

Higher order semidefinite relaxations for quadratic programming

Pablo A. Parrilo
pablo@cds.caltech.edu
Control & Dynamical Systems 107-81
California Institute of Technology
Pasadena, CA 91125-8100

Abstract

In this paper, we present improved versions of the standard semidefinite relaxation for quadratic programming, that underlies many important results in robustness analysis and combinatorial optimization. It is shown that the proposed polynomial time convex conditions are at least as strong as the standard ones, and usually better, but at a higher computational cost. Several applications of the new relaxations are provided, including less conservative upper bounds for the structured singular value μ and enhanced solutions for the MAX CUT graph partitioning problem.

1 Introduction

Many problems in systems and control theory, especially in robustness analysis and synthesis, have intrinsically “bad” computational complexity properties. These features (for example, being NP-hard) are specific to the problem class, and not associated with any particular algorithm used in its solution. In the case of NP-hardness, in particular, the practical implications are well known: unless $P=NP$, every algorithm that solves the problem will take at least an exponential number of steps, in the worst case.

For this reason, it is particularly useful to count with alternative methods, guaranteed to run in a “reasonable” time, that provide bounds on the optimal solution and/or suboptimal estimates. In the particular case of quadratic programming (QP), such a tool has been made available in the last few years. Semidefinite programming (SDP) relaxations of nonconvex QP problems have increasingly been used for a variety of problems in very diverse fields of applied mathematics. These SDP relaxations are convex optimization problems, that can be solved in polynomial time. The procedure by which a relaxed problem and its dual are obtained is known in the literature under several different names: S-procedure, Shor relaxation, covariance relaxation, lifting, etc. [19]. For certain specific cases

(such as the MAX CUT problem discussed below) these approximate solutions are provably good, as there exist hard bounds on their degree of suboptimality. However, some other problems (for instance, MAX CLIQUE, or real μ [4]) are significantly harder, since even the approximation problem within an arbitrary constant factor is NP-hard.

The standard relaxations (basically, duality) underlie a quite large number of results in many different application areas. For nonconvex problems, however, there exists the possibility of having a “duality gap,” and the original and relaxed problems will have different solutions. In this paper, we present novel convex relaxations of quadratic programming problems, that can be run in polynomial time, and provide improved gaps by combining the constraints in a *nonlinear* fashion (standard duality deals only with *linear* combination of constraints). The idea can be interpreted as finding a separating functional, not necessarily linear, that proves that the intersection of two sets is empty. We employ as a basic technique the existence of a sum of squares decomposition as a sufficient condition for nonnegativity of a multivariable form.

The generality of the problem formulation makes possible the application of the results in many different areas: the chosen examples, from robustness analysis and combinatorial optimization, are just a few cases where the presence of gaps has been recognized and analyzed in the past. As we will see, our methods provide notable improvements for these well-known examples.

The outline of the paper is as follows: in section 2 the standard semidefinite relaxation for QP problems is reviewed. In the next section, we illustrate the use of LMI methods for obtaining a sum of squares decomposition of a form in several variables, a crucial step in our development. In section 4, new convex relaxations are proposed, based on the concept of copositive functionals and the presented decomposition. There it is shown that in general it is stronger than the standard bounds, and insight about its computational complexity and its relationship with certain procedures in alge-




Figure 1: Plant M and uncertainty diagonal structure Δ

braic geometry are presented. Finally, the last section features some numerical examples, dealing with MAX CUT problems and complex μ computation, followed by conclusions and directions for future research.

2 The standard SDP relaxation

The viewpoint taken in this paper focuses on considering the SDP relaxation as a *sufficient condition* for establishing that a certain set \mathcal{A} , described by quadratic inequalities, is empty. Concretely, given m symmetric matrices $A_1, \dots, A_m \in \mathbb{R}^{n \times n}$, we define the set \mathcal{A} as the intersection of the image of \mathbb{R}^n under the quadratic forms and the positive orthant:

$$\mathcal{A} := \{ \mathbf{z} \in \mathbb{R}^m \mid z_i := \mathbf{x}^T A_i \mathbf{x} \geq 0, \mathbf{x} \in \mathbb{R}^n / \{0\} \} \quad (1)$$

For future reference, given symmetric matrices A_i , let $A_i(\mathbf{x}) := \mathbf{x}^T A_i \mathbf{x}$, and $\mathbf{a}(\mathbf{x}) := [\mathbf{x}^T A_1 \mathbf{x}, \dots, \mathbf{x}^T A_m \mathbf{x}]^T$. Both logical implications and constrained optimization problems can be posed in the framework described above, by checking for the existence of a counterexample, or a feasible point that achieves a given level of optimality, respectively. When the matrices A_i are indefinite, checking if \mathcal{A} is empty is clearly an NP-hard problem.

