Mathematical Programming manuscript No.
(will be inserted by the editor)

Pablo A. Parrilo

Semidefinite programming relaxations for semialgebraic problems

Abstract. A hierarchy of convex relaxations for semialgebraic problems is introduced. For questions reducible to a finite number of polynomial equalities and inequalities, it is shown how to construct a complete family of polynomially sized semidefinite programming conditions that prove infeasibility. The main tools employed are a semidefinite programming formulation of the sum of squares decomposition for multivariate polynomials, and some results from real algebraic geometry. The techniques provide a constructive approach for finding bounded degree solutions to the Positivstellensatz, and are illustrated with examples from diverse application fields.

Key words. Semidefinite programming – convex optimization – sums of squares – polynomial equations – real algebraic geometry.

1. Introduction

Numerous questions in applied mathematics can be formally expressed using a finite number of polynomial equalities and inequalities. Well-known examples are optimization problems with polynomial objective and constraints, such as quadratic, linear, and boolean programming. This is a fairly broad class, including problems with a combination of continuous and discrete variables, and easily seen to be NP-hard in the general case.

In this paper we introduce a new approach to the formulation of computable relaxations for this kind of problems. The crucial enabling fact is the computational tractability of the sum of squares decomposition for multivariate polynomials, coupled with powerful results from semialgebraic geometry. As a result, a whole new class of convex approximations for semialgebraic problems is obtained. The results generalize in a very natural way existing successful approaches, including the well-known semidefinite relaxations for combinatorial optimization problems.

The paper includes notions from traditionally separated research areas, namely numerical optimization and real algebra. In the interest of achieving the broadest possible communication of the main ideas, we have tried to make this article as self-contained as possible, providing a brief introduction to both semidefinite programming and real algebra. It is our belief that there is a lot of potential in the interaction between these fields, particularly with regard to practical applications. Most of the material in the pa-

Automatic Control Laboratory, Swiss Federal Institute of Technology (ETH), CH-8092 Zurich, Switzerland. e-mail: parrilo@aut.ee.ethz.ch. The majority of this research has been carried out while the author was with the Control & Dynamical Systems Department, California Institute of Technology, Pasadena, CA 91125, USA.

per is from the author's dissertation [Par00b], with the addition of new examples and references.

The paper is organized as follows: in Section 2 the problem of global nonnegativity of polynomial functions is introduced, and existing approaches are discussed. The sum of squares decomposition is presented as a sufficient condition for nonnegativity. In Section 3 a brief review of semidefinite programming is presented, and it is shown how to compute sum of squares decompositions by solving a semidefinite program. In the following section, some basic elements of semialgebraic geometry are described, and the Positivstellensatz is stated. Our main result (Theorem 5.1) follows, showing how the sum of squares decision procedure allows for the search of bounded degree solutions to the Positivstellensatz equation. We present next a refutation-based interpretation of the methodology, as well as a comparison with earlier related work. Section 6 contains some observations on the computational aspects of the implementation of the techniques. In Section 7, a sample of applications from different applied mathematics areas are presented. These include, among others, enhanced semidefinite relaxations for quadratic programming problems, and stronger conditions for matrix copositivity.

1.1. Notation

The notation is mostly standard. The inner product between two vectors in $\mathbb{R}^n$ is defined as $\langle x, y\rangle := \sum_{i=1}^n x_i y_i$. Let $\mathcal{S}^n \subset \mathbb{R}^{n\times n}$ be the space of symmetric $n \times n$ real matrices, with inner product between $X, Y \in \mathcal{S}^n$ being $\langle X, Y\rangle :=$ trace XY. A matrix $M \in \mathcal{S}^n$ is *positive semidefinite* (PSD) if $x^T M x \geq 0$, $\forall x \in \mathbb{R}^n$. Equivalently, M is positive semidefinite if all its eigenvalues are nonnegative. Let $\mathcal{S}^n_+$ be the self-dual cone of positive semidefinite matrices, with the notation $A \succeq B$ indicating that $A - B$ is positive semidefinite.

2. Global nonnegativity

A fundamental question appearing in many areas of applied mathematics is that of checking global nonnegativity of a function of several variables. Concretely, given a function F, we have the following:

Problem 2.1. Provide checkable conditions or a procedure for verifying the validity of the proposition

$$F(x_1, \ldots, x_n) \geq 0, \qquad \forall x_1, \ldots, x_n \in \mathbb{R}. \tag{2.1}$$

This is a fundamental mathematical problem, that appears in numerous application domains, and towards which considerable research efforts have been devoted. In order to study the problem from a computational viewpoint, and avoid undecidability results, it is clear that further restrictions on the class of functions F should be imposed. However, at the same time we would like to keep the problem general enough, to enable the practical applicability of the results. A good compromise is achieved by considering the case of *multivariate polynomials*.

Definition 2.2. *A* polynomial f *in* $x_1, \ldots, x_n$ *is a finite linear combination of monomials:*

$$f = \sum_\alpha c_\alpha x^\alpha = \sum_\alpha c_\alpha x_1^{\alpha_1} \ldots x_n^{\alpha_n}, \qquad c_\alpha \in \mathbb{R}, \tag{2.2}$$

where the sum is over a finite number of n-tuples $\alpha = (\alpha_1, \ldots, \alpha_n)$, $\alpha_i \in \mathbb{N}_0$. The set of all polynomials in $x_1, \ldots, x_n$ with real coefficients is written as $\mathbb{R}[x_1, \ldots, x_n]$.

The *total degree* of the monomial x^α is equal to $\alpha_1 + \cdots + \alpha_n$. The total degree of a polynomial is equal to the highest degree of its component monomials.

An important special case is that of *homogeneous polynomials* (or *forms*), where all the monomials have the same total degree.

Definition 2.3. *A* form *is a polynomial where all the monomials have the same total degree d. In this case, the polynomial is* homogeneous *of degree d, since it satisfies $f(\lambda x_1, \ldots, \lambda x_n) = \lambda^d f(x_1, \ldots, x_n)$.*

It is well-known that there is a correspondence between forms and polynomials. A form in n variables and degree m can be dehomogenized to a polynomial in $n-1$ variables, of degree less than or equal to m, by fixing any variable to the constant value 1. Conversely, given a polynomial, it can be converted into a form by multiplying each monomial by powers of a new variable, in such a way that the total degree of all monomials are the same.

The set of forms in n variables and degree m can be associated with a vector space of dimension $\binom{n+m-1}{m}$. Similarly, the set of polynomials of total degree less than or equal to m is a vector space of dimension $\binom{n+m}{m}$. These quantities will be important later in the study of the efficiency of the computational implementation of the proposed methodology.

2.1. Exact and approximate approaches

It is a fact that many problems in applied mathematics can be formulated using only polynomial equalities and inequalities, which are satisfied if and only if the problem has a solution. In this regard, Tarski's results on the existence of a decision procedure for elementary algebra over the reals, settles the decidability of Problem 2.1 for this quite large class of problems.

When F is a polynomial, the Tarski-Seidenberg decision procedure [BCR98, Mis93, Bos82] provides an explicit algorithm for deciding if (2.1) holds, so the problem is decidable. There are also a few alternative approaches to effectively answer this question, also based in decision algebra; see [Bos82] for a survey of classical available techniques, and [BPR96] for more efficient recent developments.

Regarding complexity, the general problem of testing global nonnegativity of a polynomial function is NP-hard (when the degree is at least four), as easily follows from reduction from the matrix copositivity problem; see [MK87] and Section 7.5. Therefore, unless P=NP, *any method* guaranteed to obtain the right answer *in every possible instance* will have unacceptable behavior for problems with a large number of variables. This is the main drawback of theoretically powerful methodologies such as quantifier elimination.

If we want to avoid the inherent complexity roadblocks associated with the *exact* solution, an attractive option is to settle for approximate answers, that are "reasonably close" to the original question. The issue therefore arises: are there conditions, that can be efficiently tested, that guarantee global nonnegativity of a function? As we will see in Section 3.2, one such condition is given by the existence of a sum of squares decomposition.

3. Sums of squares and SDP

Before presenting the details of our approach, we take a brief detour in the following subsection to present the basic ideas behind the convex optimization techniques used, namely semidefinite programming.

3.1. Semidefinite programming background

In this section we present a brief introduction to semidefinite programming (SDP). We refer the reader to [VB96] for an excellent survey of the theory and applications, and [WSV00] for a comprehensive treatment of the many aspects of the subject. SDP can be understood as a generalization of linear programming, where the nonnegative orthant constraint in the latter is replaced instead by the cone of positive semidefinite matrices.

A semidefinite program is defined as the optimization problem:

$$\begin{array}{ll} \text{minimize} & \langle C, X\rangle \\ \text{subject to} & \langle A_i, X\rangle = b_i \\ & X \succeq 0, \end{array} \tag{3.1}$$

where $X \in \mathcal{S}^n$ is the decision variable, $b \in \mathbb{R}^m$ and $C, A_i \in \mathcal{S}^n$ are given symmetric matrices. A geometric interpretation is the optimization of a linear functional, over the intersection of an affine subspace and the self-dual cone of positive semidefinite matrices.

The crucial feature of semidefinite programs is its *convexity*, since the feasible set defined by the constraints above is convex. For this reason, semidefinite programs have a nice duality structure, with the associated dual program being:

$$\begin{array}{ll} \text{maximize} & \langle b, y\rangle \\ \text{subject to} & \sum_{i=1}^m y_i A_i \preceq C, \end{array} \tag{3.2}$$

where $y \in \mathbb{R}^m$. Any feasible solution of the dual provides a lower bound on the achievable values of the primal; conversely, feasible primal solutions give upper bounds on dual solutions. This is known as *weak duality* and follows since:

$$\langle C, X\rangle - \langle b, y\rangle = \langle C, X\rangle - \sum_{i=1}^m y_i b_i = \langle C, X\rangle - \sum_{i=1}^m y_i \langle A_i, X\rangle = \langle C - \sum_{i=1}^m y_i A_i, X\rangle \geq 0,$$

with the last inequality being true because of self-duality of the PSD cone. Under standard constraint qualifications (for instance, existence of strictly feasible points), *strong duality* holds, and the primal and the dual problems achieve exactly the same value.

