Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization

Thesis by
Pablo A. Parrilo

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California

2000

(Defended May 18, 2000)

© 2000
Pablo A. Parrilo
All rights Reserved

To my mother

Acknowledgements

Naturally, my greatest appreciation goes to my advisor, John Doyle. By providing a fertile ground, supervision, and an unsurpassed academic environment, John made my Caltech experience unique. I will be forever thankful for this opportunity.

My gratitude extends to the other members of my thesis committee: Jerry Marsden, Richard Murray, and Mani Chandy. The academic formation at CDS has been much more than what I expected, and for this I am extremely grateful. Additionally, I would like to recognize Richard's professional attitude and uncommon equitability, that makes him an outstanding example for all of us.

This is also a good opportunity to thank Ricardo Sánchez Peña, for his mentoring during my undergrad years at the University of Buenos Aires, and Mario Sznaier, for many interesting collaborations and advice. I owe them my gratitude, not only for the professional relationship, but more importantly, for their friendship.

To all my friends here, thank you for the Red Door coffee breaks, dinners, movies, and endless raves and rantings. In particular, to my closest friends and CDS mates: Alfred Martinez, for his noble heart, Monica Giannelli, for her boundless optimism, and Sergey Pekarsky, just for being Sergey. Fecho Spedalieri and Martín Basch, with their good humor, made our frequent late night hamburgers taste better. I also want to mention Sven Khatri and Sanj Lall for showing me the ropes around Caltech, and for the many new things they taught me.

Mauricio Barahona, Ali Jadbabaie, Jim Primbs, Luz Vela, Matt West, Jie Yu, Xiaoyun Zhu, and the rest of the CDS gang, with their friendship and conversation, made the afternoons in the Steele basement much more enjoyable.

I want to thank Jos Sturm, for writing and making freely available his excellent software SeDuMi [86], and for pointing out some interesting references.

Finally, I would like to manifest my deepest gratitude and appreciation to my family, for their relentless support and unconditional love. The spirit and courage of my mother Victoria, and the affection and loyalty of my sister Marita, were always with me all these years. I owe them both much more than I would ever be able to express.

Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization

by

Pablo A. Parrilo

In Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

Abstract

In the first part of this thesis, we introduce a specific class of Linear Matrix Inequalities (LMI) whose optimal solution can be characterized exactly. This family corresponds to the case where the associated linear operator maps the cone of positive semidefinite matrices onto itself. In this case, the optimal value equals the spectral radius of the operator. It is shown that some rank minimization problems, as well as generalizations of the structured singular value (μ) LMIs, have exactly this property.

In the same spirit of exploiting structure to achieve computational efficiency, an algorithm for the numerical solution of a special class of frequency-dependent LMIs is presented. These optimization problems arise from robustness analysis questions, via the Kalman-Yakubovich-Popov lemma. The procedure is an outer approximation method based on the algorithms used in the computation of $\mathcal{H}_\infty$ norms for linear, time invariant systems. The result is especially useful for systems with large state dimension.

The other main contribution in this thesis is the formulation of a convex optimization framework for semialgebraic problems, i.e., those that can be expressed by polynomial equalities and inequalities. The key element is the interaction of concepts in real algebraic geometry (Positivstellensatz) and semidefinite programming.

To this end, an LMI formulation for the sums of squares decomposition for multivariable polynomials is presented. Based on this, it is shown how to construct sufficient Positivstellensatz-based convex tests to prove that certain sets are empty. Among other applications, this leads to a nonlinear extension of many LMI based results in uncertain linear system analysis.

Within the same framework, we develop stronger criteria for matrix copositivity, and generalizations of the well-known standard semidefinite relaxations for quadratic programming.

Some applications to new and previously studied problems are presented. A few examples are Lyapunov function computation, robust bifurcation analysis, structured singular values, etc. It is shown that the proposed methods allow for improved solutions for very diverse questions in continuous and combinatorial optimization.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Without any doubt, one of the main mathematical developments in the last century has been the introduction of the Turing computability theory and its associated computational complexity classes. Turing's pioneering work made concrete and formal the then-vague notion of *algorithm*. By proposing a specific device (a *Turing machine*) as a representative of the ambiguous notion of *computer*, a deep understanding of the power and intrinsic limitations of algorithmic approaches was achieved for the first time.

In particular, we now have a clear understanding of the notion of the *decidability* of a problem. This fundamental concept relates to the *existence of a decision algorithm* to solve a given mathematical question. Unexpectedly at first, this cannot be taken for granted. The classical example is the Turing machine *halting problem*: does there exist a general procedure that, given a computer program as an input, will correctly decide if the program terminates?

Turing's arguments conclusively established the nonexistence of such procedure. A few years earlier, Gödel had showed that incompleteness is an intrinsic feature of mathematical systems: any logic powerful enough to include arithmetic statements will necessarily contain propositions that are neither provable nor disprovable.

It is perhaps surprising that these problems are not necessarily "artificial": many interesting questions, that have arisen independently over the past decades, have this feature. For instance, some "simple" problems in control theory can be formally shown to be *not decidable*. A nice example is given by the simultaneous stabilization problem, where we look for a common controller that will stabilize a given finite set of plants. For the case of two linear time invariant systems, it is known that the problem is equivalent to that of strong stabilization, i.e., stabilization with a stable controller, and its existence can be decided with a finite number of operations. However, in the case of three or more plants, such a procedure does not exist, and the problem is rationally undecidable [11].

Fortunately, many interesting problems in systems and control theory are indeed decidable, since they can be completely solved by purely algorithmic means. As a simple example, consider the stabilization problem for linear time invariant plants. This question can be algorithmically decided, for instance, using algebraic Riccati equations.

It is a fact that a large proportion of control problems, especially in the linear case, can be formulated using only polynomial equalities and inequalities, that 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 question for this quite large class of problems. This theory has been applied in [3], for example, to show the decidability of the static output feedback problem. Since many propositions in systems theory can be formulated on a first order logic (where quantifiers only affect variables, and not other sentences in the language), its decidability is an immediate consequence of the Tarski-Seidenberg algorithm.

However, even after the decidability question is dealt with, an important issue remains: if we have an algorithm that will solve every possible instance in the problem class, what can be said about its computational complexity? The answer to this question turns out to be delicate, and the theory of *NP-completeness* [36] is the best attempt so far to answer these issues.

The foundations of the NP-completeness theory lie in the definition of "solving" a *yes/no decision problem* as a Turing machine "recognizing" a certain element of a language, namely that corresponding to the instances for which the answer is "yes." A language will be in the class P (polynomial time) if the Turing machine is only allowed to perform deterministic operations, and it always produces a correct answer in a time that is bounded by a polynomial function of the input length. If the computing device is allowed to operate nondeterministically, then a language belongs to NP (nondeterministic polynomial) if there is a Turing machine that will accept it in polynomial time. In other words, in NP we are allowed to "guess" a solution, and only required to *verify*, in polynomial time, that the answer is "yes." A language is in co-NP if its complement is in NP.

Computational complexity theory has been very successful in the classification and understanding of many relevant practical problems. However, it is only fair to say that many important questions are still unanswered. Some "basic" propositions, such as P$\neq$NP, or NP$\neq$co-NP, though almost universally believed to be true, are still lacking proof. The implications of the separation of the complexity classes are extremely important: assuming that NP$\neq$co-NP, for problems in co-NP in general there are no polynomial time verifiable certificates of infeasibility (i.e., when the answer of the decision problem is "no"). Furthermore, the important practical issue of *approximability* is just beginning to be addressed [42]. In this respect, we should emphasize that apparently similar NP-complete problems (for example, MAX CUT and MAX CLIQUE), can have completely different approximability properties.

We mentioned earlier the existence of a constructive decision procedure (actually, a *quantifier elimination* algorithm) for first order logic over the reals. Unfortunately, the complexity of this quantifier elimination procedure (Tarski-Seidenberg, or Collins' modifications) is doubly exponential in the number of variables. For this reason, the application of general quantifier elimination procedures to practical systems and control problems, while powerful in theory, does not seem to be very promising, unless algorithmic breakthroughs or the exploitation of special structure can overcome the complexity barrier.

A thorough understanding of these issues (decidability and computational complexity) is *crucial* if we want to be able to tackle complex problems. As systems get more sophisticated, the boundaries between *dynamics* and *computation* are increasingly being blurred. A prime example of this is the case of hybrid systems, where proving stability is an algorithmically undecidable problem [89]. Additionally, the sheer size of many practically interesting problems (for example, the power grid) make computational complexity issues absolutely critical.

Faced with these facts, we should ask ourselves some questions: do our current approaches and methods provide even the hope of tackling large, nonlinear problems? What are the prospects, if any, of improving over the bounds provided by standard convex relaxation procedures?

In our work, we exploit the fundamental asymmetry inherent to the definition of complexity classes. For the class of optimization problems we generally deal with, deciding the existence of a suboptimal solution (i.e., does there exist an x with $f(x) \leq \gamma$?) is usually in NP. The reason is that, if the proposition is true, there exists a good "guess" (usually x itself) by which we can check in polynomial time that the answer is actually "yes." The converse problem, deciding if $\{x|f(x) \leq \gamma\} = \emptyset$ is therefore in co-NP. This means that in general, there are no *certificates*, that can be verified in polynomial time, to show the nonexistence of solutions.

