow "simpler" than the original polynomial itself. This approach would require an induction variable that would decrease at each step.

Two possible choices are the number of variables or the degree of the polynomial under consideration. Sometimes, however, neither of the two goes down, and moreover, the dominant term may even coincide with the polynomial itself. For example, if

$$p(x) = A(x_{[3:n]})x_1^2 + B(x_{[3:n]})x_1x_2 + C(x_{[3:n]})x_2^2$$

where A, B, C are nonnegative polynomials in x_3 through x_n , then the only useful dominant term of p in the neighborhood of the set $x_1 = x_2 = 0$ is the polynomial p itself. Thus no progress whatsoever is made in this situation.

Another possibility is induction on domains but we do not yet envision how to make this idea precise, since we do not know exactly when a given polynomial is accurately evaluable on a given domain.

Further work to establish a full decision procedure is therefore highly desirable.

5 “Black-box” arithmetic

In this section we prove a necessary condition (for both the real and the complex cases) for a more general type of arithmetic, which allows for “black-box” polynomial operations. We describe the type of operations below.

Definition 5.1. *We call a black-box operation any type of operation that takes a number of inputs (real or complex) x_1, \dots, x_k and produces an output q such that q is a polynomial in x_1, \dots, x_k .*

Example 5.2. $q(x_1, x_2, x_3) = x_1 + x_2x_3$.

Remark 5.3. *Note that $+$, $-$, and \cdot are all black-box operations.*

Consider a fixed set of multivariate polynomials $\{q_j : j \in J\}$ with real or complex inputs (this set may be infinite). In our model under consideration, the arithmetic operations allowed are given by the black-box operations q_1, \dots, q_k , and negation (which will be dealt with by way of dotted edges, as in Section 4). With the exception of negation, which is exact, all the others yield a $\text{rnd}(op(a_1, \dots, a_l)) = op(a_1, \dots, a_l)(1 + \delta)$, with $|\delta| < \epsilon$ (ϵ here is the machine precision). All arithmetic operations have unit cost. We consider the same arithmetical models as in Section 2.7, with this larger class of operations.

5.1 Necessity: real and complex

In order to see how the statement of the necessity Theorem 4.8 changes, we need to introduce a different notion of allowability.

Recall that we denote by \mathcal{S} the space of variables (which may be either \mathbb{R}^n or \mathbb{C}^n). From now on we will denote the set $\{1, \dots, k\}$ by \mathcal{K} .

Definition 5.4. *Let $p(x_1, \dots, x_k)$ be a multivariate polynomial over \mathcal{S} with variety $V(p)$. Let $\mathcal{K}_Z \subseteq \mathcal{K}$, and let $\mathcal{K}_D, \mathcal{K}_S \subseteq \mathcal{K} \times \mathcal{K}$. Modify p as follows: impose conditions of the type Z_i for each $i \in \mathcal{K}_Z$, and of type D_{ij} , respectively S_{ij} , on all couples of variables in \mathcal{K}_D , respectively \mathcal{K}_S . Rewrite p subject to those conditions (e.g. set $X_i = 0$ for all $i \in \mathcal{K}_Z$), and denote it by \tilde{p} , and denote by \mathcal{K}_R the set of remaining independent variables.*

Choose a set $T \subseteq \mathcal{K}_R$, and let

$$V_{T, \mathcal{K}_Z, \mathcal{K}_D, \mathcal{K}_S}(p) = \cap_{\alpha} V(q_{\alpha}) ,$$

where the polynomials q_{α} are the coefficients of the expansion of \tilde{p} in the variables x_T :

$$\tilde{p}(x_1, \dots, x_k) = \sum_{\alpha} q_{\alpha} x_T^{\alpha} ,$$

with q_{α} being polynomials in $x_{\mathcal{K}_R \setminus T}$ only.

Finally, let \mathcal{K}_N be a subset of $\mathcal{K}_R \setminus T$. We negate each variable in \mathcal{K}_N , and let $V_{T, \mathcal{K}_Z, \mathcal{K}_D, \mathcal{K}_S, \mathcal{K}_N}(p)$ be the polynomial obtained from $V_{T, \mathcal{K}_Z, \mathcal{K}_D, \mathcal{K}_S}(p)$, with each variable in \mathcal{K}_N negated.

Remark 5.5. $V_{\emptyset, \emptyset, \emptyset, \emptyset}(p) = V(p)$. *We also note that, if we have a black-box computing p , then the set of all polynomials \tilde{p} that can be obtained from p by permuting, repeating, and negating the variables (as in the definition above) is exactly the set of all polynomials that can be evaluated with a single rounding error, using that black-box.*

Definition 5.6. For simplicity, we denote a set $(T, \mathcal{K}_Z, \mathcal{K}_D, \mathcal{K}_S, \mathcal{K}_N)$ by \mathcal{I} , and a set $(T, \mathcal{K}_S, \mathcal{K}_D, \mathcal{K}_Z)$ by \mathcal{I}_+ .

Example 5.7. Let $p(x, y, z) = x + y \cdot z$ (the fused multiply-add). We record below all possibilities for $\mathcal{I} = (T, \mathcal{K}_Z, \mathcal{K}_D, \mathcal{K}_S, \mathcal{K}_N)$, together with the obtained subvariety $V_{\mathcal{I}}(p)$.

Without loss of generality, assume that we have eliminated all redundant or complicated conditions, like $(x, y) \in \mathcal{K}_D$ and $(x, y) \in \mathcal{K}_S$ (which immediately leads to $x = y = 0$, that is, $x, y, \in \mathcal{K}_Z$). We assume thus that all variables not present in \mathcal{K}_Z cannot be deduced to be 0 from conditions imposed by \mathcal{K}_D or/and \mathcal{K}_S .

We obtain that all possibilities for $V_{\mathcal{I}}(p)$ are, up to a permutation of the variables,

- ◇ $\{x = 0\}, \{x = 1\}, \{x = -1\},$
- ◇ $\{x = 0\} \cup \{x = 1\}, \{x = 0\} \cup \{x = -1\},$
- ◇ $\{x = 0\} \cup \{y = 0\}, \{x = 0\} \cup \{y = 1\}, \{x = 0\} \cup \{y = -1\},$
- ◇ $\{x = -y^2\}, \{x = y^2\}, \{x - y \cdot z = 0\},$ and $\{x + y \cdot z = 0\}.$

Definition 5.8. We define $q_{-2}(x_1, x_2) = x_1 x_2$, $q_{-1}(x_1, x_2) = x_1 + x_2$, and $q_0(x_1, x_2) = x_1 - x_2$.

Remark 5.9. The sets

$$1. \quad Z_i = \{x : x_i = 0\} , \tag{14}$$

$$2. \quad S_{ij} = \{x : x_i + x_j = 0\} , \tag{15}$$

$$3. \quad D_{ij} = \{x : x_i - x_j = 0\} , \tag{16}$$

and unions thereof, describe all non-trivial (neither \emptyset nor \mathcal{S}) sets of type $V_{\mathcal{I}}$, for q_{-2} , q_{-1} , and q_0 .