A simple sufficient condition for the set \mathcal{A} defined in (1) to be empty is given by the existence of numbers λ_i that satisfy the condition:

$$\sum_{i=1}^m \lambda_i A_i < 0, \quad \lambda_i \geq 0. \quad (2)$$

The reasoning is very simple: if \mathcal{A} is not empty, then there exists a point $\mathbf{x} \neq 0$ such that the inner product of $\mathbf{a}(\mathbf{x})$ and λ is nonnegative, since both vectors are componentwise nonnegative. However, equation (2) makes that inner product negative. As a consequence, \mathcal{A} is empty.

As mentioned earlier, the procedure by which the condition (2) is derived has received several different names in the literature: perhaps the most correct term for this is simply *duality*, recognizing the possibility of having

a nonzero duality gap. Note that condition (2) is a linear matrix inequality (LMI), also known as an instance of a *semidefinite program* [19]. As is widely recognized today, this class of convex optimization problems can be efficiently solved, both in theory and practice.

Example 1 *As a typical example of a robustness problem that can be posed in this form, consider the case of a standard structured singular value μ problem [11]. For simplicity, let the matrix $M \in \mathbb{R}^{n \times n}$, $\Delta = \text{diag}(\delta_1, \dots, \delta_n)$ and the scalar uncertainties δ_i be real. In the notation of Figure 1, the condition that the absolute value of the uncertainties δ_i is bounded by one, is equivalent to the quadratic inequalities:*

$$\delta_i^2 \leq 1 \iff y_i^2 - x_i^2 = \mathbf{x}^T (M_i^T M_i - E_{ii}) \mathbf{x} \leq 0, \quad (3)$$

where E_{ii} is the matrix with zero elements, except for a one in the (i, i) position, and M_i is the i th row of the matrix M . Therefore, for this particular case, the matrices A_i in 1 are given by $A_i = M_i^T M_i - E_{ii}$. In this case, for example, the inexistence of nontrivial solutions could be interpreted as the robust stability of the system under uncertainty.

Applying the SDP relaxation to the system of inequalities (3) we obtain the usual μ upper bound LMI:

$$M^T D M - D < 0, D > 0.$$

where D is a diagonal matrix. \square

It is also interesting to study the dual problem, in the SDP sense, of (2). It consists of checking for the existence of a symmetric matrix $Z \neq 0$, that satisfies

$$\text{Tr} A_i Z \geq 0, \quad Z \geq 0. \quad (4)$$

This dual problem can also be obtained directly from (1), by using the cyclic properties of the trace function, and dropping the rank one condition on the matrix $Z := \mathbf{x} \mathbf{x}^T$ [19]. If the dual problem (4) does not have a solution, then neither does the original one. But at least in principle, an affirmative answer to the feasibility of (4) does not necessarily say anything about the set \mathcal{A} (in some special cases, however, it is possible to extract useful information from the matrix Z).

In the upcoming sections, we will see how to extend the standard condition (2) to a more general sufficiency test. To this end, we shortly review in the next section a useful sufficient condition for the nonnegativity of a form of several variables.

3 Positivity of a multivariable form

Let $F(x)$ be a homogeneous form of degree m in n variables $\{x_1, \dots, x_n\}$, with real coefficients. A sufficient

condition for F to be nonnegative is the existence of a sum of squares decomposition $F = \sum_i f_i^2(x)$ (of course, m has to be even). If the $f_i(x)$ are required to be rational functions, such a decomposition is always possible. This is Hilbert's 17th problem, solved in 1927 by E. Artin. When the functions $f_i(x)$ are constrained to be homogeneous forms, it is known since Hilbert that only in certain cases (depending on n and m) this decomposition is guaranteed to exist [16]. It has been shown [18, 15, 12] that this decomposition, if it exists, can be obtained by using semidefinite programming methods. We briefly review the methodology below; a complete description can be found in the references mentioned earlier.