Nevertheless, in some cases it is possible to construct such "proofs." For example, consider the problem of finding a Hamiltonian circuit in an undirected graph. If there exists a partition of the set of nodes in two disjoint subsets, connected only by one edge, then it is clear that a Hamiltonian circuit cannot exist. Such a partition, provided it exists, can be described and verified in a "small" number of operations (a polynomial function of the size of the problem). Of course, if no such partition can be found, then we do not know anything for sure about our original problem: either a Hamiltonian circuit does not exist, or the test is not powerful enough.

As we will see in the second part of this thesis, this general idea can be made concrete, and successfully applied to a class of practically interesting problems. The most important feature is that the search for *proof certificates* can be carried out in an *algorithmic way*. This is achieved by coupling efficient optimization methods and powerful theorems in semialgebraic geometry. For practical reasons, we will only be interested in the cases where we can find *"short" proofs*. A priori, there are no guarantees that a given problem has a short proof. In fact, not all problems will have short proofs, since otherwise NP=co-NP (which is not very likely). However, in general we can find short proofs that provide useful information: for instance, in the case of minimization problems, this procedure provides lower bounds on the value of the optimal solution.

The principal numerical tool used in the search for infeasibility certificates is *semidefinite programming*, a broad generalization of linear and convex quadratic optimization. Semidefinite programs, also known as Linear Matrix Inequalities (LMI) methods, are convex optimization problems, and correspond to the particular case of the convex set being the intersection of an affine family of matrices and the positive semidefinite cone. As shown in the seminal work of Nesterov and Nemirovskii [67], where a general theory of interior-point polynomial time methods for convex programming is developed, semidefinite programs can be efficiently solved both theoretically and practically. The critical ingredient there turns out to be the existence of a *computable* "self-concordant" barrier function.

The increasing popularity of LMI methods has definitely expanded the horizons of systems and control theory: it has forced the acceptance of the solution of an optimization problem as the "answer" to theoretical questions, often intractable by analytic means. It is obvious that this trend is bound to continue in the future: faster computers and enhanced algorithms will enable the application of sophisticated analysis and design methodologies, otherwise impossible to implement.

1.1 Outline and contributions

The main themes in our work are parallel, and attack simultaneously two ends of the spectrum: *special problems* with very defined characteristics, and *general tools*, that can be applied to an extremely broad class of questions.

In the first case, we show how the *special structure* in certain robustness analysis problems can be systematically exploited in order to formulate efficient algorithms. This is the motivation of Chapters 2 and 3, where a *cone invariance* property and the specific structure of the *Kalman-Yakubovich-Popov* inequalities are employed in the construction of efficient optimization procedures.

The second aspect is much more general: a framework is given for a generalization of many standard conditions and procedures in optimization and control. The central piece of the puzzle is the key role of semidefinite programming and sums of squares decompositions in the constructive application of results from semialgebraic geometry.

The main contributions of this thesis are:

- A characterization of a family of linear matrix inequalities, for which the optimal solution can be exactly described. The main feature is the notion of *cone-preserving* operators, and the associated semidefinite programs. As a consequence of a generalized version of the classical Perron-Frobenius theorem, the optimal value can be characterized as the spectral radius of an associated linear operator. It is shown that a class of robustness analysis problems are exactly of this form, and an application to some previously studied rank minimization problems is presented.

- An efficient algorithm for the solution of linear matrix inequalities arising from the Kalman-Yakubovich-Popov (KYP) lemma. This kind of LMIs are crucial in the stability and performance analysis via integral quadratic constraints (IQCs). By recasting the problem as a semi-infinite optimization problem, and the use of an *outer approximation* procedure, much more efficient solutions can be obtained.

- The sum of squares decomposition for multivariable forms is introduced, and a semidefinite programming based algorithm for its computation is presented. This makes possible the extension of LMI based methods to the analysis of a class of nonlinear systems. For example, it is shown how the new techniques enable the computation of polynomial Lyapunov functions using semidefinite programming.

- A clean and convincing description of the relationship between semialgebraic geometry results (Stellensätze) and the associated semidefinite programming sufficient tests. It is shown how the standard S-procedure can be interpreted, in the real finite dimensional case, as a Positivstellensatz refutation of fixed degree. By lifting this degree restriction, stronger sufficient conditions are derived, as shown in Chapter 6.

- The tools developed are applied in the formulation of a family of strong semidefinite relaxations of standard nonconvex quadratic programming problems. This class of relaxations provide improved bounds on the optimal solution of difficult optimization questions. The new relaxations are applied to the matrix copositivity problem, computation of the standard singular value μ, and combinatorial optimization problems such as MAX CUT. The new bounds can never be worse than those of the standard relaxation, and in many cases they are strictly better.

- As a consequence of the developed theoretical understanding, many new results and computational algorithms for different problems in control theory are presented: stability analysis of a class of differential equations, estimates for the region of attraction of Lyapunov functions, robust bifurcation analysis, etc.

In Appendix A we summarize, for the convenience of the reader, some background material in abstract algebra.

Chapter 2

Cone invariant LMIs

In this chapter, an exact solution for a special class of cone-preserving linear matrix inequalities (LMIs) is developed. By using a generalized version of the classical Perron–Frobenius theorem, the optimal value is shown to be equal to the spectral radius of an associated linear operator. This allows for a much more efficient computation of the optimal solution, using for instance power iteration-type algorithms. This particular LMI class appears in the computation of upper bounds for some generalizations of the structured singular value μ (spherical μ), and in a class of rank minimization problems previously studied. Examples and comparisons with existing techniques are provided.

2.1 Introduction

In the last few years, Linear Matrix Inequalities (LMIs, see [17, 91] for background material) have become very useful tools in control theory. Numerous control–related problems, such as $\mathcal{H}_2$ and $\mathcal{H}_\infty$ analysis and synthesis, μ-analysis, model validation, etc., can be cast and solved in the LMI framework. LMI techniques not only have provided alternative (sometimes simpler) derivations of known results, but also supplied answers for previously unsolved problems.

LMIs are convex optimization problems, that can be solved efficiently in polynomial time. The most effective computational approaches use projective or interior-point methods [67] to compute the optimal solutions.

However, for certain problems, the LMI formulation is not necessarily the most computationally efficient. A typical example of this is the computation of solutions of Riccati inequalities, appearing in $\mathcal{H}_\infty$ control. For these problems, under appropriate regularity hypotheses, the feasibility of the Riccati matrix inequality implies the solvability of the algebraic Riccati equation [34]. In this case, it is not necessary to solve LMIs, but instead just solve Riccati equations, at a lower computational cost. Similarly, the results in this chapter show that for a certain class of LMIs, the optimal solution can be computed by alternative, faster methods than general purpose LMI solvers.

An outline of the material in this chapter follows. In Section 2.2 the notation and some auxiliary facts used later are presented. In Section 2.3 a class of cone-preserving LMIs is defined, and a finite dimensional generalization of the Perron–Frobenius theorem on nonnegative matrices [9] is used to characterize the optimal solution. A brief discussion on computational approaches to the effective calculation of the solution is presented. The application of the results to the computation of the upper bound for the spherical μ problem and to a particular class of rank minimization problems follows. In the following section, some additional comments on the irreducibility conditions are made, a procedure for computing suboptimal solutions of other (non cone-preserving) classes of LMIs is outlined, and finally, some numerical examples are presented.

2.2 Preliminaries

The notation is standard. If M is a matrix, then M^T, M^* denote the transpose and conjugate transpose matrices, respectively. The identity matrix or operator will be denoted by I. A hermitian matrix $M = M^* \in \mathbb{C}^{n\times n}$ is said to be positive (semi) definite if $x^*Mx > 0 (\geq 0)$ for all nonzero $x \in \mathbb{C}^n$. The spectral radius of a finite dimensional linear operator $\mathcal{L}$ is the nonnegative real number $\rho(\mathcal{L}) = \max\{|\lambda| : \mathcal{L}(x) = \lambda x, x \neq 0\}$. The *adjoint* $\mathcal{L}^*$ of a linear operator $\mathcal{L}$ is the unique linear operator that satisfies $\langle x, \mathcal{L}(y)\rangle = \langle \mathcal{L}^*(x), y\rangle$, for all x and y, where $\langle\cdot,\cdot\rangle$ denotes an inner product. The Hadamard (or Schur) element-wise product of two matrices $A = [a_{ij}]$ and $B = [b_{ij}]$ of the same dimensions is defined as $A \circ B \equiv [a_{ij}b_{ij}]$. An important property of this product is the following:

Theorem 2.1 (Schur product theorem, [44]) *If A and B are positive semidefinite matrices, then $A \circ B$ is also positive semidefinite. Moreover, if both A and B are positive definite, so is $A \circ B$.*

A set $S \subseteq \mathbb{R}^n$ is a said to be a *cone* if $\lambda \geq 0, x \in S \Rightarrow \lambda x \in S$. A set S is *convex* if $x_1, x_2 \in S$ implies $\lambda x_1 + (1-\lambda)x_2 \in S$ for all $0 \leq \lambda \leq 1$. The *dual* of a set S is $S^* = \{y \in \mathbb{R}^n : x \in S \Rightarrow \langle x, y \rangle \geq 0\}$. A cone $\mathcal{K}$ is *pointed* if $\mathcal{K} \cap (-\mathcal{K}) = \{0\}$, and *solid* if the interior of $\mathcal{K}$ is not empty. A cone that is convex, closed, pointed and solid is called a *proper* cone. The dual set of a proper cone is also a proper cone, called the *dual cone*. An element x is in the interior of the cone K if and only if $\langle x, y \rangle > 0, \ \forall y \in K^*, y \neq 0$. A proper cone induces a *partial order* in the space, via $x \preceq y$ if and only if $y - x \in \mathcal{K}$. We also use $x \prec y$ if $y - x$ is in the interior of $\mathcal{K}$. Important examples of proper cones are the nonnegative orthant, given by $\{x \in \mathbb{R}^n, x_i \geq 0\}$, and the set of symmetric positive semidefinite matrices.