We will assume from now on that the black-box operations q_{-2}, q_{-1}, q_0 defined in 5.8, and some arbitrary extra operations ($q_j : j > 0$) are given and fixed.

Definition 5.10. We call any set $V_{\mathcal{I}}$ with $\mathcal{I} = (T, \mathcal{K}_Z, \mathcal{K}_D, \mathcal{K}_S, \mathcal{K}_N)$ as defined above basic q -allowable.

We call any set R irreducible q -allowable if it is an irreducible component of a (finite) intersection of basic q -allowable sets, i.e., when R is irreducible and

$$R \subseteq \cap_l Q_l ,$$

where each Q_l is a basic q -allowable set.

We call any set Q q -allowable if it is a (finite) union of irreducible q -allowable sets, i.e.

$$Q = \cup_j R_j ,$$

where each R_j is an irreducible q -allowable set.

Any set R which is not q -allowable we call q -unallowable.

Remark 5.11. Note that the above definition of q -allowability is closed under taking union, intersection, and irreducible components. This parallels the definition of allowability for the classical arithmetic case – in the classical case, every allowable set was already irreducible (being an intersection of hyperplanes).

Once again, we need to build the setup to state and prove our new necessity condition. To do this, we will modify the statements of the definitions and the statements and proofs of the lemmas from Section 4.1. Since most proofs just follow in the footsteps of those from Section 4.1, instead of repeating them, we will only point out the places where they differ and show how we modified them to work in the new context.

Definition 5.12. Given a polynomial p with q -unallowable variety $V(p)$, consider all sets W that are q -allowable (as in Definition 5.10), and subtract from $V(p)$ those W for which $W \subset V(p)$. We call the remaining subset of the variety points in general position and denote it by $\mathcal{G}(p)$.

Remark 5.13. Since $V(p)$ is q -unallowable, $\mathcal{G}(p)$ is non-empty.

Definition 5.14. Given $x \in \mathcal{S}$, define the set $q\text{-Allow}(x)$ as the intersection of all basic q -allowable sets going through x :

$$q\text{-Allow}(x) := \bigcap_{j=-2}^k \left(\bigcap_{x \in V_{\mathcal{I}}(q_j)} V_{\mathcal{I}}(q_j) \right),$$

for all possible choices of $T, \mathcal{K}_Z, \mathcal{K}_D, \mathcal{K}_S, \mathcal{K}_N$, and with the understanding that

$$q\text{-Allow}(x) := \mathcal{S} \quad \text{whenever} \quad x \notin V_{\mathcal{I}}(q_j), \quad \text{for all } j \in \{-2, \dots, k\}, \text{ etc.}$$

Remark 5.15. When $x \in \mathcal{G}(p)$, $q\text{-Allow}(x) \not\subset \mathcal{G}(p)$.

We can now state our necessity condition.

Theorem 5.16. Given the black-box operations $q_{-2}, q_{-1}, q_0, q_1, \dots, q_k$ and the model of arithmetic described above, let p be a polynomial defined over a domain $\mathcal{D} \subset \mathcal{S}$. Let $\mathcal{G}(p)$ be the set of points in general position on the variety $V(p)$. If there exists $x \in \mathcal{D} \cap \mathcal{G}(p)$ such that $q\text{-Allow}(x) \cap \text{Int}(\mathcal{D}) \neq \emptyset$, then p is not accurately evaluable on \mathcal{D} .

We proceed to the construction of the elements of the proof of Theorem 5.16. The algorithm will once again be represented by a DAG with input nodes, branching nodes, and output nodes. As in Section 4, for simplicity in dealing with negation (since negation is exact), we will work with *solid* edges, which convey a value unchanged, and *dotted* edges, which indicate that negation of the conveyed quantity has occurred.

From now on, unless specified, we will consider only non-branching algorithms.

We will continue to use the definition of a Zariski set (Definition 4.9) on a hypercube in δ -space, and we work with the same definition of a non-trivial computational node (recalled below).

Definition 5.17. For a given $x \in \mathcal{S}$, we say that a computational node N is of q -non-trivial type if its output is a non-constant polynomial of δ .

Recall the notation H_ϵ from Definition 4.12.

The equivalent of Proposition 4.13 becomes the following.

Proposition 5.18. *Given any algorithm, any $\epsilon > 0$, and a point x in $\mathcal{G}(p)$, there exists a Zariski open set $\Delta \subset H_\epsilon$ such that no q -non-trivial computational node has a zero output on the input x for all $\delta \in \Delta$.*

Proof. The proof of Proposition 5.18 follows the same path as that of Proposition 4.13. To each q -non-trivial node corresponds a Zariski open set in δ -space; there is a finite number of them, and their intersection provides us with the Zariski open set we are looking for. \square

We will now state and sketch the proof of the equivalent of Lemma 4.14.

Lemma 5.19. *For a given algorithm, $x \in \mathcal{G}(p)$, and $\epsilon > 0$, exactly one of the following holds:*

1. *there exists a Zariski open set $\Delta \in H_\epsilon$ such that the value $p_{comp}(x, \delta)$ computed by the algorithm is not zero when the algorithm is run with source input x and errors $\delta \in \Delta$;*
2. *$p_{comp}(y, \delta) = 0$ for all $y \in q\text{-Allow}(x)$ and all $\delta \in H_\epsilon$.*

Proof of Lemma 5.19. Give $x \in \mathcal{G}(p)$, choose δ from the Zariski open set Δ whose existence is given by Proposition 5.18. Either the output node is of q -non-trivial type (in which case $p_{comp}(x, \cdot) \neq 0$ and we are done) or there is some $\delta_0 \in \Delta$ for which $p_{comp}(x, \delta_0) = 0$. In this latter case, we trace back all the zeros again, as in the proof of Lemma 4.14, and get a set of paths of nodes that produced all 0.

Let us start from the last nodes on these paths and work our way up, level after level. The last node on such a path is either a source, or a node with all inputs from sources, or a node with at least one input which is not a source and not 0 (and hence a polynomial in δ). In the former two cases, we have traced back the zeros to basic q -allowable conditions. We will show that this is also true for the latter case.

Lemma 5.20. *If the last zero occurs at a node which computes a black-box operation q_j and which has some source inputs and some (nontrivial) polynomial inputs, then the sources lie in some $V_{\mathcal{I}}(q_j)$ (and this constraint causes the zero output of this node).*

Proof of Lemma 5.20. Label the node \tilde{N} (assume it corresponds to the black-box q_j) with output 0; then some inputs are sources, and some are polynomials of δ (since this was the last node on the path, it has no zero inputs). By the choice of δ , it follows that the output has to be the 0 polynomial in δ .