Theorem 3.1. *Consider the primal-dual SDP pair (3.1)-(3.2). If either feasible set has has a nonempty interior, then for every* $\epsilon > 0$*, there exist feasible* X, y *such that* $\langle C, X\rangle - \langle b, y\rangle < \epsilon$*. Furthermore, if both feasible sets have nonempty interiors, then the optimal solutions are achieved by some* $X_\star$*,* $y_\star$*.*

From a computational viewpoint, semidefinite programs can be efficiently solved, both in theory and in practice. In the last few years, research on SDP has experienced an explosive growth, particularly in the areas of algorithms and applications. Two of the main reasons for this practical impact are the versatility of the problem formulation, and the availability of high-quality software, such as SeDuMi [Stu99].

3.2. The sum of squares decomposition

If a polynomial F satisfies (2.1), then an obvious necessary condition is that its degree be an even number. A deceptively simple sufficient condition for a real-valued polynomial $F(x)$ to be nonnegative is the existence of a sum of squares decomposition:

$$F(x) = \sum_i f_i^2(x), \qquad f_i(x) \in \mathbb{R}[x]. \tag{3.3}$$

It is clear that if a given polynomial $F(x)$ can be written as above, for some polynomials f_i, then F is nonnegative for all values of x.

Two questions immediately arise:

- When is such decomposition possible?
- How do we compute it?

For the case of polynomials, the first question is a well-analyzed problem, first studied by David Hilbert more than a century ago. In fact, one of the items in his famous list of twenty-three unsolved problems presented at the International Congress of Mathematicians at Paris in 1900, deals with the representation of a definite form as a sum of squares of rational functions. The reference [Rez00] contains a beautiful survey by Reznick of the fascinating history of this problem, and pointers to most of the available results.

For notational simplicity, we use the notation *SOS* for "sum of squares." Hilbert himself noted that not every nonnegative polynomial is SOS. A simple explicit counterexample is the Motzkin form (here, for $n = 3$)

$$M(x, y, z) = x^4y^2 + x^2y^4 + z^6 - 3x^2y^2z^2. \tag{3.4}$$

Positive semidefiniteness can be easily shown using the arithmetic-geometric inequality (see also Example 7.3), and the nonexistence of a SOS decomposition follows from standard algebraic manipulations (see [Rez00] for details), or the procedure outlined below.

Following the notation in references [CLR95, Rez00], let $P_{n,m}$ be the set of nonnegative forms of degree m in n variables, and $\Sigma_{n,m}$ the set of forms p such that $p = \sum_k h_k^2$, where h_k are forms of degree $m/2$. Hilbert gave a complete characterization of when these two classes are equivalent.

Theorem 3.2 (Hilbert). *Let $P_{n,m}, \Sigma_{n,m}$ be as above. Then $\Sigma_{n,m} \subseteq P_{n,m}$, with equality holding only in the following cases:*

- *Bivariate forms:* $n = 2$.
- *Quadratic forms:* $m = 2$.
- *Ternary quartics:* $n = 3, m = 4$.

By dehomogenization, we can interpret these results in terms of polynomials (not necessarily homogeneous). The first case corresponds to the equivalence of the nonnegativity and SOS conditions for polynomials in one variable. This is easy to show using a factorization of the polynomial in linear and quadratic factors. The second one is the familiar case of quadratic polynomials, where the sum of squares decomposition follows from an eigenvalue/eigenvector factorization. The somewhat surprising third case corresponds to quartic polynomials in two variables.

The effective computation of the sum of squares decomposition has been analyzed from different viewpoints by several authors. In the author's opinion, there are two main sources, developed independently in unrelated fields. On the one hand, from a convex optimization perspective, the sum of squares decomposition is clearly the underlying machinery in Shor's global bound for polynomial functions (see Example 7.1), as is explicitly mentioned in [Sho87, Sho98]. On the other hand, from an algebraic perspective, it has been presented as the "Gram matrix" method and analyzed extensively by Choi, Lam and Reznick [CLR95], though undoubtedly there are traces of it in the authors' earlier papers.

An implementation of the Gram matrix method is presented in Powers and Wörmann [PW98], though no reference to convexity is made: the resulting SDPs are solved via inefficient, though exact, decision methods. In the control theory literature, related schemes appear in [BL68], and [HH96] (note also the important correction in [Fu98]). Specific connections with SDP, resembling the ones developed here, have also been explored independently by Ferrier [Fer98], Nesterov [Nes00], and Lasserre [Las01].

The basic idea of the method is the following: express the given polynomial $F(x)$ of degree $2d$ as a quadratic form in all the monomials of degree less than or equal to d given by the different products of the x variables. Concretely:

$$F(x) = z^T Q z, \qquad z = [1, x_1, x_2, \ldots x_n, x_1 x_2, \ldots, x_n^d], \tag{3.5}$$

with Q being a constant matrix. The length of the vector z is equal to $\binom{n+d}{d}$. If in the representation above the matrix Q is positive semidefinite, then $F(x)$ is clearly nonnegative. However, since the variables in z are *not algebraically independent*, the matrix Q in (3.5) *is not unique*, and Q may be PSD for some representations but not for others. By simply expanding the right-hand side of (3.5), and matching coefficients of x, it is easily shown that the set of matrices Q that satisfy (3.5) is an *affine subspace*.

If the intersection of this subspace with the positive semidefinite matrix cone is nonempty, then the original function F is guaranteed to be SOS (and therefore nonnegative). This follows from an eigenvalue factorization of $Q = T^T D T, d_i \geq 0$, which produces the sum of squares decomposition $F(x) = \sum_i d_i (Tz)_i^2$. Notice that the number of squares in the representation can always be taken to be equal to the rank of the matrix Q. For the other direction, if F can indeed be written as the sum of squares of polynomials, then expanding in monomials will provide the representation (3.5). By the above arguments, the following is true:

Theorem 3.3. *The existence of a sum of squares decomposition of a polynomial in n variables of degree $2d$ can be decided by solving a semidefinite programming feasibility problem. If the polynomial is dense (no sparsity), the dimensions of the matrix inequality are equal to $\binom{n+d}{d} \times \binom{n+d}{d}$.*

In this specific formulation, the theorem appears in [Par00b, Par00a]. As we have discussed more extensively in those works and the previous paragraphs, the crucial optimization oriented convexity ideas can be traced back to Shor, who explored them at a time before the convenient language of the SDP framework had been fully established.

Notice that the size of the resulting SDP problem is polynomial in both n or d if the other one is fixed. However, it is not jointly polynomial if both the degree and the number of variables grow: $\binom{2n}{n}$ grows exponentially with n (but in this case, the size of the problem description also blows up).

Remark 3.4. If the input polynomial $F(x)$ is *homogeneous* of degree $2d$, then it is sufficient to restrict the components of z to the monomials of degree exactly equal to d.

Example 3.5. Consider the quartic form in two variables described below, and define $z_1 := x^2, z_2 := y^2, z_3 := xy$:

$$\begin{aligned} F(x,y) &= 2x^4 + 2x^3y - x^2y^2 + 5y^4 \\ &= \begin{bmatrix} x^2 \\ y^2 \\ xy \end{bmatrix}^T \begin{bmatrix} q_{11} & q_{12} & q_{13} \\ q_{12} & q_{22} & q_{23} \\ q_{13} & q_{23} & q_{33} \end{bmatrix} \begin{bmatrix} x^2 \\ y^2 \\ xy \end{bmatrix} \\ &= q_{11}x^4 + q_{22}y^4 + (q_{33} + 2q_{12})x^2y^2 + 2q_{13}x^3y + 2q_{23}xy^3 \end{aligned}$$

Therefore, in order to have an identity, the following linear equalities should hold:

$$q_{11} = 2, \quad q_{22} = 5, \quad q_{33} + 2q_{12} = -1, \quad 2q_{13} = 2, \quad 2q_{23} = 0. \tag{3.6}$$

A positive semidefinite Q that satisfies the linear equalities can then be found using semidefinite programming. A particular solution is given by:

$$Q = \begin{bmatrix} 2 & -3 & 1 \\ -3 & 5 & 0 \\ 1 & 0 & 5 \end{bmatrix} = L^T L, \qquad L = \frac{1}{\sqrt{2}} \begin{bmatrix} 2 & -3 & 1 \\ 0 & 1 & 3 \end{bmatrix},$$

and therefore we have the sum of squares decomposition:

$$F(x,y) = \frac{1}{2}(2x^2 - 3y^2 + xy)^2 + \frac{1}{2}(y^2 + 3xy)^2.$$

Example 3.6. The following example is from [Bos82, Example 2.4], where it is required to find whether or not the quartic polynomial,

$$P(x_1, x_2, x_3) = x_1^4 - (2x_2x_3 + 1)x_1^2 + (x_2^2x_3^2 + 2x_2x_3 + 2),$$

is positive definite. In [Bos82], this property is established using decision algebra.

By constructing the Q matrix as described above, and solving the corresponding SDPs, we obtain the sums of squares decomposition:

$$P(x_1, x_2, x_3) = 1 + x_1^2 + (1 - x_1^2 + x_2x_3)^2,$$

that immediately establishes global positivity. Notice that the decomposition actually proves a stronger fact, namely that $P(x_1, x_2, x_3) \geq 1$ for all values of x_i. This lower bound is optimal, since for instance $P(0, 1, -1) = 1$.

There are two crucial properties that distinguish the sum of squares viewpoint from other approaches to the polynomial nonnegativity problem:

- The relative *tractability*, since the question now reduces to efficiently solvable SDPs.
- The fact that the approach can be easily extended to the problem of *finding* a sum of squares polynomial, in a given convex set.