A linear matrix inequality (LMI, [17]) is defined as

$$F(x) \triangleq F_0 + \sum_{i=1}^{m} x_i F_i > 0,$$

where $x \in \mathbb{R}^m$ is the variable and $F_i \in \mathbb{R}^{n \times n}$ are given symmetric matrices. The problem is to determine if there exists a vector x, that satisfies the matrix inequality. Note that this can be interpreted as a condition on the nonempty intersection of the set given by the affine function $F(x)$ and the self-dual cone of positive definite matrices. A GEVP (generalized eigenvalue problem) takes the form

$$\min\{\lambda : \lambda B(x) - A(x) > 0, B(x) > 0, C(x) > 0\}$$

where A, B and C are symmetric matrices that depend affinely on x. This is a quasiconvex optimization problem, i.e., for fixed λ, the feasible set is convex.

2.3 Problem statement and solution

A straightforward generalization of LMIs can be done by extending matrix inequalities to order inequalities for linear operators. This general abstract setting will prove to be more adequate for our purposes. The main reason why we deal with operators and not directly with their matrix representations is because the operators *act themselves on matrices* (the variables of our LMIs).

The structure of the problems we are interested in is the following:

$$\mathcal{L}(D) \prec \gamma^2 D, \quad D \succ 0 \tag{2.1}$$

where $\mathcal{L}(D)$ is a linear operator that preserves the proper cone $\mathcal{K}$, and the inequalities are to be interpreted in the sense of the partial order induced by the same cone $\mathcal{K}$. In mathematical terms, the cone-preserving assumption on $\mathcal{L}$ can be written as

$$D \in \mathcal{K} \Rightarrow \mathcal{L}(D) \in \mathcal{K}.$$

More specifically, we want to solve for the minimum value of γ, such that the generalized LMI (2.1) is feasible, i.e., the GEVP-like problem

$$\gamma_0 \stackrel{\triangle}{=} \inf\{\gamma \quad | \quad L(D) \prec \gamma^2 D, \quad D \succ 0\}. \tag{2.2}$$

The cone-preserving assumption on $\mathcal{L}$ is fundamental, since these operators have remarkable spectral properties. The most basic instance of this class of operators is the set of nonnegative matrices (i.e., real matrices with nonnegative elements). In this case, the cone $\mathcal{K}$ is the nonnegative orthant and therefore the nonnegative matrix $\mathcal{L}$ leaves $\mathcal{K}$ invariant. This is exactly the setup of the classical Perron–Frobenius theorem [44] that assures, among other things, the existence of a componentwise nonnegative eigenvector. The Perron-Frobenius theory has been extended considerably, with some generalizations to general Banach spaces (due to Krein and Rutman [55]). We are interested here in a particular finite dimensional version.

Theorem 2.2 ([9]) *Assume that the linear operator $\mathcal{L} : \mathbb{R}^n \to \mathbb{R}^n$ maps the proper cone $\mathcal{K}$ into itself. Then*

1. *$\rho(\mathcal{L})$ is an eigenvalue.*
2. *$\mathcal{K}$ contains an eigenvector of $\mathcal{L}$ corresponding to $\rho(\mathcal{L})$.*
3. *$\mathcal{K}^*$ contains an eigenvector of $\mathcal{L}^*$ corresponding to $\rho(\mathcal{L})$.*

There are several proofs of this theorem in the literature. Some use Brouwer's fixed point theorem (as in the infinite dimensional case), or properties of the Jordan canonical form (Birkhoff's proof, [10]).

In order to present the main theorem, we will have to introduce certain technical concepts, to deal with the subtleties of strict vs. nonstrict order inequalities. In particular, the concept of *irreducibility* of cone-preserving operators [9]. The original definition of irreducibility is in terms of invariant *faces*, but we will use an equivalent one, more suited to our purposes.

Definition 2.1 *A $\mathcal{K}$-cone-preserving operator $\mathcal{L}$ is* $\mathcal{K}$-irreducible *if no eigenvector of $\mathcal{L}$ lies on the boundary of the cone $\mathcal{K}$.*

The following lemma establishes a link between the irreducibility of an operator and its adjoint.

Lemma 2.3 *A $\mathcal{K}$-cone-preserving operator $\mathcal{L}$ is $\mathcal{K}$-irreducible if and only if $\mathcal{L}^*$ is $\mathcal{K}^*$-irreducible.*

2.3.1 Optimal solution

The following theorem provides a characterization of the optimal solution of the generalized eigenvalue problem (2.2).

Theorem 2.4 *Assume the operator $\mathcal{L}$ is cone-preserving. Then, the optimal solution of (2.2) has*

$$\gamma_0^2 = \rho(\mathcal{L}). \tag{2.3}$$

Furthermore, if $\gamma^2 > \gamma_0^2$, then it is always possible to find arbitrary solutions for (2.1).

Proof: Since $\mathcal{L}$ preserves the cone $\mathcal{K}$, we can apply Theorem 2.2. Let $Y \in \mathcal{K}^*$ be the eigenvector of $\mathcal{L}^*$ associated with the eigenvalue $\rho(\mathcal{L})$. Then, we can write

$$\begin{aligned}
\mathcal{L}(D) &\prec \gamma^2 D \\
\Rightarrow \langle \mathcal{L}(D), Y\rangle &< \gamma^2 \langle D, Y\rangle \\
\Rightarrow \langle D, \mathcal{L}^*(Y)\rangle &< \gamma^2 \langle D, Y\rangle \\
\Rightarrow \langle D, \rho(\mathcal{L}) Y\rangle &< \gamma^2 \langle D, Y\rangle \\
\Rightarrow \rho(\mathcal{L})\langle D, Y\rangle &< \gamma^2 \langle D, Y\rangle \\
\Rightarrow \rho(\mathcal{L}) &< \gamma^2
\end{aligned}$$

Therefore, γ^2 has to be strictly greater than the spectral radius of $\mathcal{L}$, for (2.2) to hold.

Furthermore, it is possible to get arbitrary solutions of the inequality. Just let P be any element in the interior of the cone $\mathcal{K}$, and consider the equation

$$\gamma^2 D - \mathcal{L}(D) = P$$

For fixed $\gamma^2 > \rho(\mathcal{L})$, this is a consistent system of linear equations. We only have to prove that the solution indeed satisfies $D \succ 0$. To show this, define the convergent linear iteration

$$D_{k+1} = (\mathcal{L}(D_k) + P)/\gamma^2,$$

with $D_0 = 0$. Then, since $\mathcal{L}$ is cone preserving and $P \succ 0$, the solution $D = \lim_{k\to\infty} D_k$ satisfies $D \succ 0$. □

For the nonstrict case, i.e.,

$$\inf\{\gamma \quad | \quad L(D) \preceq \gamma^2 D, \quad D \succeq 0, \quad D \neq 0\}. \tag{2.4}$$

under irreducibility assumptions, we have similarly the following theorem.

Theorem 2.5 *Assume the operator $\mathcal{L}$ is cone-preserving and irreducible. Then, the optimal solution of (2.4) is achieved, and has the value*

$$\gamma_0^2 = \rho(\mathcal{L}). \tag{2.5}$$

Proof: The proof is very similar to the previous one. For the first part, the condition $Y \succ 0$ is guaranteed by the irreducibility of $\mathcal{L}$. For the second part, the optimal D can be taken to be equal to the eigenvector of $\mathcal{L}$ associated with the spectral radius. □

The results of Theorem 2.5 above also hold without the irreducibility assumption on $\mathcal{L}$. The proof uses a continuity argument, applying the theorem to the operator $\mathcal{L} + \epsilon\mathcal{P}$, with $\mathcal{P}$ a $\mathcal{K}$-positive operator[1], i.e., one that satisfies $\mathcal{P}(\mathcal{K} - \{0\}) \subseteq \operatorname{int} \mathcal{K}$. In this case, it is easy to show that $\mathcal{L} + \epsilon\mathcal{P}$ is $\mathcal{K}$-irreducible (since $\mathcal{P}$ is). Then we just take the limit as $\epsilon \to 0$, and use continuity of the spectral radius.