Some of the non-source inputs to \tilde{N} might come from the same nodes; assume you have a total of l distinct nodes which input to \tilde{N} (non-constant) polynomials of δ , and let $I_1(\delta), I_2(\delta), \dots, I_l(\delta)$ denote these inputs. We will need the following lemma.

Lemma 5.21. *Since the DAG is acyclic, $I_1(\delta), I_2(\delta), \dots, I_l(\delta)$ are algebraically independent polynomials in δ .*

Proof of Lemma 5.21. Suppose there is some polynomial dependence among $I_1(\delta), \dots, I_l(\delta)$. Let N_1, \dots, N_l be the nodes which have computed $I_1(\delta), \dots, I_l(\delta)$, and let $\delta_1, \delta_2, \dots, \delta_l$ be the specific δ s at these nodes. Let D_1, \dots, D_l be the set of δ s present (non-trivially) in the output of each node, e.g., $\delta_i \in D_i$. At least one δ_i is not present in $\cup_{i \neq j} D_j$; otherwise we get a cycle. But then $I_i(\delta)$ is algebraically independent from the other inputs, i.e., there is some dependence among the inputs $I_j(\delta)$ with $j \neq i$. We induct on the number of remaining inputs, and exclude one input at a time, until we're left with a contradiction. This proves Lemma 5.21. \square

Replace each $I_i(\delta)$ (or $-I_i(\delta)$) in Y by the same *dummy* variable z_i (respectively $-z_i$), for each $i \in \{1, \dots, l\}$. The variables z_i are algebraically independent by Lemma 5.21.

Denote by z the new vector of inputs to \tilde{N} (both values *and* variables). The value $q_j(z) = 0$, regardless of the z_i s (since it was 0 regardless of the δ s and the z_i are algebraically independent variables). It follows that the constraints which place z on the variety of q_j are twofold: they come from constraints of the type D_{ij} and S_{ij} which describe the places where we inputted the values z_i , and they come from imposing conditions on the *other* inputs, which are *sources*. Thus the constraints on z are of the form $V_{T, \emptyset, \mathcal{K}_D, \mathcal{K}_S, \mathcal{K}_N}(q_j) =: V_{\mathcal{I}}$. This concludes the proof of Lemma 5.20. \square

Now that we have shown that the last marked vertices all provide basic q -allowable conditions, we proceed by induction: we look at a “next” marked vertex (here “next” means that all its marked ancestors have been examined already). It has some zero inputs, some source inputs, some of the inputs satisfy constraints of type D_{ij} and S_{ij} , and some of the inputs are polynomial. From here on we proceed like in Lemma 5.20, and obtain a set of new constraints to be imposed on the sources, of the type $V_{\mathcal{I}}(q_j)$, which we will intersect with the rest of the constraints obtained so far.

At the end of the examinations, we have found a set of basic q -allowable constraints which the sources must satisfy, i.e., a list of basic q -allowable sets with the property that the sources lie in their intersection; the fact that the sources satisfy these constraints is responsible for the zero output at the end of the computation.

It is not hard to see that in this case, once again, it follows that for all y in $q\text{-Allow}(x)$ and any $\delta \in \Delta$, the output is 0 (just like in Lemma 4.14). Thus we have proved Lemma 5.19. \square

Corollary 5.22. *Lemma 5.19 implies that, for any algorithm, for any $\epsilon > 0$, and any $x \in \mathcal{G}(p)$, exactly one of the following holds: the relative error of computation, $|p_{\text{comp}} - p|/|p|$, is either infinity at x for all δ in a Zariski open set or 1 at all points $y \in (q\text{-Allow}(x) \setminus V(p))$ and all $\delta \in H_\epsilon$.*

We will now consider algorithms with or without branches.

Theorem 5.23. *Given a (branching or non-branching) algorithm with output function $p_{\text{comp}}(\cdot)$, $x \in \mathcal{G}(p)$, and $\epsilon > 0$, then one of the following is true:*

1. *there exists a set $\Delta_1 \in H_\epsilon$ of positive measure such that $p_{\text{comp}}(x, \delta)$ is nonzero whenever the algorithm is run with errors $\delta \in \Delta_1$, or*
2. *there exists a set $\Delta_2 \in H_\epsilon$ of positive measure such that for every $\delta \in \Delta_2$, there exists a neighborhood $N_\delta(x)$ of x such that for every $y \in N_\delta(x) \cap (q\text{-Allow}(x) \setminus V(p))$, $p_{\text{comp}}(y, \delta) = 0$ when the algorithm is run with errors δ .*

Remark 5.24. *Just as before, this implies that, on a set of positive measure in H_ϵ , the relative accuracy of any given algorithm is either ∞ or 1.*

Proof. The proof is essentially the same as in Theorem 4.17; the only thing that needs to be examined is the existence in $q\text{-Allow}(x) \setminus V(p)$ of an infinite sequence $\{y_n\}$ with $y_n \rightarrow x$.

We will make use of the following basic result in the theory of algebraic varieties, which can for example be found as Theorem 1 in [54, Section 6.1].

Result 5.24.1. *If X and Y are polynomial varieties such that $X \subseteq Y$, then $\dim(X) \leq \dim(Y)$. If Y is irreducible and $X \subseteq Y$ is a (closed) subvariety with $\dim(X) = \dim(Y)$, then $X = Y$.*

We write $q\text{-Allow}(x)$ as a union of irreducible q -allowable components. By the way we defined $\mathcal{G}(p)$, it follows that none of these components is included in $V(p)$; by Result 5.24.1, it follows that the intersection of any irreducible q -allowable component P of $q\text{-Allow}(x)$ with $V(p)$ has a smaller dimension than P .

Choose the (unique) irreducible component P that contains x ; this component must have dimension at least 1 (since if it contained only x , the set $\{x\}$ would be q -allowable, and hence we would have extracted it from $V(p)$, which is a contradiction with the fact that $x \in \mathcal{G}(p)$). Since $P \cup V(p)$ has a smaller dimension than P , there *must* be some infinite sequence $\{y_n\}$ in $P \cup V(p)$, i.e. in $q\text{-Allow}(x) \setminus V(p)$, such that $y_n \rightarrow x$.

The rest of the argument goes through just as in Theorem 4.17. \square

Finally, as in Section 4.1, we have a corollary.

Corollary 5.25. *Let p be a polynomial over \mathcal{S} with unallowable variety $V(p)$. Given any algorithm with output function $p_{\text{comp}}(\cdot)$, a point $x \in \mathcal{G}(p)$, $\epsilon > 0$, and $\eta < 1$, there exists a set Δ_x of positive measure arbitrarily close to x and a set $\Delta \in H_\epsilon$ of positive measure, such that $|p_{\text{comp}} - p|/|p|$ is strictly larger than η when computed at a point $y \in \Delta_x$ with errors $\delta \in \Delta$.*

The proof is based on the topology of \mathcal{S} , and is identical to the proof of Corollary 4.19; we choose not to repeat it.