The basic idea of the method is the following: express the given form as a quadratic form in some new variables z . These new variables are all the monomials of degree equal to $\frac{1}{2}m$ given by the different products of the x variables. Therefore, $F(x)$ can be represented as

$$F(x) = z^T Q z, \quad (5)$$

where Q is a constant matrix. If in the representation above Q is positive semidefinite, then $F(x)$ obviously takes only nonnegative values. In principle, this condition is conservative, generally speaking. The main reason is that since the variables z_i are not independent the matrix Q in (5) might not be unique, and it may be positive semidefinite for some representations, but not for others. Fortunately, this conservativeness can be avoided by using identically satisfied constraints that relate the z_i variables among themselves (of the form $z_i z_j = z_k z_l$ or $z_i^2 = z_k z_l$). This way, it is easily shown that there is a linear subspace of matrices Q that satisfy (5). If the intersection of this subspace with the positive semidefinite matrix cone is nonempty (which can be checked by solving an LMI feasibility problem), then the original function F is guaranteed to be a sum of squares, and therefore positive semidefinite. This follows from a decomposition of $Q = L^T L$, which implies the sum of squares representation $F(x) = \sum_i (Lz)_i^2$. Conversely, if F can indeed be written as the sum of squares of forms, then expanding in monomials will provide the representation (5).

Example 2 Consider the following quartic form in two variables:

$$\begin{aligned} F(x_1, x_2) &= 2x_1^4 + 2x_1^3 x_2 - x_1^2 x_2^2 + 5x_2^4 \\ &= \begin{bmatrix} x_1^2 \\ x_2^2 \\ x_1 x_2 \end{bmatrix}^T \begin{bmatrix} 2 & 0 & 1 \\ 0 & 5 & 0 \\ 1 & 0 & -1 \end{bmatrix} \begin{bmatrix} x_1^2 \\ x_2^2 \\ x_1 x_2 \end{bmatrix} \\ &= \begin{bmatrix} x_1^2 \\ x_2^2 \\ x_1 x_2 \end{bmatrix}^T \begin{bmatrix} 2 & -\lambda & 1 \\ -\lambda & 5 & 0 \\ 1 & 0 & -1 + 2\lambda \end{bmatrix} \begin{bmatrix} x_1^2 \\ x_2^2 \\ x_1 x_2 \end{bmatrix} \end{aligned}$$

Take for instance $\lambda = 3$. In this case,

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

And therefore we have the sum of squares decomposition:

$$F(x_1, x_2) = \frac{1}{2}(2x_1^2 - 3x_2^2 + x_1 x_2)^2 + \frac{1}{2}(x_2^2 + 3x_1 x_2)^2.$$

□

4 Separating functionals and a new SDP relaxation

In order to extend the standard condition we consider the well-known interpretation of the multipliers λ_i in (2) as defining a *separating hyperplane* (or a linear functional). To see this, notice that the positivity condition on the multipliers λ_i guarantees that the linear functional $\phi(z) = \lambda^T z$ is positive in the positive orthant. Additionally, condition (2) ensures that this functional is negative on the image of \mathbb{R}^n under the map \mathbf{a} . Therefore, those two sets have empty intersection, which is what we want to prove.

Understanding this idea, the proposed method is conceptually simple: replace the linear form by a more general function. For consistency with the linear case, we keep using the term “functional” to refer to these mappings, see for example [8, Section 13.5]. For concreteness, here we consider only the case of quadratic functionals, though the extension to the general case is straightforward. The reasons are also practical: the complexity of checking nonnegativity of forms of high degree grows quite fast. Even in the relatively simple case of quartic forms (as in the case we will be analyzing), the computational requirements can be demanding.

A functional $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is *copositive* if $x_i \geq 0$ implies $\phi(\mathbf{x}) \geq 0$, i.e. is positive in the positive orthant. In this case, it is clear that a sufficient condition for \mathcal{A} being empty is the existence of a copositive functional ϕ such that:

$$\phi(\mathbf{a}(\mathbf{x})) < 0, \quad \forall \mathbf{x} \in \mathbb{R}^n / \{0\} \quad (6)$$

The reasons are exactly the ones mentioned earlier: the existence of a candidate \mathbf{x} that makes $\mathbf{a}(\mathbf{x})$ nonnegative forces the composition of the functions to be positive or zero, contradicting the condition above. Note that *the same conclusions hold if ϕ itself depends on \mathbf{x}* , as long as it is always copositive.