To see this last point, consider the polynomial family $p(x, \lambda)$, where $p(x, \lambda)$ is affine in λ, with the parameter λ belonging to a convex set $\mathcal{C} \subseteq \mathbb{R}^n$ defined by semidefinite constraints. Then, the search over $\lambda \in \mathcal{C}$ for a $p(x, \lambda)$ that is a sum of squares can be posed as a semidefinite program. The argument is exactly as before: writing $P(x, \lambda) = z^T Q z$ and expanding, we obtain linear equations among the entries of Q and λ. Since both Q are λ are defined by semidefinite constraints, the result follows.

This last feature will be the critical one in the application of the techniques to practical problems, and in extending the results to the general semialgebraic case in Section 4.

3.3. *The dual problem*

It is enlightening to analyze the dual problem, that gives conditions on when a polynomial $F(x)$ is *not* a sum of squares. Obviously, one such case is when $F(x)$ takes a negative value for some $x = x_0$. However, because of the distinction between the nonnegativity and SOS conditions, other cases are possible.

By definition, the dual of the sum of squares cone are the linear functionals that take nonnegative values on it. Obviously, these should depend only the coefficients of the polynomial, and not on the specific matrix Q in the representation $F(x) = z^T Q z$. Two possible interpretations of the dual functionals are as differential forms [PS01], or as truncated moment sequences [Las01]. As mentioned in the first paragraph, not all the elements in the dual cone will arise from pointwise function evaluation of F.

Given $F(x)$, consider any representation:

$$F(x) = z^T Q z = \text{trace } z z^T Q,$$

where z is the vector of monomials in (3.3), and Q is not necessarily positive semidefinite. The matrix zz^T has rank one, and due to the algebraic dependencies among the components of z, many of its entries are repeated. Replace now the matrix zz^T by another one W, of the same dimensions, that is positive semidefinite and satisfies the same dependencies among its entries as zz^T does. Then, by construction, the pairing $\langle W, Q\rangle = \operatorname{trace} WQ$ will not depend on the specific choice of Q, as long as it represents the same polynomial F.

Example 3.7. Consider again Example 3.5, where $z_1 = x^2, z_2 = y^2, z_3 = xy$. In this case, the dual variable is:

$$W = \begin{bmatrix} w_{11} & w_{12} & w_{13} \\ w_{12} & w_{22} & w_{23} \\ w_{13} & w_{23} & w_{33} \end{bmatrix}, \quad zz^T = \begin{bmatrix} z_1^2 & z_1 z_2 & z_1 z_3 \\ z_1 z_2 & z_2^2 & z_2 z_3 \\ z_1 z_3 & z_2 z_3 & z_3^2 \end{bmatrix}, \tag{3.7}$$

and the constraint that $z_1 z_2 = z_3^2$ translates into the condition $w_{12} = w_{33}$. We can easily verify that, after substitution with the coefficients of F, the inner product

$$\begin{aligned} \langle W, Q\rangle &= \operatorname{trace} \begin{bmatrix} w_{11} & w_{33} & w_{13} \\ w_{33} & w_{22} & w_{23} \\ w_{13} & w_{23} & w_{33} \end{bmatrix} \begin{bmatrix} q_{11} & q_{12} & q_{13} \\ q_{12} & q_{22} & q_{23} \\ q_{13} & q_{23} & q_{33} \end{bmatrix} \\ &= w_{11}q_{11} + w_{33}(q_{33} + 2q_{12}) + 2q_{13}w_{13} + w_{22}q_{22} + 2w_{23}q_{23} \\ &= 2w_{11} - w_{33} + 2w_{13} + 5w_{22}, \end{aligned}$$

does not depend on the elements of Q, but only on the coefficients themselves.

Now, given *any* Q representing $F(x)$, it is clear that a sufficient condition for F not to be a sum of squares is the existence of a matrix W as above satisfying

$$\operatorname{trace} WQ < 0, \qquad W \succeq 0.$$

The reason is the following: if $F(x)$ was indeed a sum of squares, then there would exist a $Q_{SOS} \succeq 0$ representing F. By construction, the expression above is independent of the choice of Q (as long as it represents the same polynomial), and therefore by replacing Q by the hypothetical Q_{SOS} we immediately reach a contradiction, since in that case the trace term would be nonnegative, as both W and Q_{SOS} are PSD.

The dual problem gives direct insight into the process of checking, after solving the SDPs, whether the relaxation was *exact*, since if there exists an x^* such that $F(x^*) = \sum f_i^2(x^*) = 0$, then it should necessarily be a common root of all the f_i. The simplest instance occurs when the obtained dual matrix W has *rank one*, and the components of the corresponding factorization verify the *constraints* satisfied by the z_i variables, in which case the point x^* can be directly recovered from the entries of W.

4. Real algebra

At its most basic level, algebraic geometry deals with the study of the solution set of a system of polynomial equations. From a more abstract viewpoint, it focuses on the close relationship between geometric objects and the associated algebraic structures. It is a subject with a long and illustrious history, and many links to seemingly unconnected areas of mathematics, such as number theory.

Increasingly important in the last decades, particularly from a computational viewpoint, is the fact that new algorithms and methodologies (for instance, Gröbner basis) have enabled the study of very complicated problems, not amenable to paper and pencil calculations.

In this section, a few basic elements from algebraic geometry are presented. For comparison purposes and clarity of presentation, we present both the complex and real cases, though we will be primarily concerned with the latter. An excellent introductory reference for the former is [CLO97], with [BCR98] being an advanced research treatment of the real case.

The usual name for the specific class of theorems we use is *Stellensätze*, from the German words Stellen (places) and Satz (theorem). The first such result was proved by Hilbert, and deals with the case of an algebraically closed field such as $\mathbb{C}$. Since in many problems we are interested in the real roots, we need to introduce the Artin-Schreier theory of formally real fields, developed along the search for a solution of Hilbert's 17th problem.

4.1. The complex case: Hilbert's Nullstellensatz

Let the ring of polynomials with complex coefficients in n variables be $\mathbb{C}[x_1, \ldots, x_n]$. Recall the definition of a polynomial ideal [CLO97]:

Definition 4.1. *The set $I \subseteq \mathbb{C}[x_1, \ldots, x_n]$ is an* ideal *if it satisfies:*

1. $0 \in I$.
2. *If $a, b \in I$, then $a + b \in I$.*
3. *If $a \in I$ and $b \in \mathbb{C}[x_1, \ldots, x_n]$, then $a \cdot b \in I$.*

Definition 4.2. *Given a finite set of polynomials $(f_i)_{i=1,\ldots,s}$, define the set*

$$\langle f_1, \ldots, f_s \rangle := \left\{ \sum_{i=1}^{s} f_i g_i, \quad g_i \in \mathbb{C}[x_1, \ldots, x_n] \right\}$$

It can be easily shown that the set $\langle f_1, \ldots, f_s \rangle$ is an ideal, known as the ideal *generated* by the f_i.

The result we present next is the Nullstellensatz due to Hilbert. The theorem establishes a correspondence between the set of solutions of polynomials equations (a geometric object known as an *affine variety*), and a polynomial ideal (an algebraic concept). We state below a version appropriate for our purposes:

Theorem 4.3 (Hilbert's Nullstellensatz).
Let $(f_j)_{j=1,\ldots,s}$, be a finite family of polynomials in $\mathbb{C}[x_1, \ldots, x_n]$. Let I be the ideal generated by $(f_j)_{j=1,\ldots,s}$. Then, the following statements are equivalent:

1. *The set*
$$\{x \in \mathbb{C}^n \mid f_i(x) = 0, \quad i = 1, \ldots, s\} \tag{4.1}$$
is empty.
2. *The polynomial 1 belongs to the ideal, i.e., $1 \in I$.*

3. *The ideal is equal to the whole polynomial ring:* $I = \mathbb{C}[x_1, \dots, x_n]$.
4. *There exist polynomials* $g_i \in \mathbb{C}[x_1, \dots, x_n]$ *such that:*

$$f_1(x)g_1(x) + \cdots + f_s(x)g_s(x) = 1. \tag{4.2}$$

The "easy" sufficiency direction ($4 \Rightarrow 1$) should be clear: if the identity (4.2) is satisfied for some polynomials g_i, and assuming there exists a feasible point x_0, after evaluating (4.2) at x_0 we immediately reach the contradiction 0=1. The hard part of the theorem, of course, is proving the existence of the polynomials g_i.

The Nullstellensatz can be directly applied to prove the nonexistence of *complex* solutions for a given system of polynomial equations. The polynomials g_i provide a *certificate* (sometimes called a Nullstellensatz refutation) that the set described by (4.1) is empty. Given the g_i, the identity (4.2) can be efficiently verified. There are at least two possible approaches to effectively find polynomials g_i:

Linear algebra. The first one depends on having explicit bounds on the degree of the products $f_i g_i$. A number of such bounds are available in the literature; see for instance [Bro87, Kol88, BS91]. For example, if the polynomials $f_i(x)$ have maximum degree d, and $x \in \mathbb{C}^n$, then the bound

$$\deg f_i g_i \leq \max(3, d)^n$$

holds. The bound is tight, in the sense that there exist specific examples of systems for which the expression above is an equality. Therefore, given a upper bound on the degree, and a parameterization of the unknown polynomials g_i, a solution can be obtained by solving a system of linear equations. It is also possible to attempt to search directly for low-degree solutions, since the known bounds can also be extremely conservative.

Gröbner basis. An alternative procedure uses Gröbner basis methods [CLO97, Mis93]. By Hilbert's Basis theorem, every polynomial ideal is finitely generated. Gröbner bases provide a computationally convenient representation for a set of generating polynomials of an ideal. As a byproduct of the computation of a Gröbner basis of the ideal I, explicit expressions for the polynomials g_i can be obtained.