Proof of Theorem 5.16. Follows immediately from Theorem 5.23 and Corollary 5.25. \square

5.2 Sufficiency: the complex case

In this section we obtain a sufficiency condition for the evaluability of a complex polynomial, given a black-box arithmetic with operations $q_{-2}, q_{-1}, \dots, q_k$.

Throughout this section, we assume our black-box operations include q^c , which consists of multiplication by an a complex constant: $q^c(x) = c \cdot x$. Note that this operation is natural, and that most computers perform it with relative accuracy.

We believe that the sufficiency condition we obtain here is sub-optimal in general, but it subsumes the sufficiency condition we found for the basic complex case with classical arithmetic $\{+, -, \cdot\}$.

We assume that the black-box polynomials defining the operations q_1, \dots, q_k are *irreducible*.

Lemma 5.26. *The varieties $V_{\mathcal{I}}(q_j)$ are irreducible for any $j \in \{1, \dots, k\}$ and any \mathcal{I} as in Definition 5.4 if and only if q_1, \dots, q_k are all affine polynomials.*

Proof. If q_j is an affine polynomial then any $V_{\mathcal{I}}(q_j)$ is also an affine polynomial, hence irreducible over \mathbb{C}^n . Conversely, if q_j is not an affine polynomial, then by inputting a single value x for all the variables, we obtain a one-variable polynomial of degree at least 2, which is necessarily reducible over \mathbb{C}^n . \square

We state here the best sufficiency condition for the accurate evaluability of a polynomial we were able to find in the general case, and a necessary and sufficient condition for the all-affine black-boxes case.

Theorem 5.27 (General case). *Given a polynomial $p : \mathbb{C}^n \rightarrow \mathbb{C}$ with $V(p)$ a union of intersections of hyperplanes Z_i, S_{ij}, D_{ij} , and varieties $V(q_j)$, for $j \in \{1, \dots, k\}$, then p is accurately evaluable.*

Theorem 5.28 (Affine case). *If all black-boxes q_j are affine, then a polynomial $p : \mathbb{C}^n \rightarrow \mathbb{C}$ is accurately evaluable iff $V(p)$ is a union of intersections of hyperplanes Z_i, S_{ij}, D_{ij} , and varieties $V_{\mathcal{I}}(q_j)$, for $j \in \{1, \dots, k\}$ and \mathcal{I} as in Definition 5.4.*

We will begin by proving Theorem 5.27. We will once again make use of Theorem 4.22 and of Theorem 5.24.1.

Lemma 5.29. *If $V(p)$ is as in Theorem 5.27, then $V(p)$ is a simple union of hyperplanes Z_i, S_{ij}, D_{ij} and varieties $V(q_j)$ (with no intersections).*

Proof. Indeed, if that were not the case, then some irreducible q -allowable component P of $V(p)$ would be an intersection of two or more sets described in Theorem 5.27. If P were contained in the intersection of two or more (distinct) hyperplanes, its dimension would be smaller than $n - 2$, and we would get a contradiction.

Suppose now that P was contained in the intersection of a $V(q_j)$ with some other variety or hyperplane. All such varieties, by Theorem 4.22, must have dimension $n - 1$, and since all such varieties and hyperplanes are irreducible, by Theorem 5.24.1, their intersection must have dimension strictly smaller than $n - 2$. Contradiction; we have thus proved that the variety $V(p)$ is a *simple* union of hyperplanes Z_i, D_{ij}, S_{ij} , and varieties $V(q_j)$. \square

Corollary 5.30. *If $p : \mathbb{C}^n \rightarrow \mathbb{C}$ is a polynomial whose variety $V(p)$ is q -allowable, then it is a product $p = c \prod_j p_j$, where each p_j is a power of $x_i, (x_i - x_j), (x_i + x_j)$, or q_j , and c is a complex constant.*

Proof. By Lemma 5.29, the variety $V(p)$ is a union of basic q -allowable hyperplanes and varieties $V(q_j)$.

Choose an irreducible q -allowable set in the union. If this set is a hyperplane, then by following the same argument as in Corollary 4.24, we obtain that p factors into some \tilde{p} and some power of either $x_i, (x_i - x_j)$, or $(x_i + x_j)$.

Suppose now that the irreducible q -allowable set were a variety $V(q_j)$; since p is 0 whenever q_j is 0 and q_j is irreducible, it follows that q_j divides p . We factor then p into the largest power of q_j which divides it, and some other polynomial \tilde{p} .

In either of the two cases, we proceed by factoring \tilde{p} in the same fashion, until we encounter a polynomial \tilde{p} of degree 0. That polynomial is the constant c . \square

Proof of Theorem 5.27. By Corollary 4.24, $p = c \prod_j p_j$, with each p_j a power of $x_k, (x_k \pm x_l)$, or q_l .

Since each of the factors is accurately evaluable, the algorithm that forms their product evaluates p accurately. Multiplication by c (corresponding to the black-box q^c) is also accurate, hence p is accurately evaluable. \square

The proof of Theorem 5.28 follows the path described above; we sketch it here.

Proof of Theorem 5.28. The key fact in obtaining this condition is the irreducibility of all sets $V_{\mathcal{I}}(q_j)$, which is guaranteed by Lemma 5.26, together with the result of Lemma 5.29. Once again, we can write the polynomial as a product of powers of $x_i, (x_i \pm x_j)$, or $V_{\mathcal{I}}(q_j)$, times a constant; this takes care of the sufficiency part, while the necessity follows from Theorem 5.16. \square

Remark 5.31. *Note that Theorem 5.28 is a more general necessary and sufficient condition than Theorem 4.21, which corresponds to having three affine black-boxes (i.e., q_{-2}, q_{-1}, q_0), and we restrict the polynomials to have integer coefficients (thus eliminating the need for q^c).*

6 Accurate Linear Algebra in Rounded Arithmetic

Now we describe implications of our results to the question of whether we can accurately do numerical linear algebra on structured matrices. By a *structured matrix* we mean a family of n -by- n matrices M whose entries $M_{ij}(x)$ are simple polynomial or rational functions of parameters x . Typically there are only $O(n)$ parameters, and the polynomials $M_{ij}(x)$ are closely related (for otherwise little can be said). Typical examples include Cauchy matrices ($M_{ij}(x, y) = 1/(x_i + y_j)$), Vandermonde matrices ($M_{ij}(x) = x_i^{j-1}$), generalized Vandermonde matrices ($M_{ij}(x) = x_i^{j-1+\lambda_j}$, where the λ_j are a nondecreasing sequence of nonnegative integers), Toeplitz matrices ($M_{ij}(x) = x_{i-j}$), totally positive matrices (where M is expressed as a product of simple nonnegative bidiagonal matrices arising from its Neville factorization), acyclic matrices, suitably discretized elliptic partial differential operations, and so on [20, 24, 23, 21, 19, 25, 27, 26, 40, 41].