Two questions immediately arise: How do we characterize copositive functionals, and how do we check condition (6)? From a complexity viewpoint, these two questions are as intractable as the original problem. It turns out that for the case of polynomial (or rational) functionals ϕ , a partial answer to both questions can be obtained by using the sum of squares decomposition presented in the last section.

For the exact copositivity problem, existing results show that checking if a quadratic form is not copositive is an NP-complete problem [10]. A simple sufficient condition for copositivity of a matrix Φ (see [13] for stronger SDP-based copositivity tests) is given by the existence of a decomposition of Φ as the sum of two matrices, one positive definite and the other one componentwise nonnegative, i.e.:

$$\Phi = P + N, \quad P \geq 0, \quad n_{ij} \geq 0.$$

Note that without loss of generality, we can always take the diagonal elements of N to be zero.

Therefore, we can consider quadratic copositive functionals ϕ of the form above (i.e. $\phi(\mathbf{v}) := \mathbf{v}^T \Phi_{\mathbf{x}} \mathbf{v}$), applied to the vector $[1, \mathbf{a}(\mathbf{x})]^T$, since we want to allow for linear terms too. Mainly for computational reasons, we would like the left-hand side of (6) to be a homogeneous form. This imposes certain constraints in the structure of ϕ . It can be verified that the positive definite part of ϕ cannot help in making the form negative definite. Based on all these facts, a sufficient condition for \mathcal{A} being empty is presented next, where we also consider the case of equality constraints.

Theorem 1 *Assume there exist solutions $Q_i, T_i \in \mathbb{R}^{n \times n}, r_{ij} \in \mathbb{R}$ to the inequality:*

$$\sum_{i=1}^{n_a} Q_i(x) A_i(x) + \sum_{1 \leq i < j \leq n_a} r_{ij} A_i(x) A_j(x) + \sum_{j=1}^{n_b} T_j(x) B_j(x) < 0, \quad \forall x \in \mathbb{R}^n / \{0\}. \quad (7)$$

where $Q_i \geq 0$ and $r_{ij} \geq 0$. Then, the only solution of

$$\begin{aligned} A_i(x) &\geq 0, & i = 1, \dots, n_a \\ B_i(x) &= 0, & i = 1, \dots, n_b \end{aligned}$$

is $x = 0$.

Proof: *It basically follows from the same arguments as in the linear case: the existence of a nontrivial x implies a contradiction. Therefore, the set \mathcal{A} is necessarily empty. \square*

Note that the left-hand side of the equation above 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) can be written (except for the sign change) as a *sum of squares*. As we have seen in section 3, 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 (2) is feasible, then the left-hand side of (7) is obviously a sum of squares (recall that positive definite quadratic forms are always sums of squares).

Remark 1 *It is interesting to compare this condition with the Nullstellensatz and Positivstellensatz used in algebraic geometry [2, Theorem 4.4.2]. These are necessary and sufficient algebraic tests for checking whether or not a given semialgebraic set is empty. For example, for the equality-only case, Hilbert's Nullstellensatz says that if we have polynomials f_1, \dots, f_m such that there is no common solution of the equations $f_i(\mathbf{x}) = 0$, then there exists polynomials g_i such that*

$$g_1 f_1 + g_2 f_2 + \dots + g_m f_m = 1.$$

The converse proposition, of course, always holds. In other words, the constant polynomial belongs to the ideal generated by the f_i . There are known extensions to the inequality case, involving so-called cones and other algebraic objects. These involve sums of squares and the products of the functions defining the inequalities, in a manner very similar to (7).

The connections between these conditions are further explored in [14], where Positivstellensatz-based methods are employed in the formulation of strong semidefinite programming conditions for many problems in systems and control theory.

It is often the case that one of the quadratic forms, say A_1 , depends on a certain parameter γ , and we are interested in finding the smallest (or largest) value of γ for which the set $\mathcal{A}(\gamma)$ is empty. In this case, when we take into account the γ dependence of A_1 , the problem of testing feasibility of (7) is no longer an LMI, since we have products of γ and the decision variables Q_i and q_{ij} . There are two possible remedies to this problem: the first one is to remember that even though (7) is not a semidefinite program, it is still a quasiconvex problem, since for fixed γ the level sets are convex. The alternative is to fix some of the variables (for example, $Q_1 = I$, and $q_{1j} = 1$), to make the left-hand side of (7) linear in γ . In principle, this last technique can be conservative, when comparing to the case where all the variables are free.