Example 4.4. As an example of a Nullstellensatz refutation, we prove that the following system of polynomial inequalities does not have solutions over $\mathbb{C}$.

$$\begin{aligned} f_1(x) &:= x^2 + y^2 - 1 = 0 \\ f_2(x) &:= x + y = 0 \\ f_3(x) &:= 2x^3 + y^3 + 1 = 0. \end{aligned}$$

To show this, consider the polynomials

$$\begin{aligned} g_1(x) &:= \tfrac{1}{7}(1 - 16x - 12y - 8xy - 6y^2) \\ g_2(x) &:= \tfrac{1}{7}(-7y - x + 4y^2 - 16 + 12xy + 2y^3 + 6y^2 x) \\ g_3(x) &:= \tfrac{1}{7}(8 + 4y). \end{aligned}$$

After simple algebraic manipulations, we verify that

$$f_1 g_1 + f_2 g_2 + f_3 g_3 = 1,$$

proving the nonexistence of solutions over $\mathbb{C}$.

4.2. The real case: Positivstellensatz

The conditions in the Nullstellensatz are necessary and sufficient only in the case when the field is algebraically closed (as in the case of $\mathbb{C}$). When this requirement does not hold, only the sufficiency argument is still valid. A simple example is the following: over the reals, the equation

$$x^2 + 1 = 0$$

does not have a solution (i.e., the associated real variety is empty). However, the corresponding polynomial ideal does not include the element 1.

When we are primarily interested in real solutions, the lack of algebraic closure of $\mathbb{R}$ forces a different approach, and the theory should be modified accordingly. This led to the development of the Artin-Schreier theory of *formally real* fields; see [BCR98, Raj93] and the references therein.

The starting point is one of the intrinsic properties of $\mathbb{R}$:

$$\sum_{i=1}^{n} x_i^2 = 0 \Longrightarrow x_1 = \ldots = x_n = 0. \tag{4.3}$$

A field is *formally real* if it satisfies the above condition. The theory of formally real fields has very strong connections with the sums of squares that we have seen at the beginning of Section 3.2. For example, an alternative (but equivalent) statement of (4.3) is that a field is formally real if and only if the element -1 is not a sum of squares.

In many senses, real algebraic geometry still lacks the full maturity of its counterpart, the algebraically closed case (such as $\mathbb{C}$). Fortunately, many important results are available: crucial to our developments will be the Real Nullstellensatz, also known as Positivstellensatz [Ste74, BCR98].

Before proceeding further, we need to introduce a few concepts. Given a set of polynomials $p_i \in \mathbb{R}[x_1, \ldots, x_n]$, let $M(p_i)$ be the *multiplicative monoid* generated by the p_i, i.e., the set of finite products of the elements p_i (including the empty product, the identity). The following definition introduces the ring-theoretic concept of *cone*.

Definition 4.5. *A* cone P *of* $\mathbb{R}[x_1, \ldots, x_n]$ *is a subset of* $\mathbb{R}[x_1, \ldots, x_n]$ *satisfying the following properties:*

1. $a, b \in P \Rightarrow a + b \in P$
2. $a, b \in P \Rightarrow a \cdot b \in P$
3. $a \in \mathbb{R}[x_1, \ldots, x_n] \Rightarrow a^2 \in P$

Given a set $S \subseteq \mathbb{R}[x_1, \ldots, x_n]$, let $P(S)$ be the smallest cone of $\mathbb{R}[x_1, \ldots, x_n]$ that contains S. It is easy to see that $P(\emptyset)$ corresponds to the polynomials that can be expressed as a sum of squares, and is the smallest cone in $\mathbb{R}[x_1, \ldots, x_n]$. For a finite set $S = \{a_1, \ldots, a_m\} \subseteq \mathbb{R}[x_1, \ldots, x_n]$, its associated cone can be expressed as:

$$P(S) = \{p + \sum_{i=1}^{r} q_i b_i \mid p, q_1, \ldots, q_r \in P(\emptyset),\ b_1, \ldots, b_r \in M(a_i)\}.$$

The Positivstellensatz, due to Stengle [Ste74], is a central theorem in *real* algebraic geometry. It is a common generalization of linear programming duality (for linear inequalities) and Hilbert's Nullstellensatz (for an algebraically closed field). It states that, for a system of polynomial equations and inequalities, either there exists a solution in $\mathbb{R}^n$, or there exists a certain polynomial identity which bears *witness* to the fact that no solution exists. For concreteness it is stated here for $\mathbb{R}$, instead of the general case of real closed fields.

Theorem 4.6 ([BCR98, Theorem 4.4.2]). *Let $(f_j)_{j=1,\ldots,s}$, $(g_k)_{k=1,\ldots,t}$, $(h_\ell)_{\ell=1,\ldots,u}$ be finite families of polynomials in $\mathbb{R}[x_1,\ldots,x_n]$. Denote by P the cone generated by $(f_j)_{j=1,\ldots,s}$, M the multiplicative monoid generated by $(g_k)_{k=1,\ldots,t}$, and I the ideal generated by $(h_\ell)_{\ell=1,\ldots,u}$. Then, the following properties are equivalent:*

1. *The set*
$$\left\{ x \in \mathbb{R}^n \;\middle|\; \begin{array}{ll} f_j(x) \geq 0, & j = 1,\ldots,s \\ g_k(x) \neq 0, & k = 1,\ldots,t \\ h_\ell(x) = 0, & j = 1,\ldots,u \end{array} \right\} \tag{4.4}$$
is empty.
2. *There exist $f \in P, g \in M, h \in I$ such that $f + g^2 + h = 0$.*

Proof. We show only the sufficiency part, i.e., $2 \Rightarrow 1$. We refer the reader to [BCR98] for the other direction.

Assume that the set is not empty, and consider any element x_0 from it. In this case, it follows from the definitions that:

$$f(x_0) \geq 0, \qquad g^2(x_0) > 0, \qquad h(x_0) = 0$$

This implies that $f(x_0) + g^2(x_0) + h(x_0) > 0$, in contradiction with the assumption that $f + g^2 + h = 0$. ☐

The Positivstellensatz guarantees the existence of *infeasibility certificates* or *refutations*, given by the polynomials f, g and h. For complexity reasons these certificates cannot be polynomial time checkable for every possible instance, unless NP=co-NP. While effective bounds on the degrees do exist, their expressions are at least triply exponential.

Example 4.7. To illustrate the differences between the real and the complex case, and the use of the Positivstellensatz, consider the very simple case of the standard quadratic equation

$$x^2 + ax + b = 0.$$

By the fundamental theorem of algebra (or in this case, just the explicit formula for the solutions), the equation always has solutions on $\mathbb{C}$. For the case when $x \in \mathbb{R}$, the solution set will be empty if and only if the discriminant D satisfies

$$D := b - \frac{a^2}{4} > 0.$$

In this case, taking

$$\begin{aligned} f &:= [\tfrac{1}{\sqrt{D}}(x+\tfrac{a}{2})]^2 \\ g &:= 1 \\ h &:= -\tfrac{1}{D}(x^2+ax+b), \end{aligned}$$

the identity $f + g^2 + h = 0$ is satisfied.

5. Finding refutations using SDP

Theorem 4.6 provides the basis for a hierarchy of sufficient conditions to verify that a given semialgebraic set is empty. Notice that it is possible to affinely parameterize a family of candidate f and h, since from Section 3.2, the sum of squares condition can be expressed as an SDP. Restricting the degree of the possible multipliers, we obtain semidefinite programs, that can be efficiently solved.

Our main result provides therefore a constructive approach to solutions of the Positivstellensatz equations:

Theorem 5.1 ([Par00b]). *Consider a system of polynomial equalities and inequalities of the form (4.4). Then, the search for bounded degree Positivstellensatz refutations can be done using semidefinite programming. If the degree bound is chosen to be large enough, then the SDPs will be feasible, and the certificates obtained from its solution.*

It is convenient to contrast this result with the Nullstellensatz analogue, where the search for bounded-degree certificates could be done using just linear algebra.

Proof. Given a degree d, choose g in the following way: if $t = 0$, i.e., the set of inequations is empty, then $g = 1$. Otherwise, let $g = \prod_{i=1}^{t} g_i^{2m}$, choosing m such that the degree of g is greater than or equal to d. For the cone of inequalities, choose a degree $d_2 \geq d$, $d_2 \geq \deg(g)$. Write

$$f = p_0 + p_1 f_1 + \cdots + p_s f_s + p_{12} f_1 f_2 + \cdots + p_{12\ldots s} f_1 \cdots f_s$$

and give a parameterization of the polynomials p_i of degree less than or equal to d_2. Similarly, for the polynomial h in the ideal of equations, write

$$h = q_1 h_1 + \cdots + q_u h_u,$$

parameterizing the polynomials q_i of degree less than or equal to d_2.

Consider now the SDP feasibility problem:

$$p_i \ \text{ are sums of squares,}$$

with the equality constraints implied by the equation $f + g^2 + h = 0$, the decision variables being the coefficients of the p_i, q_i.

If the set defined by (4.4) is empty, then by the Positivstellensatz, polynomial certificates $f_\star, g_\star, h_\star$ do exist. By construction of the SDP problem above, there exists a finite number d_0, such that for every $d \geq d_0$ the semidefinite program is feasible, since there exists at least one feasible point, namely $f_\star, g_\star, h_\star$. Therefore, a set of infeasibility certificates of the polynomial system can directly be obtained from a feasible point of the SDP. □

Remark 5.2. The procedure as just described contains some considerable overparametrization of the polynomials, due to the generality of the statement and the need to deal with special cases. In general, once the problem structure is known (for instance, no inequations), much more compact formulations can be given, as in the case of quadratic programming presented in Section 7.4.

The presented formulation deals only with the case of proving that semialgebraic sets are empty. Nevertheless, it can be easily applied to more general problems, such as checking nonnegativity of a polynomial over a semialgebraic set. We describe two simple cases, more being presented in Section 7.