2.3.2 Computation

In the previous subsection a characterization of the optimal value as the spectral radius of an operator was provided. Here we describe some approaches to the problem of effectively computing the value of γ_0.

The most straightforward way (although not the most efficient), is to compute a matrix representation of the operator, and use a general purpose algorithm to compute its eigenvalues. This is clearly not very convenient for large scale problems, where Lanczos or Arnoldi methods are usually the best choice, especially if we are interested only in a few eigenvalues/eigenvectors (as in the present case).

The use of a matrix representation also allows for "squaring"-type methods, where a sequence of matrices A^{2^k} is used. This can be computed using the iteration $A_{k+1} = A_k^2$, with $A_0 = A$ and a suitable normalization scheme at each step. The

[1] Examples of positive operators for the nonnegative orthant and the positive semidefinite cone are the matrix with all its elements equal to one, and the operator $\mathcal{P}(A) = \mathrm{trace}(A)I$, respectively.

effect of the squaring procedure is a separation of the eigenvalues depending on their absolute value (since $\rho(A^2) = \rho^2(A)$).

Under a mild hypothesis ($\mathcal{K}$-primitivity, a subset of $\mathcal{K}$-irreducibility), it is possible to use power iteration-type methods to compute the spectral radius. Primitivity is equivalent to requiring $\rho(\mathcal{L})$ to be strictly greater than the magnitude of any other eigenvalue [9]. It is always possible to obtain a primitive operator by small perturbations of a non primitive one.

In this case, the simple iteration

$$D_{k+1} = \mathcal{L}(D_k)/\|\mathcal{L}(D_k)\|$$

is guaranteed to converge to the eigenvector associated with the spectral radius (and its norm to the optimal value), for every initial value $D_0 \succ 0$. Note also that in the primitive case the squaring procedure describe above result in a very efficient and compact algorithm, since in this case A_k tends to a rank one matrix, from where the spectral radius can be obtained immediately.

It should also be remarked that this power iteration approach to solve a particular type of LMIs has no relationship with the power-type algorithms usually employed in the computation of μ lower bounds.

2.3.3 Applications

Lyapunov inequalities

A simple example, presented here mainly as an illustration of the results, is given by the discrete time Lyapunov inequality, also known as the *Stein* inequality. This is the LMI used to check stability of discrete time linear systems.

It takes the form

$$M^* X M - X < 0, \qquad X > 0, \tag{2.6}$$

and it clearly has the required structure. Using the theory above, we obtain an alternative proof of the well-known result that says that the LMI (2.6) is feasible if and only if the spectral radius of M is less than one.

This example also shows an important point: even if the LMI we are directly interested in does *not* have the cone-invariance property, if may be possible to apply the preceding theory to an equivalent problem. As an illustration, consider for example the continuous time Lyapunov LMI. It is well known that it can be converted into the Stein equation, by the following transformations ($\beta > 0$ is not an eigenvalue of A).

$$
\begin{aligned}
& A^*P + PA < 0 \\
\iff \quad & (A^* + \beta I)P(A + \beta I) - (A^* - \beta I)P(A - \beta I) < 0 \\
\iff \quad & (A^* - \beta I)^{-1}(A^* + \beta I)P(A + \beta I)(A - \beta I)^{-1} - P < 0
\end{aligned}
$$

This is equivalent to defining $M = (A + \beta I)(A - \beta I)^{-1}$, the usual *bilinear* transformation between continuous and discrete domains, and checking for discrete stability.

It is also possible to study Riccati inequalities under a similar framework, at least in the semidefinite case. The theory in this case requires some extensions of the Perron–Frobenius setting to *nonlinear* operators (available in the literature). This approach is not pursued further here.

Spherical μ upper bound LMI

It is possible to directly apply the results developed above to the computation of the LMI upper bound for the generalizations of μ known as spherical μ [52]. In this problem, Frobenius-like constraints are put on the uncertainty block Δ, as opposed to induced norm constraints on each block. For simplicity, we will only refer only to the scalar case.

More concretely, we want to obtain conditions that guarantee the well-posedness of the feedback interconnection of a constant matrix M and a diagonal uncertainty block $\Delta = \operatorname{diag}\{\delta_1, \delta_2, \ldots, \delta_n\}$, $\delta_i \in \mathbb{C}$, that satisfies $\sum_{i=1}^n |\delta_i|^2 \leq 1$. As in the standard case [69], necessary and sufficient conditions are computationally hard, and therefore approximation methods should be used instead. Sufficient conditions (given by μ upper bounds) are usually computed using LMI methods.

In this case, the underlying linear vector space is now the set of hermitian matrices, and $\mathcal{K}$ will be the self-dual cone of positive semidefinite matrices. Note that all the "vectors" in the preceding abstract setting are now matrices.

In the spherical μ upper bound case, the LMI to be solved is very similar to the standard μ LMI upper bound (2.13).

$$M^*(P \circ D)M - \gamma^2 D < 0, \qquad D > 0, \tag{2.7}$$

where P is a positive definite matrix (equal to the identity, in the restricted case presented above).

Lemma 2.6 *Let P be positive semidefinite. Then, the operator $\mathcal{L}(D) = M^*(P \circ D)M$ preserves the cone $\mathcal{K}$ of positive semidefinite matrices.*

Proof: $\mathcal{L}$ is the composition of the two operators $\mathcal{L}_1(D) = P \circ D$ and $\mathcal{L}_2(D) = M^*DM$. The first one is cone-preserving by Theorem 2.1. The second one has the same property, since $x^*M^*DMx < 0$ implies $y^*Dy < 0$, with $y = Mx$. $\square$

In the particular case where P is the identity, we obtain the following corollary:

Corollary 1 *Let γ_0 be the optimal solution of the GEVP:*

$$\gamma_0 \triangleq \inf\{\gamma \quad | \quad M^*(I \circ D)M - \gamma^2 D < 0, \quad D > 0\}. \tag{2.8}$$

Then,

$$\gamma_0^2 = \rho(M^T \circ M^*).$$

Proof: A matrix representation of the nontrivial part (i.e., after removing the trivial kernel) of the operator $M^*(I \circ D)M$ can easily be obtained by elementary algebra (or, somewhat easier, using Kronecker products), to show the equality

$$\mathrm{diag}(M^*(I \circ D)M) = (M^T \circ M^*)\mathrm{diag}(D),$$

where $\mathrm{diag}(D)$ is the operator that maps the diagonal elements of a matrix into a vector. $\square$

The corollary shows that both the optimal value of γ and D can be obtained by just solving one eigenvalue problem, with dimensions equal to those of M. Note that the matrix $M^T \circ M^*$ is simply the matrix whose elements are the square of the absolute value of the elements of M.

Rank minimization problem

In [63, 62], Mesbahi and Papavassilopoulos show that for certain special cases, the rank minimization problem (which is computationally hard in general) can be reduced to a semidefinite program (an LMI). The structure of their problem can be shown to be basically equivalent to the one presented here. Theorem 2.4 above can be used to show that it is *not even necessary* to solve the resulting LMI, just solving a linear system (using direct or iterative techniques, for example) will provide the optimal solution. As in the previous subsection, the cone $\mathcal{K}$ in this problem is the self-dual cone of positive semidefinite matrices.

The problem considered in [63, 62] is stated as:

$$\begin{aligned} & \min \operatorname{rank} X \\ \text{subject to: } Q + \mathcal{M}(X) &\succeq 0 \\ X &\succeq 0, \end{aligned}$$

where Q is a negative semidefinite matrix and $\mathcal{M}$ is a linear map of the structure (called "type $\mathcal{Z}$")

$$\mathcal{M}(X) = X - \sum_{i=1}^{k} M_i X M_i'.$$

Under these hypotheses, it is possible to prove [63, 62] that a solution can be obtained by solving the associated LMI:

$$\begin{aligned} & \min \operatorname{trace} X \\ \text{subject to: } Q + \mathcal{M}(X) &\succeq 0 \\ X &\succeq 0. \end{aligned}$$

Let $P = -Q$ (therefore P is positive semidefinite, i.e., $P \succeq 0$), and $P \neq 0$, to avoid the trivial solution $X = 0$. Defining $\mathcal{L}(X) := X - \mathcal{M}(X) = \sum_{i=1}^{k} M_i X M_i'$, we obtain the equivalent formulation:

$$\min \operatorname{trace} X \tag{2.9}$$

$$\text{subject to: } X - \mathcal{L}(X) \succeq P \tag{2.10}$$

$$X \succeq 0. \tag{2.11}$$

It is clear from its definition (and the proof of Lemma 2.6) that $\mathcal{L}(X)$ preserves the cone of semidefinite positive matrices.

Theorem 2.7 *If the LMIs (2.10-2.11) are feasible, then $\rho(\mathcal{L}) \leq 1$.*

Proof: The proof is essentially similar to that of Theorem 2.4, taking $\gamma = 1$ and using the condition $P \succeq 0$. □

In the case $\rho(\mathcal{L}) < 1$, then the constraint (2.11) is not binding at optimality, and the solution can be obtained by solving the consistent linear system

$$X - \mathcal{L}(X) = P, \tag{2.12}$$

as the following theorem shows.