4.1 Computational complexity

A few words on the complexity of the proposed procedure are in order. When solving the relaxation using standard software, the main burden lies in the computation of the solution of the resulting system of LMIs, in particular due to the need of checking if the resulting quartic form is a sum of squares. The LMI corresponding to this condition has a matrix size equal to $\binom{n+1}{2}$, provided no simplifications occur. However, the main difficulty is really caused by the large number of variables arising from the redundant constraints, since its number is $O(n^4)$. Even though it is polynomial (and therefore the whole procedure runs in polynomial time), this rapid growth rate is not quite acceptable. In many special cases, symmetry considerations can help




Figure 2: The Petersen Graph

reduce the number significantly. However, for the general case with a large number of variables, alternative approaches are certainly needed. Some concrete possibilities, currently under study, are to exploit problem structure, and to incorporate only a certain subset of variables into the optimization.

5 Examples

In this section, we present some examples from robustness analysis and combinatorial optimization, that show the advantages of the improved relaxations, when compared to the standard procedures.

5.1 Structured singular value upper bound

As mentioned in Remark 1, the standard upper bound of the structured singular value μ [11] can be interpreted as the result of applying the standard relaxation to the quadratic forms defining the uncertainty structure. It is therefore a natural test problem for the techniques introduced in this paper.

Consider the counterexample, due to Morton and Doyle, to the proposition that μ is equal to its standard upper bound in the case with four scalar uncertainties [11, Section 9.2]. This corresponds to a certain rank two matrix $M \in \mathbb{C}^{4 \times 4}$. This matrix has a value of $\mu(M) \approx 0.8723$. However, the standard μ upper bound has an exact value of 1. For this problem, with the improved relaxation, we are able to prove an upper bound of 0.895 by solving a semidefinite program.

5.2 The MAX CUT problem

The maximum cut (MAX CUT) problem consists in finding a partition of the nodes of a graph in two disjoint subsets V_1 and V_2 , in such a way to maximize the number of edges that have an endpoint in V_1 and the other in V_2 . It has important practical applications, such as optimal circuit layout. The decision version of this problem (does there exist a cut with value greater than or equal to K ?) is known to be NP-complete [5].

By casting the problem as a boolean maximization, we

can write the MAX CUT problem as an equality constrained quadratic program. One standard formulation [6] is the following:

$$\max_{y_i \in \{-1, 1\}} \frac{1}{2} \sum_{i,j} w_{ij} (1 - y_i y_j), \quad (8)$$

where w_{ij} is the weight corresponding to the (i, j) edge, and is zero if the nodes i and j are not connected. The constraints $y_i \in \{-1, 1\}$ are equivalent to the quadratic constraints $y_i^2 = 1$.

We can obtain useful upper bounds on the optimal value of (8) using semidefinite programming. Removing the constant term, and changing the sign, the original problem is clearly equivalent to:

$$\min_{y_i^2=1} \sum_{i,j} w_{ij} y_i y_j. \quad (9)$$

The corresponding semidefinite relaxation is given by:

$$\min_{Y \geq 0, Y_{ii}=1} \text{trace } WY, \quad (10)$$

and its dual

$$\max_{D \leq W} \text{trace } D, \quad (11)$$

where D is a diagonal matrix. Any feasible solution of the dual (11) provides a lower bound on the optimal value of (10), and therefore on that of (9).

A simple case where both the exact problem and the standard SDP relaxation can be analyzed is that of the n -cycle C_n . This is a graph with n nodes and n edges, where the edges form a closed chain. In other words, if the vertices are numbered from v_1 to v_n , then all the edges have the form (v_i, v_{i+1}) , where $v_{n+1} = v_1$. For this graph, the exact value for the unweighted MAX CUT problem can easily be shown to be equal to n if n is even, or $n - 1$ otherwise.

In the case of even n , the standard relaxation provides a bound that is exact, i.e., equal to n . For the odd n case, we have the upper bound

$$MC(C_n) \leq n \cos^2 \frac{\pi}{2n}.$$

For this class of graphs, the gap is maximal in the case of the 5-cycle ($n = 5$). The optimal solution is 4, but the computed upper bound is equal to $\frac{5}{8}(5 + \sqrt{5}) \approx 4.5225$. When applying the procedure developed in this paper to the n -cycle, we recover the *optimal* solution, i.e., the new relaxation has *zero gap*.