Example 5.3. Consider the problem of verifying if the implication

$$a(x) = 0 \Rightarrow b(x) \geq 0 \tag{5.1}$$

holds. The implication is true if and only if the set

$$\{x \mid -b(x) \geq 0,\ b(x) \neq 0,\ a(x) = 0\}$$

is empty. By the Positivstellensatz, this holds iff there exist polynomials s_1, s_2, t and an integer k such that:

$$s_1 - s_2 b + b^{2k} + ta = 0,$$

and s_1 and s_2 are sums of squares. A special case, easy to verify, is obtained by taking $s_1(x) = 0$, $k = 1$, and $t(x) = b(x)r(x)$, in which case the expression above reduces to the condition:

$$b(x) + r(x)a(x) \qquad \text{is a sum of squares,} \tag{5.2}$$

which clearly implies that (5.1) holds. Since this expression is affine in $r(x)$, the search for such an $r(x)$ can be posed as a semidefinite program.

Example 5.4. Let $f(x)$ be a polynomial function, to be minimized over a semialgebraic set S. Then, γ is a lower bound of $\inf_{x \in S} f(x)$ if and only if the semialgebraic set $\{x \in S, f(x) < \gamma\}$ is empty. For fixed γ, we can search for certificates using SDP. It is also possible, at the expense of fixing some of the variables, to search for the best possible γ for the given degree.

In the case of basic *compact* semialgebraic sets, i.e., compact sets of the form $K = \{x \in \mathbb{R}^n, f_1(x) \geq 0, \ldots, f_s(x) \geq 0\}$, a stronger version of the Positivstellensatz, due to Schmüdgen [Sch91] can be applied. It says that a polynomial $f(x)$ that is strictly positive on K, actually belongs to the cone generated by the f_i. The Positivstellensatz presented in Theorem 4.6 only guarantees in this case the existence of g, h in the cone such that $fg = 1 + h$. An important computational drawback of the Schmüdgen formulation is that, due to the cancellations that must occur, the degrees of the infeasibility certificates can be significantly larger than in the standard Positivstellensatz [Ste96].

5.1. Interpretation and related work

The main idea of Positivstellensatz refutations can be easily summarized. If the constraints $h_i(x_0) = 0$ are satisfied, we can then generate by multiplication and addition a whole class of expressions, namely those in the corresponding ideal, that should also vanish at x_0. For the inequation case ($g_i \neq 0$), multiplication of the constraints g_i provides new functions that are guaranteed not to have a zero at x_0. For the constraints $f_i \geq 0$, new *valid inequalities*, nonnegative at x_0, are derived by multiplication with other constraints and nonnegative functions (actually, sums of squares). By simultaneously searching over all these possibilities, and combining the results, we can obtain a proof of the infeasibility of the original system. These operations are simultaneously carried over by the optimization procedure.

It is interesting to compare this approach with the standard duality bounds in convex programming. In that case, *linear* combinations of constraints (basically, linear functionals) are used to derive important information about the feasible set. The Positivstellensatz formulation instead achieves improved results by combining the constraints in an arbitrary *nonlinear* fashion, by allowing multiplication of constraints and products with nonnegative functions.

There are many interesting links with foundational questions in logic and theoretical computer science. The Positivstellensatz constitutes a complete algebraic proof system (see [GV02] and the references therein), so issues about proof length are very relevant. For many practical problems, very concise (low degree) infeasibility certificates can be constructed, even though in principle there seems to be no reason to expect so. This is an issue that clearly deserves much more research.

Related ideas have been explored earlier in “lift-and-project” techniques used to derive valid inequalities in zero-one combinatorial optimization problems, such as those introduced by Lovász-Schrijver [LS91,Lov94] and Sherali-Adams [SA90]. In particular, the latter authors develop the so-called Reformulation-Linearization technique (RLT), later extended by Tuncbilek to handle general polynomial problems [SA99, Chapter 9]. In this procedure, products of explicit upper and lower bounds on the variables are formed, which are later linearized by the introduction on new variables, resulting in a relaxation that can be formulated as a linear program. Both approaches can be used to develop tractable approximations to the convex hull of zero-one points in a given convex set. A typical application is the case of integer linear or polynomial programs, which are known NP-hard problems. An important property of these approaches is the exactness of the procedure after an *a priori* fixed number of liftings.

Some common elements among these procedures are the use of new variables and constraints, defined as products of the original ones, and associated linear (in RLT) or semidefinite constraints (in the Lovász-Schrijver N_+ relaxation). In the author’s opinion, an important asset of the approach introduced in the current paper as opposed to earlier work is that it focuses on the algebraic-geometric structure of the solution set itself, rather than on that of the describing equations. Additionally, and similar to the results mentioned above, it can be shown using simple algebraic properties that *a priori* bounded finite termination always holds for the case of zero-dimensional ideals [Par02a].

The recent independent work of Lasserre [Las01] has several common elements with the one presented here, while focusing more closely on polynomial optimization and the dual viewpoint of truncated moment sequences. Using additional compactness assumptions, and invoking results on linear representations of strictly positive polynomials over semialgebraic sets, a sequence of SDPs for approximating global minima is proposed. In the framework of the present paper, it can be shown that Lasserre's approach corresponds to a very specific kind of affine Positivstellensatz certificates. This restriction, while somewhat attractive from a theoretical viewpoint, in some cases can produce significantly weaker computational bounds than using the full power of the Positivstellensatz [Ste96], an issue explored in more detail elsewhere [Par02b]. An advantage of the techniques presented here is that exact convergence is *always* achieved in a finite number of steps (regardless of compactness, or any other assumption), while the results in [Las01] can only produce a monotone sequence of bounds, converging to the optimal. It should be mentioned that Laurent [Lau] has recently analyzed in detail the relationship of the Lasserre approach in the 0-1 case with the earlier relaxation schemes mentioned above.

The bottom line is that *all* the procedures mentioned above can be understood as particular cases of Positivstellensatz refutations, where either restrictions in the verification (linear programming in RLT) or structural constraints on the certificates themselves (linearity, in Lasserre's approach) are imposed *a priori*. In a very concrete sense, reinforced by the connections with proof systems alluded to above, the Positivstellensatz represents the most general deductive system for which inferences from the given equations can be made, and for which proofs can be efficiently search over and found.

6. Computational considerations

6.1. Implementation

In this section, we briefly discuss some aspects of the computational implementation of the sum of squares decision procedure. As we have seen in Section 3, for semidefinite programs, just like in the linear programming case, there are two formulations: primal and dual. In principle, it is possible to pose the sum of squares problem as either of them, with the end results being mathematically equivalent. However, for reasons to be described next, one formulation may be numerically more efficient than the other, depending on the dimension of the problem.

As mentioned in Section 3, a semidefinite program can be interpreted as an optimization problem over the intersection of an affine subspace $\mathcal{L}$ and the cone $\mathcal{S}^n_+$. Depending on the dimension of $\mathcal{L}$, it may be computationally advantageous to describe the subspace using either a set of generators (an explicit, or *image* representation) or the defining linear equations (the implicit, or *kernel* representation).

When the dimension of $\mathcal{L}$ is small relative to the ambient space, then an efficient representation will be given by a set of generators (or a basis), i.e.,

$$X = G_0 + \sum_i \lambda_i G_i.$$

$2d \setminus n$	1	3	5	7	9	11	13	15
2	2	4	6	8	10	12	14	16
4	3	10	21	36	55	78	105	136
6	4	20	56	120	220	364	560	816
8	5	35	126	330	715	1365	2380	3876
10	6	56	252	792	2002	4368	8568	15504
12	7	84	462	1716	5005	12376	27132	54264

Table 6.1. Dimension of the matrix Q as a function of the number of variables n and the degree $2d$. The corresponding expression is $\binom{n+d}{d}$.

On the other hand, if $\mathcal{L}$ is nearly full dimensional, then a more concise description is given instead by the set of linear relations satisfied by the elements of $\mathcal{L}$, that is,

$$\langle X, A_i \rangle = b_i.$$

While the resulting problems are formally the same, there are usually significant differences in the associated computation times.

Consider the problem of checking if a dense polynomial of total degree $2d$ in n variables is a sum of squares, using the techniques described earlier. The number of coefficients is, as we have seen, equal to $\binom{n+2d}{2d}$. The dimension of the corresponding matrix Q is $\binom{n+d}{d}$ (see Table 6.1).

If we use an explicit representation the total number of additional variables we need to introduce can be easily be shown to be:

$$N_1 = \frac{1}{2}\left[\binom{n+d}{d}^2 + \binom{n+d}{d}\right] - \binom{n+2d}{2d}.$$

On the other hand, in the implicit formulation the number of equality constraints (i.e., the number of matrices A_i in (3.1)) is exactly equal to the number of coefficients, i.e.

$$N_2 = \binom{n+2d}{2d}.$$

Example 6.1. We revisit Example 3.5, where an implicit (or kernel) representation of the one dimensional subspace of matrices Q was given. An explicit (image) representation of the same subspace is given by:

$$Q = \begin{bmatrix} 2 & -\lambda & 1 \\ -\lambda & 5 & 0 \\ 1 & 0 & 2\lambda - 1 \end{bmatrix}.$$

The particular matrix in Example 3.5 corresponds to the choice $\lambda = 3$. Notice that the free variable λ corresponds to the algebraic dependency among the entries of z: $(x^2)(y^2) = (xy)^2$.

For fixed d, both quantities N_1, N_2 are $O(n^{2d})$; however, the corresponding constants can be vastly different. In fact, the following expressions hold:

$$N_1 \approx \left(\frac{1}{2(d!)^2} - \frac{1}{(2d)!}\right) n^{2d}, \qquad N_2 \approx \left(\frac{1}{(2d)!}\right) n^{2d}.$$

For large values of d, the second expression is much smaller than the first one, making the implicit formulation preferable. For small values of n and d, however, the situation is not clear-cut, and the explicit one can be a better choice.