Theorem 2.8 *Let X_e be the solution of (2.12). Then, X_e is an optimal solution of the LMI (2.9-2.11).*

Proof: Let's show first that $X_e \succeq 0$. As in the proof of Theorem 2.4, consider the sequence X_i, with $X_0 = 0$ and $X_{i+1} = \mathcal{L}(X_i) + P$. All the elements in the sequence belong to the cone $\mathcal{K}$. The sequence converges (due to the spectral radius condition), and $\lim_{i\to\infty} X_i = X_e$. Closedness of $\mathcal{K}$ implies $X_e \in \mathcal{K}$.

Let X be any feasible solution of the LMI. Therefore, we have:

$$X_e - \mathcal{L}(X_e) = P,$$

$$X - \mathcal{L}(X) \ \succeq \ P.$$

Subtracting, we obtain

$$X - X_e \succeq \mathcal{L}(X - X_e),$$

and by repeated application of $\mathcal{L}$ to both sides of the inequality

$$X - X_e \succeq \mathcal{L}^k(X - X_e), \qquad \forall k \geq 1.$$

Since $\rho(\mathcal{L}) < 1$, the right-hand side of the preceding inequality vanishes as $k \to \infty$. This implies $X - X_e \succeq 0$, and therefore $\mathrm{trace}(X) \geq \mathrm{trace}(X_e)$. $\square$

Note: The case $\rho(\mathcal{L}) = 1$ can also be analyzed, via perturbation arguments.

2.4 Additional comments and examples

In this section we give some examples on the irreducibility notion mentioned above, and mention some of the applications of the results in the computation of approximate solutions for other LMIs that are not necessarily cone-preserving.

2.4.1 More on irreducibility

To explain a little bit more of the irreducibility concept introduced above, we will present a couple of examples. In what follows, we will consider the GEVP problem (2.8).

For the first case, take M to be

$$M = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$$

According to Corollary 1, the optimal solution γ of the GEVP (2.7) (with $P = I$) is given by the spectral radius of $M^* \circ M^T$, which is $\gamma_0 = 1$. In this case, the eigenvector (really a matrix) associated with this eigenvalue is

$$X = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}.$$

Clearly, this matrix is in the boundary of the cone of positive semidefinite matrices. Therefore, the operator associated with this problem is *not irreducible.* The optimal value of γ cannot be achieved by any positive definite D, although we can approximate the solution as closely as we want, as explained in the proof of Theorem 2.4.

For an example of an irreducible operator, although not a primitive one, consider

$$M = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

The eigenvalues of the associated operator are 0, 1 and -1, and the eigenvector corresponding to the spectral radius is the identity matrix, which lies in the interior of the cone of positive semidefinite matrices. Therefore, it is irreducible. However, it is not primitive, and therefore it is not possible to directly apply power iteration to compute the spectral radius.

2.4.2 Suboptimal solutions of LMIs

The cone-preserving requirement for the LMI is a strict one, since it implies that the optimal solution actually achieves an equality in the limit. Many of the common LMIs appearing in control problems do not necessarily give an equality at optimality. A typical example is the standard μ LMI, where the decision variable D is not full, but structured. In other words, the partial order induced by the inequality is not the same as the one induced by the variable D.

However, the methodology presented above can be used as a fast method for computing suboptimal feasible solutions for certain problems. These suboptimal values can often be used as starting points for more general LMI solvers.

For example, for the standard μ upper bound LMI

$$M^*(I \circ D)M - \gamma^2 (I \circ D) < 0, \qquad D > 0, \tag{2.13}$$

it is possible to compute an approximate solution by using the following procedure:

1. Compute the exact solution γ_1^2, D_1 of the spherical μ LMI (2.7).

2. Compute the smallest η that satisfies

$$D_1 \leq \eta^2 (I \circ D_1). \tag{2.14}$$

   This is a generalized eigenvalue problem, that can be easily reduced to the computation of the maximum eigenvalue of a hermitian matrix. It is also possible to show, since D is positive definite, that $\eta^2 \leq n$ [52].

3. Therefore, a suboptimal solution of the LMI is given by $I \circ D_1$, and the optimal value is $\gamma = \eta\gamma_1 \leq \sqrt{n}\gamma_1$.

Effectively, we have

$$M^*(I \circ D_1)M \leq \gamma_1^2 D_1 \leq \eta^2 \gamma_1^2 (I \circ D_1).$$

It is possible to (almost) achieve the worst case difference between the optimal solution and the approximate one ($\sqrt{n}$). For example, for the matrix

$$M = \begin{bmatrix} 1 & \varepsilon & \cdots & \varepsilon \\ \varepsilon & \varepsilon & \cdots & \varepsilon \\ \vdots & \vdots & \ddots & \vdots \\ \varepsilon & \varepsilon & \cdots & \varepsilon \end{bmatrix},$$

with ε small, the optimal value of the LMI (2.13) is $1 + O(\varepsilon)$, but the fast upper bound is approximately $\sqrt{n}$.

Another available procedure for computing fast solutions of the μ LMI is the one due to Osborne [68]. A preliminary comparison made with random, normally distributed matrices gives a slight advantage to the Osborne procedure. However, the algorithm proposed can give better upper bounds (the opposite is also possible), as the following example shows. For the matrix

$$M = \begin{bmatrix} 0 & -9 & -4 \\ 2 & 6 & 6 \\ -3 & -1 & 6 \end{bmatrix}$$

the μ upper bound computed by Osborne preconditioning is 10.321, and the bound of the proposed procedure is 9.69 (the value of the LMI upper bound is 9.6604, and is in fact equal to μ since there are three blocks).

2.4.3 Examples

As a simple example of the computational advantages of the proposed formulation, we will compare the effort required to compute solutions of the spherical μ LMI upper bound(2.7), for a given problem.

We take M to be a 16×16 complex matrix, randomly generated. The computation of the optimal value of the LMI (2.7) with a general purpose LMI solver for MATLAB [35] and a tolerance set to 10^{-4} requires (on a Sun Ultra 1 140) approximately 160 seconds. By using the procedure presented here, either by power iteration or explicitly computing the eigenvalues, the result can be obtained in less than one second.

Chapter 3

Efficient solutions for KYP-based LMIs

The semidefinite programs appearing in linear robustness analysis problems usually have a very particular structure. This special form is a consequence of both the linearity and the time invariance of the underlying system. In this chapter, we will see how this special structure can be exploited in the formulation of efficient algorithms.

The KYP lemma (Kalman-Yakubovich-Popov [93], see [77] for an elementary proof) establishes the equivalence between a frequency domain inequality and the feasibility of a particular kind of LMI (linear matrix inequality). It is an important generalization of classical linear control results, such as the bounded real and positive real lemma. It is also a fundamental tool in the practical application of the IQC (integral quadratic constraints) framework [61] to the analysis of uncertain systems. The theorem replaces an infinite family of LMIs, parameterized by ω, by a finite dimensional problem. This is extremely useful from a practical viewpoint, since it allows for the use of standard finite dimensional LMI solvers.

However, in the case of systems with large state dimension n, the KYP approach is not very efficient, since the matrix variable P appearing in the LMI (3.2) has $(n^2+n)/2$ components, and therefore the computational requirements are quite big, even for medium sized problems. For example, for a problem with a plant having 100 states (which is not uncommon in certain applications), the resulting problem has more than 5000 variables, beyond the limits of what can be currently solved with reasonable time and space requirements using general-purpose LMI software.

In this chapter, we present an efficient algorithm for the solution of this type of inequalities. The approach is an *outer approximation* method [72], and is based on the algorithms used in the computation of $\mathcal{H}_\infty$ system norms. The idea is to impose the frequency domain inequality (3.1) only at a discrete number of frequencies. These frequencies are then updated by a mechanism reminiscent of those used in $\mathcal{H}_\infty$ norm computation.

Previous related work includes of course the literature on the computation of $\mathcal{H}_\infty$ system norms. In particular, references [16, 20, 15] developed quadratically convergent algorithms, based explicitly on the Hamiltonian approach. Also, a somewhat related approach in [60] implements a cutting-plane based algorithm, where linear constraints are imposed on the optimization variables.

3.1 The KYP lemma

In this section we review some basic linear algebra facts, and also present a version of the KYP lemma. The notation is standard.

A $2n \times 2n$ real matrix is said to be *Hamiltonian* (or *infinitesimally symplectic*) if it satisfies $H^T J + JH = 0$, where

$$J \triangleq \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix}.$$

Hamiltonian matrices have a spectrum that is symmetric with respect to the origin. That is, λ is an eigenvalue iff $-\lambda^*$ is. It can be shown that a partitioned matrix

$$H = \begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix}$$

is Hamiltonian if and only if H_{12} and H_{21} are both symmetric and $H_{11}^T + H_{22} = 0$.