It has been recently shown that all the matrices on the above list (except Toeplitz and non-totally-positive generalized Vandermonde matrices) admit accurate algorithms in rounded arithmetic for many or all of the problems of numerical linear algebra:

- computing the determinant
- computing all the minors
- computing the inverse
- computing the triangular factorization from Gaussian elimination, with various kinds of pivoting
- computing eigenvalues
- computing singular values

We have gathered together these results in Table 1.

The proliferation of these accurate algorithms for some but not all matrix structures motivates us to ask for which structures they exist.

To convert this to a question about polynomials, we begin by noting that being able to compute the determinant accurately is a *necessary* condition for most of the above computations. For example, if the diagonal entries of a triangular factorization of A , or its eigenvalues, are computable with small relative error, then so is their product, the determinant.

It is also true that being able to compute all the minors accurately is a *sufficient* condition for many of the above computations. For the inverse and triangular factorization, this follows from Cramer's rule and Sylvester's theorem, resp., and for the singular values an algorithm is described in [21, 23].

Table 1: General Structured Matrices

Type of matrix	$\det A$	A^{-1}	Any minor	LDU	SVD	Sym EVD
Acyclic (bidiagonal and other)	n [21]	n^2 [21]	n [21]	$\leq n^2$ [21]	n^3 [21]	N/A
Total Sign Compound (TSC)	n [21]	n^3 [21]	n [21]	n^4 [21]	n^4 [21]	n^4 [29]
Diagonally Scaled Totally Unimodular (DSTU)	n^3 [21]	n^5	n^3 [21]	n^3 [21]	n^3 [21]	n^3 [29]
Weakly diagonally dominant M-matrix	n^3 [48]	n^3 [4, 3]	No [23]	n^3 [25]	n^3 [25]	n^3 [29]
Displace- ment Rank One	Cauchy	n^2	n^2	$\leq n^3$ [19]	n^3 [19]	n^3 [29]
	Vandermonde	n^2	No [23]	No [23]	n^3 [19, 27]	n^3 [29]
	Polynomial Vandermonde	n^2 [33]	No Section 6.1	No Section 6.1	No Section 6.1	* [27]

Thus, if the determinants $p_n(x) = \det M^{n \times n}(x)$ of a class of n -by- n structured matrices M do not satisfy the necessary conditions described in Theorem 5.16 for *any* enumerable set of black-box operations (perhaps with other properties, like bounded degree), then we can conclude that accurate algorithms of the sort described in the above citations are impossible.

In particular, to satisfy these necessary conditions would require that the varieties $V(p_n)$ be allowable (or q -allowable). For example, if V is a Vandermonde matrix, then $\det(V) = \prod_{i < j} (x_i - x_j)$ satisfies this condition, using only subtraction and multiplication.

The following theorem states a condition which guarantees the impossibility of an algorithm using *any* enumerable set of black-box operations of bounded degree:

Theorem 6.1. *Let $M(x)$ be an n -by- n structured complex matrix with determinant $p_n(x)$ as described above. Suppose $p_n(x)$ has an irreducible factor $\hat{p}_n(x)$ whose degree goes to infinity as n goes to infinity. Then for any enumerable set of black-box arithmetic operations of bounded degree, for sufficiently large n it is impossible to accurately evaluate $p_n(x)$ over the complex numbers.*

Proof. Let q_1, \dots, q_m be any finite set of black-box operations. To obtain a contradiction, suppose the complex variety $V(p_n)$ satisfies the necessary conditions of Theorem 5.16, i.e., that $V(p_n)$ is allowable. This means that $V(p_n)$, which includes the hypersurface $V(\hat{p}_n)$ as an irreducible component, can be written as the union of irreducible q -allowable sets (by Def. 5.10). This means that $V(\hat{p}_n)$ must itself be equal to an irreducible q -allowable set (a hypersurface), since representations as unions of irreducible sets are unique. The irreducible q -allowable sets of codimension 1 are defined by single irreducible polynomials, which are in turn derived by the process of setting variables equal to one another, to one another's negation, or zero (as described in Defs. 5.4 and 5.10), and so have bounded degree. This contradicts the unboundedness of the degree of $V(\hat{p}_n)$. \square

In the next theorem we apply this result to the set of complex Toeplitz matrices. We use the following notation. Let T be an n -by- n Toeplitz matrix, with x_j on the j -th diagonal, so x_0 is on the main diagonal, x_{n-1} is in the top right corner, and x_{1-n} is in the bottom left corner.

Theorem 6.2. *The determinant of a Toeplitz matrix T is irreducible over any field.*

Corollary 6.3. *The determinants of the set of complex Toeplitz matrices cannot be evaluated accurately using any enumerable set of bounded-degree black-box operations.*

Proof of Theorem 6.2. We use induction on n . We note that $\det T$ depends on every variable x_j , because $\det T$ includes the monomials $\pm x_j^{n-j} x_{j-n}^j$ for $j > 0$, as well as x_0^n , and these monomials contain the maximum powers of x_j and x_{j-n} appearing in the determinant. Now x_{n-1} appears exactly once in T , so $\det T$ must be an affine function of x_{n-1} , say $\det T = x_{n-1} \cdot p_{1n} + p_{2n}$. By expanding $\det T$ along the first row or column, we see that p_{1n} is itself the determinant of a Toeplitz matrix with diagonals x_{1-n}, \dots, x_{n-3} , and p_{2n} depends on x_{1-n}, \dots, x_{n-2} but not x_{n-1} . If $\det T = x_{n-1} \cdot p_{1n} + p_{2n}$ were reducible, its factorization would have to look like $x_{n-1} \cdot p_{1n} + p_{2n} = (x_{n-1} \cdot p_{3n} + p_{4n})p_{5n}$, where all the subscripted p polynomials are independent of x_{n-1} , implying either that $p_{1n} = p_{3n}p_{5n}$ were reducible, a contradiction by our induction hypothesis, or $p_{3n} = 1$ and so $p_{1n} | p_{2n}$. Now we can write $p_{2n} = x_{n-2}^2 q_{1n} + x_{n-2} q_{2n} + q_{3n}$ where $q_{1n} \neq 0$, since $\det T$ includes the monomial $\pm x_{n-2}^2 x_{n-2}^{n-2}$ and no higher powers of x_{n-2} . Furthermore q_{1n} is independent of x_{n-1} , x_{n-2} and x_{n-3} , and q_{2n} and q_{3n} are independent of x_{n-1} and x_{n-2} . Since p_{1n} is independent of x_{n-2} , the only way we could have $p_{1n} | p_{2n}$ is to have $p_{1n} | q_{1n}$, $p_{1n} | q_{2n}$, and $p_{1n} | q_{3n}$. But since p_{1n} depends on x_{n-3} and q_{1n} is independent of x_{n-3} , this is a contradiction. So the determinant of a Toeplitz matrix must be irreducible. \square