We consider next three representative examples:

1. The case of a quartic univariate polynomial ($n = 1, 2d = 4$). Notice that this is equivalent, by dehomogenization, to the quartic bivariate form in Examples 3.5 and 6.1. The resulting matrix Q has dimensions 3×3, and the number of variables for the explicit and implicit formulation are $N_1 = 1$ and $N_2 = 5$, respectively.
2. A trivariate polynomial of degree 10 ($n = 3, 2d = 10$). The corresponding matrix has dimensions 56×56, and the number of variables is $N_1 = 1310$ and $N_2 = 286$. The advantages of the second approach are clear.
3. A quartic polynomial in 15 variables ($n = 15, 2d = 4$). The corresponding matrix has dimensions 136×136, and the number of variables is $N_1 = 5440$ and $N_2 = 3876$.

A minor inconvenience of the implicit formulation appears when the optimization problem includes additional variables, for which no a priori bounds are known. Most current SDP implementations do not easily allow for an efficient mixed primal-dual formulation, where some variables are constrained to be in the PSD cone and others are free. This is a well-known issue already solved in the linear programming setting, where current software allows for the efficient simultaneous handling of both nonnegative and unconstrained variables.

6.2. Exploiting structure

If the polynomials are sparse, in the sense that only a few of the monomials are nonzero, then it is usually possible to considerably simplify the resulting SDPs. To do this, we can use a result by Reznick, first formulated in [Rez78], that characterizes the monomials that can appear in a sum of squares representation, in terms of the Newton polytope of the input polynomial.

Another property that can be fully exploited for algorithmic efficiency is the presence of symmetries. If the problem data is invariant under the action of a symmetry group, then the computational burden of solving the optimization problem can be substantially reduced. This aspect has strong connections with representation and invariant theories, and is analyzed in much more detail in [GP01].

In practice, the actual performance will be affected by other elements in addition to the number of variables in the chosen formulation. In particular, the extent to which the specific problem-dependent structure can be exploited is usually the determining factor in the application of optimization methods to medium or large-scale problems.

7. Applications

In this section we outline some specific application areas to which the developed techniques have shown a great potential, when compared to traditional tools. The descriptions are necessarily brief, with more detailed treatments appearing elsewhere. Computational implementations for most of these examples are provided with the freely available SeDuMi-based SOSTOOLS software package [PPP02], that fully implements the approach developed in this paper. Needless to say, the generality of the semialgebraic problem formulation makes possible the use of the presented methods in numerous other areas.

7.1. Global bounds for polynomial functions

It is possible to apply the technique to compute global lower bounds for polynomial functions [Sho87, Sho98, Las01]. For an in-depth analysis of this particular problem, including numerous examples and a comparison with traditional algebraic techniques, we refer the reader to [PS01].

The condition

$$F(x) - \gamma \quad \text{is a sum of squares}$$

is affine in γ, and therefore it is possible to efficiently compute the maximum value of γ for which this property holds. For every feasible γ, $F(x) \geq \gamma$ for all x, so γ is a lower bound on the global minimum. In many cases, as in the Example below, the resulting bound is optimal, i.e., equal to the global minimum, and a point $x_\star$ achieving the global minimum can be recovered from a factorization of the dual solution.

Example 7.1. Consider the function

$$F(x, y) = 4x^2 - \frac{21}{10}x^4 + \frac{1}{3}x^6 + xy - 4y^2 + 4y^4,$$

cited in [Mun99, p. 333] as a test example for global minimization algorithms, since it has several local extrema. Using the techniques described earlier, it is possible to find the largest γ such that $F(x) - \gamma$ is a sum of squares.

Doing so, we find $\gamma_* \approx -1.03162845$. This turns out to be the exact global minimum, since that value is achieved for $x \approx 0.089842, y \approx -0.7126564$.

However, for the reasons mentioned earlier in Section 3.2, it is possible to obtain a lower bound that is strictly less than the global minimum, or even no useful bound at all.

Example 7.2. As examples of a problem with nonzero gaps, we compute global lower bounds of dehomogenizations of the Motzkin polynomial $M(x, y, z)$ presented in (3.4). Since $M(x, y, z)$ is nonnegative, its dehomogenizations also have the same property. Furthermore, since $M(1, 1, 1) = 0$, they always achieve its minimum possible value.

Fixing the variable y, we obtain

$$F(x, z) := M(x, 1, z) = x^4 + x^2 + z^6 - 3x^2z^2.$$

To obtain a lower bound, we search for the maximum γ for which $F(x, z) - \gamma$ is a sum of squares.

Solving the corresponding SDPs, the best lower bound that can be obtained this way can be shown to be $-\frac{729}{4096} \approx -0.177978$, and follows from the decomposition:

$$F(x,z) + \tfrac{729}{4096} = (-\tfrac{9}{8}z + z^3)^2 + (\tfrac{27}{64} + x^2 - \tfrac{3}{2}z^2)^2 + \tfrac{5}{32}x^2$$

The gap can also be infinite, for some particular problems. Consider the dehomogenization in z:

$$G(x,y) := M(x,y,1) = x^4y^2 + x^2y^4 + 1 - 3x^2y^2.$$

While $G(x,y) \geq 0$, it can be shown that $G(x,y) - \gamma$ is not a sum of squares for *any* value of γ, and therefore no useful information can be obtained in this case. This can be fixed (using the Positivstellensatz, or the approach in Example 7.3 below) at the expense of more computation.

As we have seen, the method can sometimes produce suboptimal bounds. This is to be expected, for computational complexity reasons and because the class of nonnegative polynomials is not equal to the SOS ones. It is not clear yet how important this is in practical applications: for example, for the class of random instances analyzed in [PS01], *no example* was produced on which the obtained bound does not coincide with the optimal value. In other words, even though bad examples do indeed exist, they seem to be "rare," at least for some particular ensembles.

In any case, there exist possible workarounds, at a higher computational cost. For a nonnegative $F(x)$, Artin's positive answer to Hilbert's 17th problem assures the existence of a polynomial $G(x)$, such that $F(x)G^2(x)$ can be written as a sum of squares. In particular, Reznick's results [Rez95] show that if F is *positive definite* it is always possible to take $G(x) = (\sum x_i^2)^r$, for sufficiently large r.

Example 7.3. Consider the case of the Motzkin form given in equation (3.4). As mentioned earlier, it cannot be written as a sum of squares of polynomials. Even though it is only semidefinite (so in principle we cannot apply Reznick's theorem), after solving the SDPs we obtain the decomposition:

$$\begin{aligned}(x^2+y^2+z^2)\,M(x,y,z) &= (x^2yz - yz^3)^2 + (xy^2z - xz^3)^2 + (x^2y^2 - z^4)^2 + \\ &\quad + \frac{1}{4}(xy^3 - x^3y)^2 + \frac{3}{4}(xy^3 + x^3y - 2xyz^2)^2,\end{aligned}$$

from where nonnegativity is obvious. Since the polynomials in Example 7.2 are dehomogenizations of $M(x,y,z)$, it follows that this method yields exact solutions for those examples.

To give a rough idea of the large scale problems to which we have applied the techniques in [PS01], we mention that the SOS lower bound for a dense quartic polynomial in thirteen variables (i.e., with 2380 monomials) can be solved on a standard desktop machine, using off-the-shelf software, in approximately 30 minutes.

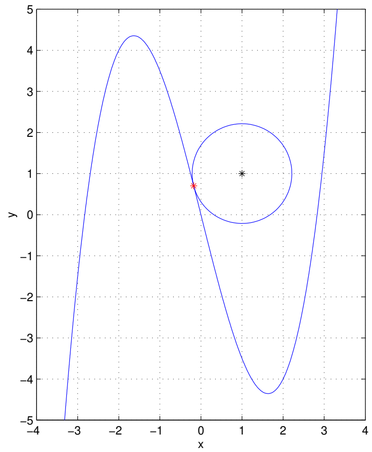

Fig. 7.1. The curve $C(x,y) = 0$ and the minimum distance circle.

7.2. Geometric problems

Many problems in computational geometry can be fully described using a semialgebraic formulation. Properties such as intersection of geometric objects easily reduce to the real feasibility of sets of polynomial equations. In the following very simple example, we use the Positivstellensatz to compute a lower bound on the distance between a point and an algebraic curve.

Example 7.4. In this problem, we compute a lower bound on the distance between a given point (x_0, y_0) and an algebraic curve $C(x,y) = 0$. Take $(x_0, y_0) = (1, 1)$, and let the algebraic curve be

$$C(x,y) := x^3 - 8x - 2y = 0.$$

In this case, we can formulate the optimization problem

$$\min_{C(x,y)=0} (x - x_0)^2 + (y - y_0)^2 \tag{7.1}$$

A lower bound on the optimal value can be obtained as described earlier. Restricting the degree of the auxiliary polynomials to a simple linear expression in x, we can compute the maximum value of γ^2 that satisfies

$$(x-1)^2 + (y-1)^2 - \gamma^2 + (\alpha + \beta x)(x^3 - 8x - 2y) \quad \text{is a sum of squares.} \tag{7.2}$$

It should be clear that if condition (7.2) holds, then every pair of points (x, y) in the curve are at a distance at least equal to γ from (x_0, y_0). To see this, note that if the point (x, y) is in the curve $C(x, y) = 0$, then the last term in (7.2) vanishes, and therefore $(x - 1)^2 + (y - 1)^2 \geq \gamma$. The expression is affine in α, β, and γ^2, and so the problem can be directly solved using SDP.

The optimal solution of the SDPs is:

$$\alpha \approx -0.28466411, \quad \beta \approx 0.07305057, \quad \gamma \approx 1.47221165.$$

The obtained bound γ is sharp, since it is achieved by the values

$$x \approx -0.176299246, \quad y \approx 0.702457168.$$

In Figure 7.1 a plot of $C(x)$ and the optimal solution is presented.

Notice that the original optimization formulation (7.1) is *not* a convex program, and has other local extrema. Nevertheless, the procedure always computes a bound, and in this case we actually recover the global minimum.