A basic fact about determinants of matrices, easy to prove using an Schur-like matrix decomposition, is the following:

Lemma 3.1 *Let Q be a partitioned matrix*

$$Q = \begin{bmatrix} Q_{11} & Q_{12} \\ Q_{21} & Q_{22} \end{bmatrix}$$

with Q_{11} and Q_{22} invertible. Then, we have the identity:

$$\det Q = \det Q_{11} \det(Q_{22} - Q_{21}Q_{11}^{-1}Q_{12}) = \det Q_{22} \det(Q_{11} - Q_{12}Q_{22}^{-1}Q_{21})$$

A fairly general version of the KYP lemma, as presented in [77] is the following:

Theorem 3.2 *Let $A \in \mathbb{R}^{n\times n}, B \in \mathbb{R}^{n\times m}, M = M^T \in \mathbb{R}^{(n+m)\times(n+m)}$, with A having no purely imaginary eigenvalues. Then, the two following statements are equivalent:*

1.

$$F(j\omega) \triangleq \begin{bmatrix} (j\omega I - A)^{-1}B \\ I \end{bmatrix}^* M \begin{bmatrix} (j\omega I - A)^{-1}B \\ I \end{bmatrix} < 0, \quad \forall \omega \in \mathbb{R} \cup \{\infty\} \tag{3.1}$$

2. *There exists a symmetric $n \times n$ matrix P that satisfies*

$$\begin{bmatrix} A^TP + PA & PB \\ B^TP & 0 \end{bmatrix} + M < 0 \tag{3.2}$$

Proof: *We present a proof of (2) $\Rightarrow$ (1), to show the connection with the methods of Chapter 4. The second condition guarantees $M_{22} < 0$, so the case $\omega = \infty$ holds. In what follows, we analyze the case $\omega \neq \infty$.*

An equivalent statement of (3.1) is the implication

$$j\omega x = Ax + Bu \quad \Longrightarrow \quad \begin{bmatrix} x \\ u \end{bmatrix}^* M \begin{bmatrix} x \\ u \end{bmatrix} < 0.$$

Let P be a symmetric matrix. Clearly, a condition that guarantees that the expression above holds is:

$$x^*P(Ax+Bu-j\omega x)+(Ax+Bu-j\omega x)^*Px+\begin{bmatrix} x \\ u \end{bmatrix}^* M \begin{bmatrix} x \\ u \end{bmatrix} < 0,$$

for all $x \in \mathbb{R}^n, u \in \mathbb{R}^m, (x,u) \neq 0$. It can be easily verified that the terms containing ω cancel, and the expression can be rewritten as (3.2).

For a proof of the other direction ($1 \Rightarrow 2$), see [77]. □

In the application of this result to the stability analysis of uncertain systems, the matrix M depends affinely on some parameter vector x. These are the variables of the LMI optimization problem, where we try to minimize some linear function of x over the feasible set (for example, a bound on the $\mathcal{L}_2$-induced norm). In what follows, the dependence on x is usually omitted, for notational reasons.

Here we will deal only with the strict version of the KYP lemma, i.e., with a strict inequality in (3.1), (3.2). The reason is twofold: in the first place, no controllability/stabilizability assumptions are necessary, simplifying the proofs. Secondly, since the resulting LMIs will in general be solved using interior-point methods, the existence of a strictly feasible solution is usually guaranteed.

3.2 The Algorithm

The basic idea is to replace the semi-infinite optimization problem (3.1) by a finite dimensional relaxation. We choose to impose the constraint only at a finite number of frequencies $\omega_k \in \Omega$ (see [50] for a related approach). For a given ω, equation (3.1) is an LMI in M.

A high-level description of the algorithm follows:

Algorithm 1

1. *Initialize the set of frequencies $\Omega \triangleq \{0\}$.*

2. *Solve (3.1) with the current Ω set.*

3. *Find a frequency ω_k where the constraint (3.1) is violated (up to an ϵ tolerance). If no such frequency exists, exit.*

4. *Add that frequency to the set Ω, and go to step 2.*

As we can see, the underlying idea of an outer approximation algorithm is a generalization of a *cutting plane* method [72]. We replace the description of the feasible set by a convenient relaxation. If the resulting solution does not satisfy the original constraints, a cutting plane (in this case, a possibly curved hypersurface) that separates that solution from the true feasible set is added. The process is repeated until the desired tolerance is reached.

As in the case of $\mathcal{H}_\infty$ norm computation [16, 20], the effectiveness of the algorithm hinges on the possibility of detecting in an efficient manner the frequencies at which the inequality is violated. To this end, define the $2n \times 2n$ Hamiltonian matrix:

$$H = \begin{bmatrix} A - BM_{22}^{-1}M_{21} & -BM_{22}^{-1}B^T \\ -M_{11} + M_{12}M_{22}^{-1}M_{21} & -A^T + M_{12}M_{22}^{-1}B^T \end{bmatrix} \tag{3.3}$$

It can be shown (see for example [93]) that the conditions (3.1), (3.2) are satisfied if and only if $M_{22} < 0$ and H has no imaginary eigenvalues. In this case, it is possible to obtain a solution P of the LMI (3.1) by computing a solution of the Riccati equation associated with the Hamiltonian (or a suitable perturbation, if the subspace complementarity condition is not satisfied). If the eigenvalue condition is violated, then there is a relationship between the critical frequencies, as the following theorem shows.

Theorem 3.3 *Assume $M_{22} < 0$. Then, $F(j\omega_0)$ is singular, if and only if $j\omega_0$ is an imaginary eigenvalue of H.*

Proof: *Consider the partitioned matrix*

$$Q \triangleq \left[\begin{array}{cc|c} j\omega I - A & 0 & -B \\ M_{11} & j\omega I + A^T & M_{12} \\ \hline M_{21} & B^T & M_{22} \end{array}\right].$$

The diagonal submatrices are invertible, since A has no imaginary eigenvalues and $M_{22} < 0$. Applying Lemma 3.1, we immediately have the identity

$$\det(j\omega I - H)\det M_{22} = \det(j\omega I + A^T)\det F(j\omega)\det(j\omega I - A)$$

from where the result follows. □

Special cases of this theorem are the ones used in [16] to compute the $\mathcal{H}_\infty$ norm or the minimum dissipation of a transfer function.

Several options are available for the choice of the frequency to add to the set Ω. A particularly good one is to choose ω_k as the frequency at which $F(j\omega)$ is maximally positive (i.e., where its first singular value achieves its maximum over frequency). This can be obtained at a computational cost similar to that of an $\mathcal{H}_\infty$ norm. In the following section we present a convergence argument for the procedure resulting from this choice. A cheaper alternative is to pick a criterion similar to the one used in [20]. Given the imaginary eigenvalues of H, consider the midpoint frequencies, and choose the one where the constraint is most violated. The computational requirements of this step are minimal, compared to the one required to solve the LMIs.

An important difference of the LMI case discussed here with the simpler $\mathcal{H}_\infty$ norm case (where the *only* LMI variable is the KYP one) is that at optimality more than one constraint can be active. In fact, the results in [50] show that at most $n+1$ frequencies are active, where n is the number of IQCs.

In the algorithm as described, no constraint dropping occurs. That is, we keep adding constraints, until convergence. Since we know *a priori* a bound on the number of active constraints, dropping old, non-binding constraints seems a natural idea.

The distinctive feature of the algorithm is that the KYP variable P, never appears explicitly in the procedure. Nevertheless, as mentioned before, it is possible to compute its value after the problem is solved, at a computational cost similar to solving a Riccati equation.

A somewhat related approach is used in [60], where the eigenvectors of the Hamiltonian are used to construct linear constraints for the elements of M. In our approach, the constraints are matrix valued (not linear) and we do not impose the restrictions directly at the critical frequencies, but at other points where they are more violated. This way, convergence should be improved (in the $\mathcal{H}_\infty$ case, it is even quadratic).

3.2.1 Convergence

It is possible to prove convergence of the first version of the algorithm. This corresponds to the choice of ω_k as the point at which the frequency domain inequality is maximally violated. In fact, for this variation we can apply the results on the convergence of more abstract version of the outer approximation method (Conceptual Algorithm 3.5.19 in [72]).

It is possible to show (see [72]) that if the algorithm produces a infinite sequence of solutions, then any accumulation point of this sequence is a global solution of the original problem. The infinite set of frequency constraints can be "compactified" either by considering the extended real line or by a standard bilinear transformation.

Currently we do not have explicit, nonconservative expressions for the convergence rate. This seems to be a general feature of the outer approximation class of algorithms, since even for cutting plane methods the known theoretical bounds are usually extremely conservative, when compared to the actual performance.

3.3 Using the dual

A not so convenient feature of the presented approach is that a new *constraint* is added at each iteration. This implies that the previous solution will not be primal feasible, forcing a restart of the optimization, unless an infeasible start method is used.

This can be solved by focusing instead on the dual optimization problem, as is well known from the linear programming literature, for instance. In this case, new

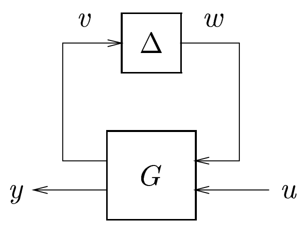

Figure 3.1: Standard block diagram.

variables are added to the problem at each iteration. Note that this can also be interpreted as having a dual feasible starting point, which is useful in case we are using a primal-dual LMI solver (such as SDPSOL [18]).