In the real case, irreducibility of p_n is not enough to conclude that p_n cannot be evaluated accurately, because $V_{\mathbb{R}}(p_n)$ may still be allowable (and even vanish). So we consider another necessary condition for allowability: Since all black-boxes have a finite number of arguments, their associated codimension-1 irreducible components must have the property that whether $x \in V_{\mathcal{I}}(q_j)$ depends on only a finite number of components of x . Thus to prove that the hypersurface $V_{\mathbb{R}}(p_n)$ is not allowable, it suffices to find at least one regular point x^* in $V_{\mathbb{R}}(p_n)$ such that the tangent hyperplane at x^* is not parallel to sufficiently many coordinate directions, i.e., membership in $V_{\mathbb{R}}(p_n)$ depends on more variables than any $V_{\mathcal{I}}(q_j)$. This is easy to do for real Toeplitz matrices.

Theorem 6.4. *Let V be the variety of the determinant of real singular Toeplitz matrices. Then V has codimension 1, and at almost all regular points, its tangent hyperplane is parallel to no coordinate directions.*

Corollary 6.5. *The determinants of the set of real Toeplitz matrices cannot be evaluated accurately using any enumerable set of bounded-degree black-box operations.*

Proof of Theorem 6.4. Let $\text{Toep}(i, j)$ denote the Toeplitz matrix with diagonal entries x_i through x_j ; thus $\text{Toep}(i, j)$ has dimension $(j - i)/2 + 1$. Let U be the Zariski open set where $\det \text{Toep}(i, j) \neq 0$ for all $1 - n \leq i \leq j < n - 1$ and $j - i$ even. Then $\det T$ is a nonconstant affine function of x_{n-1} , and so for any choice of x_{1-n}, \dots, x_{n-2} in U , $\det T$ is zero for a unique choice of x_{n-1} . This shows that $V_{\mathbb{R}}(\det T)$ has real codimension 1.

Furthermore, $\det T$ has highest order term in each x_i , $0 < i \leq n - 1$, equal to $\pm \text{Toep}(1 - n, 2i - n - 1) x_i^{n-i}$, i.e., with nonzero coefficient on U . It also has the highest order term in each x_i , $1 - n \leq i < 0$, equal to $\pm \text{Toep}(n + 2i + 1, n - 1) x_i^{n+i}$, i.e., with nonzero coefficient on U . Finally, the highest order term in x_0 is x_0^n , with coefficient 1. Thus the gradient of $\det T$ has all nonzero components on a Zariski open set, and whether $\det T = 0$ depends on all variables. \square

6.1 Vandermonde matrices and generalizations

In this section we will explain the entries filled “**No**” in Table 1. First we will show that polynomial Vandermonde matrices do not have algorithms for computing accurate inverses, by proving that certain minors needed in the expression cannot be accurately computed (this will also explain the “**No**” in the *Any minor* column). Finally, we will show that the LDU factorization for polynomial Vandermonde matrices cannot be computed accurately.

First we consider the class of generalized Vandermonde matrices V , where $V_{ij} = P_{i-1}(x_j)$ is a polynomial function of x_j , with $1 \leq i, j \leq n$. This class includes the standard Vandermonde (where $P_{i-1}(x_j) = x_j^{i-1}$) and many others.

First consider a generalized Vandermonde matrix where $P_{i-1}(x_j) = x_j^{i-1+\lambda_{n-i+1}}$, where $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ with $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 0$. The tuple λ is called a *partition*. Any square submatrix of such a generalized Vandermonde matrix is also a generalized Vandermonde matrix. A generalized Vandermonde matrix is known to have determinant of the form $s_\lambda(x) \prod_{i < j} (x_i - x_j)$ where $s_\lambda(x)$ is a polynomial of degree $|\lambda| = \sum_i \lambda_i$, and called a Schur function [43]. In infinitely many variables (not our situation) the Schur function is irreducible [31], but in finitely many variables, the Schur function is sometimes irreducible and sometimes not [57, Exer. 7.30]. But there are irreducible Schur functions of arbitrarily high degree. Thus we conclude by Theorem 6.1 that no enumerable set of black-box operations of bounded degree can compute all Schur functions accurately when the x_i are complex.

If we restrict the domain \mathcal{D} to be nonnegative real numbers, then the situation changes: The nonnegativity of the coefficients of the Schur functions shows that they are positive in \mathcal{D} , and indeed the generalized Vandermonde matrix is totally positive [38]. Combined with the homogeneity of the Schur function, Theorem 3.2 implies that the Schur function, and so determinants (and minors) of totally positive generalized Vandermonde matrices can be evaluated accurately in classical arithmetic. For accurate algorithms that are more efficient than the one in Theorem 3.2, see [28].

Now consider a polynomial Vandermonde matrix V_P defined by a family $\{P_k(x)\}_{k \in \mathbb{N}}$ of polynomials such that $\deg(P_k) = k$, and $V_P(i, j) = P_{i-1}(x_j)$.

Note that any V_P can be written as $V_P = VC$, with V being a regular Vandermonde matrix, and C being an upper triangular matrix of coefficients of the polynomials P_k , i.e.,

$$P_{j-1}(x) = \sum_{i=1}^j C(i, j) x^{i-1}, \quad \forall 1 \leq j \leq n.$$

Furthermore assume we can write $C = \tilde{D}\tilde{C}$, with \tilde{D} being the diagonal matrix of highest-order coefficients, i.e., $\tilde{D}(i, i) = C(i, i)$ for all $1 \leq i \leq n$. We will assume that the matrices C and \tilde{D} are given to us exactly.

Denote by $c_{i-1} := \tilde{D}(i, i)$, for all $1 \leq i \leq n$ the highest-order coefficients of the polynomials $P_0(x), \dots, P_{n-1}(x)$.

To compute the inverse of the matrix V_P , we need to compute the minors that result from deleting a row and a column, i.e., (in MATLAB notation) $\det(V_P([1:i-1, i+1:n], [1:j-1, j+1:n]))$. We will focus our attention on the computation of the $(i, n-1)$ minors, i.e., the ones that result from deleting any of the rows and the $(n-1)$ st column.