7.3. The discriminant of symmetric matrices

The following example illustrates the sum of squares techniques, and deals with the discriminant of symmetric matrices. It has been previously analyzed in [Ily92,Lax98]. Given a symmetric matrix $A \in \mathcal{S}^n$, define its characteristic polynomial $p(\lambda)$ as:

$$p(\lambda) := \det(\lambda I - A).$$

This is a polynomial in λ, of degree n. Its discriminant D (see for instance [Mis93]) is a homogeneous polynomial of degree $n(n-1)$ in the $\binom{n+1}{2}$ coefficients of A. Since A is symmetric, its eigenvalues (the roots of p) are real, and therefore the discriminant D takes only nonnegative values, i.e., $D \geq 0$. The results in [Ily92,Lax98] show that additionally the polynomial p is always a sum of squares. For instance, when $n = 2$, we have:

$$A = \begin{bmatrix} a & b \\ b & c \end{bmatrix}, \quad p(\lambda) = \lambda^2 + (-a-c)\lambda + ac - b^2, \quad D = 4b^2 + a^2 + c^2 - 2ac,$$

and the SOS property holds since D can be alternatively expressed as

$$D = (a-c)^2 + (2b)^2.$$

An explicit expression for the discriminant as a sum of squares is presented in [Ily92]. An interesting unsolved problem is finding a representation with the minimum possible number of squares. For the case $n = 3$, i.e.,

$$M = \begin{bmatrix} a & b & d \\ b & c & e \\ d & e & f \end{bmatrix},$$

after solving the SDPs, using as objective function the trace of the matrix Q as a heuristic for the rank, we obtain the following decomposition using *seven* squares:

$$\begin{aligned}
D &= f_1^2 + f_2^2 + f_3^2 + f_4^2 + 15(f_5^2 + f_6^2 + f_7^2) \\
f_1 &= e^2 f + b^2 c + d^2 a - cf^2 - ac^2 - fa^2 - ce^2 - ab^2 - fd^2 + c^2 f + a^2 c + f^2 a \\
f_2 &= 2d^3 - de^2 - b^2 d - 2dc^2 + 2dcf - bef + 2bce - 2adf - abe + 2acd
\end{aligned}$$

$$
\begin{aligned}
f_3 &= 2e^3 - eb^2 - d^2e - 2ea^2 + 2eac - dbc + 2dab - 2fec - fdb + 2fae \\
f_4 &= 2b^3 - bd^2 - e^2b - 2bf^2 + 2bfa - eda + 2efd - 2cba - ced + 2cfb \\
f_5 &= be^2 - dce - bd^2 + ade \\
f_6 &= db^2 - eab - de^2 + feb \\
f_7 &= ed^2 - bfd - eb^2 + cbd.
\end{aligned}
$$

For the case $n = 3$, the expressions in [Ily92] produce a decomposition with ten distinct square terms.

7.4. Quadratic programming

In this section we specialize the results presented earlier to the common case of quadratic inequalities. Concretely, given m symmetric matrices $A_1, \ldots, A_m \in \mathcal{S}^n$, define the set $\mathcal{A}$ as:

$$\mathcal{A} := \left\{ x \in \mathbb{R}^n \middle| \ \mathbf{x}^T A_i \mathbf{x} \geq 0, \quad \|\mathbf{x}\| = 1 \right\} \tag{7.3}$$

A well-known sufficient condition for the set $\mathcal{A}$ to be empty is given by the existence of scalars λ_i that satisfy the condition:

$$\sum_{i=1}^{m} \lambda_i A_i \preceq -I, \qquad \lambda_i \geq 0. \tag{7.4}$$

The reasoning is very simple: assume $\mathcal{A}$ is not empty, and multiply (7.4) left and right by any $\mathbf{x} \in \mathcal{A}$. In this case, the left-hand side of (7.4) is nonnegative, since all terms are nonnegative, but the right-hand side is -1. This is a contradiction, so $\mathcal{A}$ is empty.

The condition (7.4) is the basis of many results in semidefinite relaxations for quadratic programming problems, such as the one underlying the Goemans-Williamson MAX-CUT algorithm [GW95], and many others. For instance, for MAX-CUT both the objective and the boolean constraints on the decision variables can be modeled by quadratic expressions, namely $\frac{1}{2}\sum_{ij} w_{ij}(1 - x_i x_j)$ and $x_i^2 - 1 = 0$, respectively. Applying the condition (7.4) above to the homogenized system, we obtain a semidefinite program exactly equivalent to the standard SDP MAX-CUT relaxation. It is well-known (and obvious from a complexity standpoint) that this condition can be conservative, in the sense that only bounds on the optimal value are obtained in the worst case.

In the framework of this paper, a good interpretation of condition (7.4) is as a Positivstellensatz refutation, with the multipliers restricted to be a constant. By removing the degree restrictions, more powerful tests can be devised. In the following theorem [Par00b], the case of quadratic multipliers is stated. The generalizations to higher degrees are straightforward, following directly from Theorem 5.1.

Theorem 7.5. *Assume there exist solutions $Q_i \in \mathcal{S}^n, r_{ij} \in \mathbb{R}$ to:*

$$\sum_{i=1}^{m} Q_i(x) A_i(x) + \sum_{1 \leq i < j \leq m} r_{ij} A_i(x) A_j(x) < 0, \qquad \forall x \in \mathbb{R}^n / \{0\}. \tag{7.5}$$

where $Q_i(x) := x^T Q_i x, Q_i \succeq 0$ and $r_{ij} \geq 0$. Then, the set $\mathcal{A}$ is empty.

Proof. It basically follows from the same arguments as in the Positivstellensatz case: the existence of a nontrivial x implies a contradiction. □

Note that the left-hand size of (7.5) is a homogeneous form of degree four. Checking the full condition as written would be again a hard problem, so we check instead a sufficient condition: that the left-hand side of (7.5) can be written (except for the sign change) as a sum of squares. As we have seen in Section 3.2, this can be checked using semidefinite programming methods.

The new relaxation is always at least as powerful as the standard one: this can be easily verified, just by taking $Q_i = \lambda_i I$ and $r_{ij} = 0$. Then, if (7.4) is feasible, then the left hand side of (7.5) is obviously a sum of squares (recall that positive definite quadratic forms are always sums of squares).

In [Par00b], we have applied the new procedure suggested by Theorem 7.5 to a few instances of the MAX-CUT problem where the standard relaxation is known to have gaps, such as the n-cycle and the Petersen graph. For these instances, the new relaxations are exact, i.e., they produce the optimal solution.

7.5. Matrix copositivity

A symmetric matrix $M \in \mathbb{R}^{n\times n}$ is said to be *copositive* if the associated quadratic form takes only nonnegative values on the nonnegative orthant, i.e., if $x_i \geq 0 \Rightarrow x^T M x \geq 0$. As opposed to positive definiteness, which can be efficiently verified, checking if a given matrix is not copositive is an NP-complete problem [MK87].

There exist in the literature explicit necessary and sufficient conditions for a given matrix to be copositive. These conditions are usually expressed in terms of principal minors (see [Väl86,CPS92] and the references therein). However, the complexity results mentioned above imply that in the worst case these tests can take an exponential number of operations (unless P = NP). Thus, the need for efficient sufficient conditions to guarantee copositivity.

Example 7.6. We briefly describe an application of copositive matrices [QDRT98]. Consider the problem of obtaining a lower bound on the optimal solution of a linearly constrained quadratic optimization problem:

$$f^* = \min_{Ax\geq 0,\ x^T x=1} \ x^T Q x$$

If there exists a solution C to the SDP:

$$Q - A^T C A \succeq \gamma I$$

where C is a copositive matrix, then it immediately follows that $f^* \geq \gamma$. Thus, having semidefinite programming tests for copositivity allows for enhanced bounds for this type of problems.

The main difficulty in obtaining conditions for copositivity is dealing with the constraints in the variables, since each x_i has to be nonnegative. While we could apply the general Positivstellensatz construction to this problem, we opt here for a more natural, though equivalent, approach. To check copositivity of M, we can consider $x_i = z_i^2$ and study the global nonnegativity of the fourth order form given by:

$$P(\mathbf{z}) := \mathbf{z}^T M \mathbf{z} = \sum_{i,j} m_{ij} z_i^2 z_j^2$$

where $\mathbf{z} = [z_1^2, z_2^2, \ldots, z_n^2]^T$. It is easy to verify that M is copositive if and only if the form $P(\mathbf{z})$ is positive semidefinite. Therefore, sufficient conditions for $P(\mathbf{z})$ to be nonnegative will translate into sufficient conditions for M being copositive.

If we use the sum of squares sufficient condition, then this turns out to be equivalent to the original matrix M being decomposed as the sum of a positive semidefinite and an elementwise nonnegative matrix, i.e.

$$M = P + N, \qquad P \succeq 0, \quad n_{ij} \geq 0. \tag{7.6}$$

This is a well-known sufficient condition for copositivity (see for example [Dia62]). The equivalence between these two tests has also been noticed in [CL77, Lemma 3.5].

The advantage of the approach is that stronger conditions can be derived. By considering higher order forms, a hierarchy of increasingly powerful tests is obtained. Of course, the computational requirements increase accordingly.

Take for example the family of $2(r+2)$-forms given by

$$P_r(\mathbf{z}) = \left(\sum_{i=1}^{n} z_i^2 \right)^r P(\mathbf{z}).$$

Then it is easy to see that if P_i is a sum of squares, then P_{i+1} is also a sum of squares. The converse proposition does not necessarily hold, i.e. P_{i+1} can be a sum of squares, while P_i is not. Additionally, if $P_r(\mathbf{z})$ is nonnegative, then so is $P(\mathbf{z})$. So, by testing if $P_r(\mathbf{z})$ is a sum of squares (which can be done using SDP methods, as described), we can guarantee the nonnegativity of $P(z)$, and as a consequence, copositivity of M.