For the frequency domain inequalities arising from IQC optimization, the dual problem has been extensively analyzed in [50]. It has been shown there that upper bounds, or even the optimal value, of the quantities of interest (for example, $\mathcal{L}_2$-induced norms) can be obtained from a finite number of frequencies. However, no procedure to compute or approximate these frequencies was available, other than a standard gridding.

The algorithm presented here provides an explicit methodology for the update of the frequencies. This way, better bounds can be obtained in an iterative fashion, with an arbitrarily small error.

3.4 Example

In this section two examples of the application of the proposed algorithm are presented. The first one is very simple, and mainly for illustration purposes. In the second one, the performance is compared with a standard LMI solver for a medium scale problem. Both examples are solved using MATLAB's LMI toolbox, with the default options.

Example 3.1 *Consider the standard block diagram in Fig. 3.1. We will use the proposed algorithm to compute the worst case $\mathcal{L}_2$ induced norm between u and y, for*

Frequencies	Obj. Value	Imag. Eigs. of H
0	2.0012	0.0353 1.9984
0 1.0169	2.7282	1.0171 1.2073
0 1.0169 1.1122	2.7474	-

Table 3.1: Numerical values for Example 3.1.

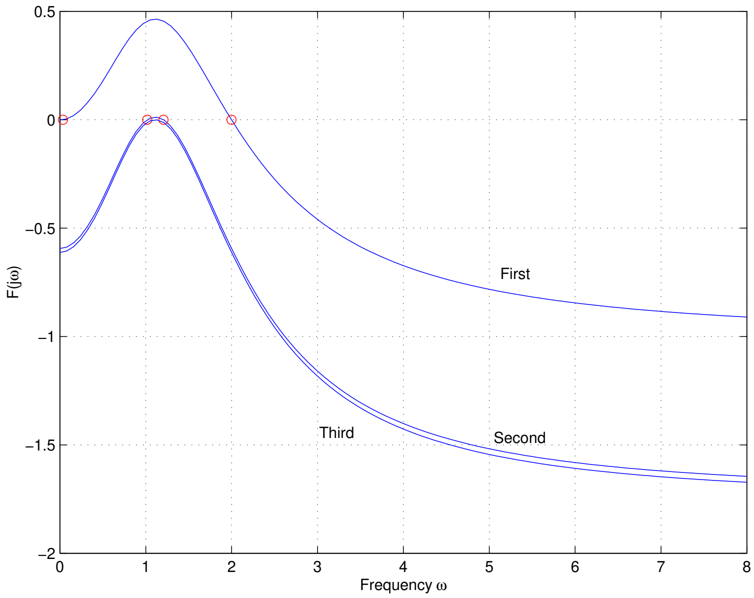

Figure 3.2: Frequency domain plots corresponding to Example 3.1.

the plant given by

$$G = \begin{bmatrix} \frac{s+1}{s^2+2s+2} & 1 \\ 1 & 0 \end{bmatrix}.$$

The Δ block is an uncertain contractive LTV operator, and therefore satisfies the IQC given by

$$\Pi(j\omega) = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

The results of the sequence of subproblems are shown in Table 3.1 and Fig. 3.2.

As we can see, on the third and last iteration we obtain a value of the parameters that makes the frequency domain inequality to be satisfied. That makes possible, if

Frequencies	Obj. Value	Time (sec.)
0	64.33	14.8
0 2.9	77.3456	30.29
0 2.9 2.7353	77.5511	54.87

Table 3.2: Numerical values for Example 3.2.

desired, to recover the value of the optimal KYP variable P, by solving a Riccati equation. In this case, we obtain

$$P = \begin{bmatrix} 3.4849 & 0.6674 \\ 0.6674 & 0.6644 \end{bmatrix}.$$

This is within numerical error of the solution obtained by directly solving the LMI (3.1).

In the next example, we show the numerical advantages of using the outlined procedure for solving the LMIs appearing in analysis problems with systems of large state dimension.

Example 3.2 *The system is again in the standard form of Fig. 3.1. The plant G, chosen randomly, has 50 states, and the signals u, y, v, w are vector-valued, with each having 10 components. The uncertainty Δ corresponds to a diagonal gain bounded LTV operator, and therefore there are 10 IQCs associated with it.*

For this example, we have chosen as the new frequency to be added to the set Ω the one at which the constraints are maximally violated, as explained before. Though more expensive, it seems to have faster convergence properties. A straightforward solution of the LMIs with the KYP variable takes 996 sec., on a Sun Ultra 10/300Mhz. On the same hardware, the total time required by the presented procedure is less than 120 sec. Note that here we are solving the primal problem, and the MATLAB LMI toolbox uses a projective algorithm, and does not use any dual information. This implies that each subproblem is solved from scratch. The time spent in computing the maximum over frequencies (analog to an $\mathcal{H}_\infty$ norm) is negligible.

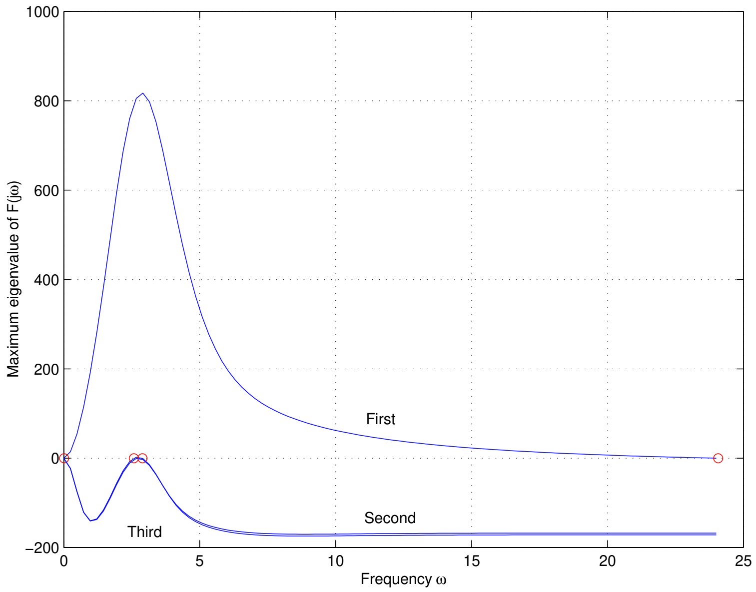

Figure 3.3: Frequency domain plots corresponding to Example 3.2.

Note that in this last example, as opposed to the previous one, more than one constraint is active at optimality. A result from [50] is that at most $n+1$ frequencies are active, so this is consistent with the expected behavior.

Finally, we remark that even though we are currently using a relatively inefficient implementation (since we are not using the information obtained in earlier stages in the solution of the subproblems), the algorithm still outperforms the standard approach.

Chapter 4

Sums of squares and algebraic geometry

This chapter presents our approach to the formulation of stronger convex conditions for a large class of optimization and systems and control problems. The fundamental feature is the computational tractability of the sum of squares decomposition for multivariable polynomials. As shown below, the problem can be solved via semidefinite programming methods.

Complementing this formulation with results in semialgebraic geometry (the Positivstellensatz), a whole class of convex approximations for optimization problems is developed. In subsequent chapters, we specialize the techniques to some specific problems.

4.1 Global nonnegativity

A basic problem that appears in many areas of mathematics is that of checking global nonnegativity of a function of several variables. Concretely, the problem is to give equivalent conditions or a procedure for checking the validity of the proposition

$$F(x_1, \ldots, x_n) \geq 0, \qquad \forall x_1, \ldots, x_n \in \mathbb{R}. \tag{4.1}$$

This is a very important problem, and lots of research efforts have been devoted to it. In order to study the problem from an algorithmic approach, we need to put further restrictions on the class of functions F, since the general question can be shown to be undecidable. To illustrate this, consider Richardson's theorem, as quoted in [71].

Theorem 4.1 *Let $\mathcal{R}$ consist of the class of expressions generated by*

1. *The rational numbers and the two real numbers π and $\ln 2$.*

2. *The variable x.*

3. *The operations of addition, multiplication, and composition.*

4. *The sine, exponential, and absolute value functions.*

If $E \in \mathcal{R}$, the predicate "$E = 0$" is recursively undecidable.

It is clear then that we necessarily need to limit the structure of the possible functions F, while at the same time making the problem general enough to guarantee the applicability of the results. A good compromise is achieved by considering the case of *polynomial functions.*

Definition 4.1 *A* polynomial *f in $x_1, \ldots, x_n$ with coefficients in a field k is a finite linear combination of monomials:*

$$f = \sum_{\alpha} c_\alpha x^\alpha = \sum_{\alpha} c_\alpha x_1^{\alpha_1} \ldots x_n^{\alpha_n}, \qquad c_\alpha \in k, \tag{4.2}$$

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 coefficients in k is denoted $k[x_1, \ldots, x_n]$.