The resulting matrices look like

$$M_{P,i} := V_P([1:i-1, i+1:n], [1:n-2, n]) = \begin{bmatrix} c_0 & P_1(x_1) & \dots & P_{n-3}(x_1) & P_{n-1}(x_1) \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ c_0 & P_1(x_{i-1}) & \dots & P_{n-3}(x_{i-1}) & P_{n-1}(x_{i-1}) \\ c_0 & P_1(x_{i+1}) & \dots & P_{n-3}(x_{i+1}) & P_{n-1}(x_{i+1}) \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ c_0 & P_1(x_n) & \dots & P_{n-3}(x_n) & P_{n-1}(x_n) \end{bmatrix} .$$

Hence, we can manipulate the columns of $\det(M_{P,i})$ by subtracting from them linear combinations of other columns, to obtain

$$\det(M_{P,i}) = \begin{vmatrix} c_0 & c_1x_1 & c_2x_1^2 & \dots & c_{n-3}x_1^{n-2} & c_{n-1}x_1^{n-1} + C(n-1, n)x_1^{n-2} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ c_0 & c_1x_{i-1} & c_2x_{i-1}^2 & \dots & c_{n-3}x_{i-1}^{n-2} & c_{n-1}x_{i-1}^{n-1} + C(n-1, n)x_{i-1}^{n-2} \\ c_0 & c_1x_{i+1} & c_2x_{i+1}^2 & \dots & c_{n-3}x_{i+1}^{n-2} & c_{n-1}x_{i+1}^{n-1} + C(n-1, n)x_{i+1}^{n-2} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ c_0 & c_1x_n & c_2x_n^2 & \dots & c_{n-3}x_n^{n-2} & c_{n-1}x_n^{n-1} + C(n-1, n)x_n^{n-2} \end{vmatrix}$$

By expanding on the last column, and using the results from [23], we obtain that there are constants E and F , specifically $E = \prod_{i=1}^n c_{i-1}$ and $F = E \frac{C(n-1, n)}{c_{n-1}}$, such that

$$\det(M_{P,i}) = E \prod_{\substack{k < j \\ k, j \neq i}} (x_j - x_k) + F \cdot s_{[1]}(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n) ,$$

with $s_{[1]}$ being the Schur function corresponding to the partition $\lambda = [1]$, i.e.,

$$\det(M_{P,i}) = E \prod_{\substack{k < j \\ k, j \neq i}} (x_j - x_k) + F \cdot (x_1 + \dots + x_{i-1} + x_{i+1} + \dots + x_n) ,$$

and it is not hard to see that for any $n \geq 4$, the above polynomial in $x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n$ *does not have an allowable variety*, and hence the inverse *cannot be evaluated accurately*.

Moreover, if the above determinant contains irreducible factors of arbitrarily high degree, then, by Theorem 6.1, this family of matrices has inverses that cannot be evaluated accurately even with the addition of any enumerable set of bounded-degree black-boxes.

This explains why we have filled in with “**No**” the entries corresponding to columns “ A^{-1} ” and “*Any minor*” in the Polynomial Vandermonde row of Table 1. Below we explain the “**No**” in the same of row of column “ LDU ”.

If we let $V_P = L_P D_P U_P$ and $V = LDU$, it follows that

$$\begin{aligned} L_P &= L ; \\ D_P &= D\tilde{D} ; \\ U_P &= \tilde{D}^{-1}UC . \end{aligned}$$

Since we cannot compute L accurately in the general Vandermonde case, it follows that we cannot compute L_P accurately in the polynomial Vandermonde case.

Finally, we explain the “*” entries in the polynomial Vandermonde row. These depend on special properties of the polynomial. In general, neither the SVD nor the symmetric eigenvalue decomposition (EVD) are computable accurately, but if the polynomials are certain orthogonal polynomials, then the accurate SVD is possible [27], and symmetric EVD may also be possible [29].

Acknowledgments

We are grateful to Bernd Sturmfels, Gautam Bharali, Plamen Koev, William (Velvel) Kahan, and Gregorio Malajovich for interesting discussions and to Jonathan Dorfman for his help with the graphs.

References

- [1] A. Aho, J. Hopcroft, and J. Ullman. *The design and analysis of computer algorithms*. Addison-Wesley, 1974.
- [2] G. Alefeld and J. Herzberger. *Introduction to interval computations*. Academic Press, 1983.
- [3] S. A. Attahiru, X. Junggong, and Q. Ye. Accurate computation of the smallest eigenvalue of a diagonally dominant M-matrix. preprint, 2002.
- [4] S. A. Attahiru, X. Junggong, and Q. Ye. Entrywise perturbation theory for diagonally dominant M-matrices with applications. *Numer. Math.*, 90:401–414, 2002.
- [5] D. H. Bailey. Multiprecision translation and execution of Fortran programs. *ACM Trans. Math. Soft.*, 19(3):288–319, Sept. 1993.
- [6] D. H. Bailey. A Fortran-90 based multiprecision system. *ACM Trans. Math. Soft.*, 21(4):379–387, Dec. 1995.
- [7] L. Blum. Computing over the reals: Where Turing meets Newton. *Notices of the AMS*, 51(9), October 2004.
- [8] L. Blum, F. Cucker, M. Shub, and S. Smale. Complexity and real computation: A manifesto. *Intern. J. Bifurcation and Chaos*, 6(1):3–26, 1996.
- [9] L. Blum, F. Cucker, M. Shub, and S. Smale. *The Mathematics of Numerical Analysis*, volume 32 of *Lecture Notes in Applied Mathematics*, chapter Algebraic Settings for the Problem “P \neq NP?”. AMS, 1996.
- [10] L. Blum, F. Cucker, M. Shub, and S. Smale. *Complexity and Real Computation*. Springer, 1997.
- [11] L. Blum, M. Shub, and S. Smale. On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *AMS Bulletin (New Series)*, 21(1):1–46, July 1989.

- [12] F. Chatelin and V. Frayssé. *Lectures on Finite Precision Computations*. SIAM, Philadelphia, PA, 1996.
- [13] F. Cucker and J. P. Dedieu. Decision problems and round-off machines. *Theory of Computing Systems*, 34:433–452, 2001.
- [14] F. Cucker and D. Grigoriev. Complexity lower bounds for approximate algebraic computation trees. *J. Complexity*, 15:499–512, 1999.
- [15] F. Cucker and S. Smale. Complexity estimates depending on condition and roundoff error. *J. ACM*, 46(1):113–184, Jan 1999.
- [16] J. Demmel. Underflow and the reliability of numerical software. *SIAM J. Sci. Stat. Comput.*, 5(4):887–919, Dec 1984.
- [17] J. Demmel. On error analysis in arithmetic with varying relative precision. In *Proceedings of the 8th Symposium on Computer Arithmetic*, Como, Italy, May 1987. IEEE Computer Society Press.
- [18] J. Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.
- [19] J. Demmel. Accurate SVDs of structured matrices. *SIAM J. Mat. Anal. Appl.*, 21(2):562–580, 1999.
- [20] J. Demmel. The complexity of accurate floating point computation. In *Proceedings of the 2002 International Congress of Mathematicians*, Beijing, 2002.
- [21] J. Demmel, M. Gu, S. Eisenstat, I. Slapničar, K. Veselić, and Z. Drmač. Computing the singular value decomposition with high relative accuracy. *Lin. Alg. Appl.*, 299(1–3):21–80, 1999.
- [22] J. Demmel and Y. Hida. Accurate and efficient floating point summation. *SIAM J. Sci. Comp.*, 25(4):1214–1248, 2003.
- [23] J. Demmel and P. Koev. Necessary and sufficient conditions for accurate and efficient rational function evaluation and factorizations of rational matrices. In V. Olshevsky, editor, *Special Issue on Structured Matrices in Mathematics, Computer Science and Engineering*, volume 281 of *Contemporary Mathematics*, pages 117–145. AMS, 2001.
- [24] J. Demmel and P. Koev. Accurate and efficient algorithms for floating point computation. In *Proceedings of the 2003 International Congress of Industrial and Applied Mathematics*, Sydney, 2004.
- [25] J. Demmel and P. Koev. Accurate SVDs of weakly diagonally dominant M-matrices. *Numer. Math.*, 98:99–104, 2004.
- [26] J. Demmel and P. Koev. The accurate and efficient solution of a totally positive generalized Vandermonde linear system. *SIAM J. Matrix Anal. Appl.*, 27(1):142–152, 2005.
- [27] J. Demmel and P. Koev. Accurate SVDs of polynomial Vandermonde matrices involving orthonormal polynomials. to appear in *SIAM J. Matrix Anal. Appl.*; www-math.mit.edu/~plamen/files/chebvand.ps, 2005.