For concreteness, we will analyze in some detail the case $r = 1$, i.e., the sixth order form

$$P_1(\mathbf{z}) := \sum_{i,j,k} m_{ij} z_i^2 z_j^2 z_k^2.$$

The associated SDP test can be presented in the following theorem:

Theorem 7.7. *Consider the SDPs:*

$$\begin{aligned} M - \Lambda^i &\succeq 0, & i &= 1, \ldots, n \\ \Lambda^i_{ii} &= 0, & i &= 1, \ldots, n \\ \Lambda^i_{jj} + \Lambda^j_{ji} + \Lambda^j_{ij} &= 0, & i &\neq j \\ \Lambda^i_{jk} + \Lambda^j_{ki} + \Lambda^k_{ij} &\geq 0, & i &\neq j \neq k \end{aligned} \tag{7.7}$$

where the n matrices $\Lambda^i \in \mathcal{S}^n$ are symmetric ($\Lambda^i_{jk} = \Lambda^i_{kj}$). If there exists a feasible solution, then $P_1(\mathbf{z})$ is nonnegative, and therefore M is copositive. Furthermore, this test is at least as powerful as condition (7.6).

This hierarchy of enhanced conditions for matrix copositivity, together with results by Powers and Reznick [PR01], has been recently employed by De Klerk and Pasechnik [dKP02] in the formulation of strengthened bounds for the stability number of a graph. A very interesting result in that paper is an explicit example of a copositive matrix $M \in \mathcal{S}^{12}$, for which the test corresponding to $r = 1$ is not conclusive.

Acknowledgements. I would like to acknowledge the useful comments of my advisor John Doyle, Stephen Boyd, and Bernd Sturmfels. In particular, Bernd suggested the example in Section 7.3. I also thank the reviewers, particularly Reviewer #2, for their useful and constructive comments.

References

[BCR98] J. Bochnak, M. Coste, and M-F. Roy. *Real Algebraic Geometry*. Springer, 1998.

[BL68] N. K. Bose and C. C. Li. A quadratic form representation of polynomials of several variables and its applications. *IEEE Transactions on Automatic Control*, 14:447–448, 1968.

[Bos82] N. K. Bose. *Applied multidimensional systems theory*. Van Nostrand Reinhold, 1982.

[BPR96] S. Basu, R. Pollack, and M.-F. Roy. On the combinatorial and algebraic complexity of quantifier elimination. *J. ACM*, 43(6):1002–1045, 1996.

[Bro87] W. D. Brownawell. Bounds for the degrees in the Nullstellensatz. *Annals of Mathematics*, 126:577–591, 1987.

[BS91] C. A. Berenstein and D. C. Struppa. Recent improvements in the complexity of the effective Nullstellensatz. *Linear Algebra and its Applications*, 157(2):203–215, 1991.

[CL77] M. D. Choi and T. Y. Lam. An old question of Hilbert. *Queen's papers in pure and applied mathematics*, 46:385–405, 1977.

[CLO97] D. A. Cox, J. B. Little, and D. O'Shea. *Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra*. Springer, 1997.

[CLR95] M. D. Choi, T. Y. Lam, and B. Reznick. Sums of squares of real polynomials. *Proceedings of Symposia in Pure Mathematics*, 58(2):103–126, 1995.

[CPS92] R. W. Cottle, J. S. Pang, and R. E. Stone. *The linear complementarity problem*. Academic Press, 1992.

[Dia62] P. H. Diananda. On non-negative forms in real variables some or all of which are non-negative. *Proceedings of the Cambridge Philosophical Society*, 58:17–25, 1962.

[dKP02] E. de Klerk and D.V. Pasechnik. Approximating the stability number of a graph via copositive programming. *SIAM Journal on Optimization*, 12(4):875–892, 2002.

[Fer98] Ch. Ferrier. Hilbert's 17th problem and best dual bounds in quadratic minimization. *Cybernetics and Systems Analysis*, 34(5):696–709, 1998.

[Fu98] M. Fu. Comments on "A procedure for the positive definiteness of forms of even order". *IEEE Transactions on Automatic Control*, 43(10):1430, 1998.

[GP01] K. Gatermann and P. A. Parrilo. Symmetry groups, semidefinite programs, and sums of squares. In preparation, 2001.

[GV02] D. Grigoriev and N. Vorobjov. Complexity of Null- and Positivstellensatz proofs. *Annals of Pure and Applied Logic*, 113(1-3):153–160, 2002.

[GW95] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995.

[HH96] M. A. Hasan and A. A. Hasan. A procedure for the positive definiteness of forms of even order. *IEEE Transactions on Automatic Control*, 41(4):615–617, 1996.

[Ily92] N. V. Ilyushechkin. The discriminant of the characteristic polynomial of a normal matrix. *Math. Zameski*, 51:16–23, 1992. English translation in Math. Notes 51, 1992, pp. 230-235.

[Kol88] J. Kollár. Sharp effective Nullstellensatz. *J. Amer. Math. Soc.*, 1:963–975, 1988.

[Las01] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2001.

[Lau] M. Laurent. A comparison of the Sherali-Adams, Lovász-Schrijver and Lasserre relaxations for 0-1 programming. Technical Report PNA-R0108, CWI, Amsterdam, 2001.

[Lax98] P. D. Lax. On the discriminant of real symmetric matrices. *Communications on Pure and Applied Mathematics*, LI:1387–1396, 1998.

[Lov94] L. Lovász. Stable sets and polynomials. *Discrete mathematics*, 124:137–153, 1994.

[LS91] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991.

[Mis93] B. Mishra. *Algorithmic Algebra*. Springer-Verlag, 1993.

[MK87] K. G. Murty and S. N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39:117–129, 1987.

[Mun99] N. Munro, editor. *The Use of Symbolic Methods in Control System Analysis and Design*, volume 56 of *Control engineering series*. IEE Books, 1999.

[Nes00] Y. Nesterov. Squared functional systems and optimization problems. In J.B.G. Frenk, C. Roos, T. Terlaky, and S. Zhang, editors, *High Performance Optimization*, pages 405–440. Kluwer Academic Publishers, 2000.

[Par00a] P. A. Parrilo. On a decomposition for multivariable forms via LMI methods. In *Proceedings of the American Control Conference*, 2000.

[Par00b] P. A. Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, May 2000. Available at `http://www.cds.caltech.edu/~pablo/`.

[Par02a] P. A. Parrilo. An explicit construction of distinguished representations of polynomials nonnegative over finite sets. Technical Report IfA Technical Report AUT02-02. Available from `http://www.aut.ee.ethz.ch/~parrilo`, ETH Zürich, 2002.

[Par02b] P. A. Parrilo. Positivstellensatz and distinguished representations approaches for optimization: a comparison. In preparation, 2002.

[PPP02] S. Prajna, A. Papachristodoulou, and P. A. Parrilo. *SOSTOOLS: Sum of squares optimization toolbox for MATLAB*, 2002. Available from `http://www.cds.caltech.edu/sostools` and `http://www.aut.ee.ethz.ch/~parrilo/sostools`.

[PR01] V. Powers and B. Reznick. A new bound for Pólya's theorem with applications to polynomials positive on polyhedra. *J. Pure Appl. Algebra*, 164(1-2):221–229, 2001.

[PS01] P. A. Parrilo and B. Sturmfels. Minimizing polynomial functions. Submitted to the DIMACS volume of the Workshop on Algorithmic and Quantitative Aspects of Real Algebraic Geometry in Mathematics and Computer Science. Available from `arXiv:math.OC/0103170`, 2001.

[PW98] V. Powers and T. Wörmann. An algorithm for sums of squares of real polynomials. *Journal of pure and applied algebra*, 127:99–104, 1998.

[QDRT98] A. J. Quist, E. De Klerk, C. Roos, and T. Terlaky. Copositive relaxation for general quadratic programming. *Optimization methods and software*, 9:185–208, 1998.

[Raj93] A. R. Rajwade. *Squares*, volume 171 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 1993.

[Rez78] B. Reznick. Extremal PSD forms with few terms. *Duke Mathematical Journal*, 45(2):363–374, 1978.

[Rez95] B. Reznick. Uniform denominators in Hilbert's seventeenth problem. *Math Z.*, 220:75–97, 1995.

[Rez00] B. Reznick. Some concrete aspects of Hilbert's 17th problem. In *Contemporary Mathematics*, volume 253, pages 251–272. American Mathematical Society, 2000.

[SA90] H. D. Sherali and W. P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM J. Disc. Math.*, 3(3):411–430, 1990.

[SA99] H. D. Sherali and W. P. Adams. *A Reformulation-Linearization Technique for Solving Discrete and Continuous Nonconvex Problems*. Kluwer Academic Publishers, 1999.

[Sch91] K. Schmüdgen. The k-moment problem for compact semialgebraic sets. *Math. Ann.*, 289:203–206, 1991.

[Sho87] N. Z. Shor. Class of global minimum bounds of polynomial functions. *Cybernetics*, 23(6):731–734, 1987. (Russian orig.: Kibernetika, No. 6, (1987), 9–11).

[Sho98] N. Z. Shor. *Nondiferentiable Optimization and Polynomial Problems*. Kluwer Academic Publishers, 1998.

[Ste74] G. Stengle. A Nullstellensatz and a Positivstellensatz in semialgebraic geometry. *Math. Ann.*, 207:87–97, 1974.

[Ste96] G. Stengle. Complexity estimates for the Schmüdgen Positivstellensatz. *J. Complexity*, 12(2):167–174, 1996.

[Stu99] J. Sturm. *SeDuMi version 1.03.*, September 1999. Available from `http://www.unimaas.nl/~sturm/software/sedumi.html`.

[Väl86] H. Väliaho. Criteria for copositive matrices. *Linear Algebra and its applications*, 81:19–34, 1986.

[VB96] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, March 1996.

[WSV00] H. Wolkowicz, R. Saigal, and L. Vandenberghe, editors. *Handbook of Semidefinite Programming*. Kluwer, 2000.