Consider now the Petersen graph, shown in Figure 2. This nonplanar graph has ten nodes and fifteen edges, and has very interesting theoretical properties [7]. For the unit weight case described (i.e., when we only count the number of edges cut), the optimal solution can be

shown to be 12. The solution of the standard semidefinite relaxation for this problem is equal to 12.5. When applying the new relaxation to this problem, we are able to obtain the exact value 12.

In the paper [1], a different strengthened SDP relaxation for the specific case of MAX CUT is presented. Even though the results in that paper provide improved bounds over the standard relaxation, in neither the case of the 5-cycle nor the Petersen graph the obtained bounds are exact. Of course, a fair comparison should also take into account the computational requirements, which are higher in our proposed method than in that of [1].

6 Conclusions

A new polynomial time scheme for computing bounds on the optimal solution of hard nonconvex problems was introduced. The resulting estimates are always at least as strong as the ones obtained by the traditional semidefinite relaxation. The key ideas are to combine constraints *nonlinearly*, as opposed to standard duality, and to use the sum of squares decomposition as a sufficient condition for nonnegativity of a function. The results obtained from its application to a few test problems are certainly encouraging: tighter (or even exact) bounds can be obtained. Of course, more study is needed in order to fully assess its potential relevance.

While it is perfectly possible to immediately apply the developed tools to small problems, an important issue is certainly the computational feasibility of applying these relaxations to medium or large scale problems. Though in principle all the relaxations are polynomial time algorithms, for large size instances even the standard semidefinite relaxation can be troublesome, only in part due to the relative immaturity of current SDP solvers. Clearly, more research is needed in the implementation aspects, especially on the issue of exploiting problem structure. Some recent interesting approaches, such as the work in [3] on MAX CUT, show that there is lot of room for improvement, especially when working in specific problem classes.

An important research topic is to understand the connections with related ideas that have been explored in the context of “lift-and-project” methods [9, 17] for deriving valid inequalities in zero-one combinatorial optimization problems. Some common elements of the approaches are the use of new variables and constraints, defined as products of the original ones, and the use of semidefinite constraints (in the Lovász-Schrijver N_+ relaxation). The main feature of our work, however, is the extension to the general continuous case via the use of the sum of squares decomposition.

References

- [1] M. F. Anjos and H. Wolkowicz. An strengthened SDP relaxation via a second lifting for the MAXCUT problem. Preprint, available at <http://orion.math.uwaterloo.ca/~hwolkowi/henry/reports/ABSTRACT> 1999.
- [2] J. Bochnak, M. Coste, and M.-F. Roy. *Real Algebraic Geometry*. Springer, 1998.
- [3] S. Burer and R. Monteiro. An efficient algorithm for solving the MAXCUT SDP relaxation. Preprint, available at http://www.isye.gatech.edu/~monteiro/tech_reports/maxcut.dvi.g 1998.
- [4] M. Fu. The real structured singular value is hardly approximable. *IEEE Transactions on Automatic Control*, 42(9):1286–1288, 1997.
- [5] M. R. Garey and D. S. Johnson. *Computers and Intractability: A guide to the theory of NP-completeness*. W. H. Freeman and Company, 1979.
- [6] 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.
- [7] D. A. Holton and J. Sheehan. *The Petersen graph*, volume 7 of *Australian Mathematical Society Lecture Series*. Cambridge University Press, 1993.
- [8] A. N. Kolmogorov and S. V. Fomin. *Introductory Real Analysis*. Dover, 1970.
- [9] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991.
- [10] K. G. Murty and S. N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39:117–129, 1987.
- [11] A. Packard and J. C. Doyle. The complex structured singular value. *Automatica*, 29(1):71–109, 1993.
- [12] P. A. Parrilo. On a decomposition for multivariable forms via LMI methods. In *Proceedings of the American Control Conference*, 2000.
- [13] P. A. Parrilo. Semidefinite programming based tests for matrix copositivity. In *Proceedings of the 39th IEEE Conference on Decision and Control*, 2000.
- [14] P. A. Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, 2000.
- [15] V. Powers and T. Wörmann. An algorithm for sums of squares of real polynomials. *Journal of pure and applied algebra*, 127:99–104, 1998.
- [16] A. R. Rajwade. *Squares*, volume 171 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, 1993.
- [17] 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.
- [18] 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).
- [19] L. Vandenberghé and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, Mar. 1996.