- [28] J. Demmel and P. Koev. Efficient and accurate evaluation of Schur and Jack polynomials. to appear in *Math. Comp.*, math.mit.edu/~plamen/files/s2.pdf, 2005.
- [29] Froilán M. Dopico, Juan M. Molera, and Julio Moro. An orthogonal high relative accuracy algorithm for the symmetric eigenproblem. *SIAM J. Matrix Anal. Appl.*, 25(2):301–351 (electronic), 2003.
- [30] A. Edelat and P. Sünderhauf. A domain-theoretic approach to real number computation. *Theoretical Computer Science*, 210:73–98, 1998.
- [31] H. K. Farahat. On Schur functions. *Proc. London Math. Soc. (3)*, 8:621–630, 1958.
- [32] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, resultants, and multi-dimensional determinants*. Birkhäuser Boston, 1994.
- [33] N. J. Higham. Stability analysis of algorithms for solving confluent Vandermonde-like systems. *SIAM J. Mat. Anal. Appl.*, 11(1):23–41, 1990.
- [34] N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, Philadelphia, PA, 2nd edition, 2002.
- [35] K. Hulek. *Elementary algebraic geometry*. American Math. Soc., 2003. translated by H. Verrill.
- [36] W. Kahan. www.cs.berkeley.edu/~wkahan.
- [37] W. Kahan. The Improbability of Probabilistic Error Analysis. UC Berkeley Statistics Colloquium, www.cs.berkeley.edu/~wkahan/improber.pdf, 1998.
- [38] S. Karlin. *Total Positivity*. Stanford University Press, 1968.
- [39] K.-I. Ko. *Complexity theory of real functions*. Birkhäuser, 1991.
- [40] P. Koev. Accurate computations with totally nonnegative matrices. submitted to *SIAM J. Matrix Anal. Appl.* www-math.mit.edu/~plamen/files/acctp.pdf, 2004.
- [41] P. Koev. Accurate eigenvalues and SVDs of totally nonnegative matrices. *SIAM J. Matrix Anal. Appl.*, 27(1):1–23, 2005.
- [42] D. Lozier and F. Olver. Closure and precision in level-index arithmetic. *SIAM J. Num. Anal.*, 27(5), 1990.
- [43] I. G. MacDonald. *Symmetric functions and Hall polynomials*. Oxford University Press, 2nd edition, 1995.
- [44] Ezra Miller and Bernd Sturmfels. *Algebraic Combinatorics*. Springer-Verlag, 2005.
- [45] R. K. Montoye, E. Hokenek, and S. L. Runyon. Design of the IBM RISC System 6000 floating point execution unit. *IBM Journal of Research and Development*, 34(1):59–70, 1990.
- [46] R. E. Moore. *Methods and applications of interval analysis*. SIAM, Philadelphia, 1979.
- [47] A. Neumaier. *Interval Methods for Systems of Equations*. Cambridge University Press, Cambridge, England, 1990.

- [48] C. O’Cinneide. Relative-error for the LU decomposition via the GTH algorithm. *Numer. Math.*, 73:507–519, 1996.
- [49] M. Pour-El and J. Richards. *Computability in Analysis and Physics*. Springer-Verlag, 1989.
- [50] D. Priest. Algorithms for arbitrary precision floating point arithmetic. In P. Kornerup and D. Matula, editors, *Proceedings of the 10th Symposium on Computer Arithmetic*, pages 132–145, Grenoble, France, June 26-28 1991. IEEE Computer Society Press.
- [51] J. Renegar. On the computational complexity and geometry of the first-order theory of the reals: Parts I, II and III. *J. Symb. Comp.*, 13, 1992.
- [52] J. Renegar. Is it possible to know a problem instance is ill-posed? *J. Complexity*, 10:1–56, 1994.
- [53] B. Reznick. *Some concrete aspects of Hilbert’s 17th problem*, volume 253 of *Contemporary Mathematics*. Amer. Math. Society, 2000.
- [54] I. R. Shafarevich. *Basic Algebraic Geometry*. Springer, 2nd edition, 1994.
- [55] Jonathan Richard Shewchuk. Adaptive Precision Floating-Point Arithmetic and Fast Robust Geometric Predicates. *Discrete & Computational Geometry*, 18:305–363, 1997.
- [56] D. Spielman and S.-H. Teng. Smoothed analysis of algorithms. In *Proceedings of the 2002 International Congress of Mathematicians*, Beijing, 2002.
- [57] Richard Stanley. *Enumerative Combinatorics*, volume 2. Cambridge Univ. Press, 1999.
- [58] A. Tarski. *A decision method for elementary algebra and geometry*. University of California Press, Berkeley, 1951.
- [59] J. Taylor. *Several complex variables with connections to algebraic geometry and Lie groups*. AMS Series on Graduate Studies in Mathematics. American Math Society, 2004.
- [60] A. Turing. Rounding-off errors in matrix processes. *Quart. J. Mech. Appl. Math.*, 1:287–308, 1948.
- [61] J. Vignes. A stochastic arithmetic for reliable computation. *Math. and Comp. in Sim.*, 35:233–261, 1993.
- [62] J. von Neumann and H. Goldstine. Numerical inverting matrices of high order. *Bull. AMS*, 53:1021–1099, 1947.
- [63] J. H. Wilkinson. *Rounding Errors in Algebraic Processes*. Prentice Hall, Englewood Cliffs, 1963.
- [64] Günter Ziegler. *Lectures on Polytopes*. Springer-Verlag, 1995.