Definition 4.2 *A* form *is a polynomial where all the monomials have the same degree $d := \sum_i \alpha_i$. In this case, the polynomial is* homogeneous *of degree d, since it satisfies $f(\lambda x_1, \ldots, \lambda x_n) = \lambda^d f(x_1, \ldots, x_n)$.*

Many concrete problems, particularly in systems and control, can be reduced to the verification of the global nonnegativity of a polynomial function [13]. Some examples, presented in Chapter 7, are Lyapunov function computation, output feedback stabilization, multidimensional system stability, etc.

As mentioned in Chapter 1, the Tarski-Seidenberg decision procedure [12, 64, 13] provides in this case an explicit algorithm for deciding if (4.1) holds, so we know that the problem is decidable. There are also a few alternative approaches, also based in decision algebra; see [13] for a survey of existing techniques.

It is possible to show that the general problem of testing global positivity of a polynomial function is in fact NP-hard (when the degree is at least four). Therefore, (unless P=NP) *any method* guaranteed to obtain the right answer *in every possible instance* will have unacceptable behavior for a problem with a large number of variables. This is the main drawback of theoretically powerful methodologies such as quantifier elimination [31, 47].

If we want to avoid the inherent complexity problems related with the *exact* solution, the question arises: are there any conditions, that can be tested in polynomial time, to guarantee global positivity of a function? As we will shortly see, one such condition is given by the existence of a sum of squares decomposition.

4.2 Sums of squares

If a polynomial F satisfies (4.1), then an obvious necessary condition is that the degree of the polynomial (or form, in the homogeneous case) be even. A deceptively simple sufficient condition for a real-valued function $F(x)$ to be nonnegative everywhere is given by the existence of a sum of squares decomposition:

$$F(x) = \sum_i f_i^2(x)$$

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

However, the question immediately arises: when is such decomposition possible? Naturally, in order for the problem to make sense, some restriction on the class of functions f_i has to be imposed again. Otherwise, we can always define f_1 to be the square root of F, making the condition both useless and trivial.

For the case of polynomials, this is a well-analyzed problem, first studied by

David Hilbert more than a century ago. In fact, one of the questions 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.

For notational simplicity, we will use the notation *psd* for "positive semidefinite" and *sos* for "sum of squares." Following the notation in references [24, 80], let $P_{n,m}$ be the set of psd 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 himself noted that not every psd polynomial (or form) is sos. A simple, more modern counterexample is the Motzkin form (here, for $n = 3$)

$$M(x, y, z) = x^4y^2 + x^2y^4 + z^6 - 3x^2y^2z^2 \tag{4.3}$$

Positive semidefiniteness can be easily shown using the arithmetic-geometric inequality, and the nonexistence of a sos decomposition follows from standard algebraic manipulations (see [80] for details), or the procedure outlined below (Example 4.5).

Hilbert gave a complete characterization of when these two classes are equivalent. There are three cases for which the equality holds. The first one, is the case of forms in two variables ($n = 2$), which are equivalent by dehomogenization to 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 forms (i.e., $m = 2$) where the sum of squares decomposition follows from the eigenvalue/eigenvector factorization. There is also a surprising third case, where $P_{3,4} = \Sigma_{3,4}$, corresponding to quartic forms in three variables.

The sum of squares decomposition is the underlying machinery in Shor's global bound for polynomial functions [91], as is explicitly mentioned in [83]. It has also been presented as the "Gram matrix" method in [24] and more recently in [74], although no mention to interior point methods is made: the resulting LMIs are solved via decision methods. A related scheme also appears in [41] (note also the important correction in [33]).

The basic idea of the method is the following: express the given polynomial as a quadratic form in some new variables z. These new variables are the original x ones, plus all the monomials of degree less than or equal to $m/2$ given by the different products of the x variables. Therefore, $F(x)$ can be represented as:

$$F(x) = z^T Q z \tag{4.4}$$

where Q is a constant matrix. If in the representation above Q is positive semidefinite, then $F(x)$ is also psd. This is the idea in [14], for example, and it can be shown to be conservative, generally speaking. The main reason is that since the variables z_i *are not independent*, the representation (4.4) *might not be unique*, and Q may be psd for some representations but not for others. Similar issues appear in the analysis of quasi-LPV systems; see [45]. 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 $z_i^2 = z_k z_l$), it is easily shown that there is a linear subspace of matrices Q that satisfy (4.4). 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 psd). This follows from an eigenvalue decomposition of $Q = T^T D T, d_i \geq 0$, which implies the sos $F(x) = \sum_i d_i (Tz)_i^2$. Conversely, if F can indeed be written as the sum of squares of polynomials, then expanding in monomials will provide the representation (4.4).

Example 4.1 *Consider the quartic form in two variables described below, and define* $z_1 := x_1^2, z_2 := x_2^2, z_3 := x_1 x_2$*:*

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

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

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

And therefore we have the sum of squares decomposition:

$$F(x_1, x_2) = \frac{1}{2}((2x_1^2 - 3x_2^2 + x_1x_2)^2 + (x_2^2 + 3x_1x_2)^2).$$

□

Example 4.2 *The following example is from [13, 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 [13], this property is established using decision algebra.

By constructing the Q *matrix as described above, and solving the corresponding LMIs, 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. *In fact, the bound can be shown to be exact, since for example* $P(0, 1, -1) = 1$. □

If the polynomial F is sparse, in the sense that many of the monomials are zero, then it is usually possible to considerably simplify the resulting LMIs. To do this, we can use a result, first formulated in [78], that characterizes the monomials that can appear in a sum of squares representation. Define the *cage* (or *Newton polytope* [87]) of F as the integer lattice points in the convex hull of the degrees α (in 4.2), considered as vectors in $\mathbb{R}^n$. Then, it can be shown that the only monomials x^β that can appear in a sum of squares representation are those such that 2β is in the cage of F.

The most important properties that distinguish the semidefinite programming condition from other approaches to the polynomial nonnegativity problem are its relative *tractability*, and the fact that it can be easily extended to the uncertain case (i.e., when we are looking for a psd F, subject to additional conditions). This last feature will prove to be critical in the application to the theory to many control related problems.

As an example, we can apply the technique to compute global lower bounds for polynomial functions [83]. Since the condition

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

is affine in γ, then it is possible to efficiently compute the maximum value of γ for which this property holds. In some cases, as in Example 4.2 above, the resulting bound is optimal. However, for the reasons mentioned earlier, it is also possible to obtain conservative results.

Example 4.3 *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 (4.3) above. 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 LMIs, 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:

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

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. Fortunately, techniques are available to deal with such cases, as we will shortly see.* $\square$

4.3 The dual problem

It is enlightening to analyze the dual problem, that gives conditions on when $F(x)$ is *not* a sum of squares. Given $F(x)$, consider a representation

$$F(x) = z^TQz = \text{trace } zz^TQ,$$

for all vectors z. Relaxing the rank one constraint on the matrix $W := zz^T$ (now W is only positive semidefinite), it is clear that a sufficient condition for $F(x)$ not to be a sum of squares is the existence of a matrix W satisfying

$$\text{trace } WQ < 0, \qquad W \geq 0.$$

The non uniqueness of Q in the quadratic representation now translates into equality constraints between the elements of W. These equality constraints ensure that products between the newly defined variables that are supposed to be identical actually have the same value.

Example 4.4 *Consider again Example 4.1. In this case, the dual variable is:*

$$W = \begin{bmatrix} w_{11} & w_{12} & w_{13} \\ w_{12} & w_{22} & w_{23} \\ w_{13} & w_{23} & w_{33} \end{bmatrix} = \begin{bmatrix} z_1^2 & z_1z_2 & z_1z_3 \\ z_1z_2 & z_2^2 & z_2z_3 \\ z_1z_3 & z_2z_3 & z_3^2 \end{bmatrix}, \tag{4.5}$$

and the constraint that $z_1 z_2 = z_3^2$ *translates into the condition* $w_{12} = w_{33}$.

The dual problem gives direct insight in the process of checking, after solving the LMIs, if the relaxation was *exact*. In this case, under no degeneracies, the optimal W matrix will have *rank one*, and the components of the corresponding eigenvector will verify the *constraints* satified by the z_i variables.

It should be noted that, at least in principle, the method has some degree of conservativeness. As explained above, this is because the class of psd polynomials is not equal to the sos ones. It is not clear yet how relevant this gap is in practical terms. After all, almost every time the positivity of a function needs to be established (for example, backstepping methods in control theory), this is usually done by constructing a sos representation, either implicitly or explicitly. In any case, there are possible workarounds, at some computational cost. For a psd $F(x)$, 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 [79] 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 4.5 *Consider the case of the Motzkin form given in equation (4.3). As mentioned before, 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 LMIs 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. When applying this improved method to the problems with nonzero gaps in Example 4.3, exact solutions are obtained. □

Additional stronger conditions can be obtained using the Positivstellensatz presented in Section 4.4.2 below.

4.3.1 Computational considerations

The computational cost of the procedure clearly depends on both the degree of the polynomial, and the number of variables. The number of monomials of degree less than or equal to $m/2$ (m is even) is $N_z := \binom{n-1+m/2}{m/2}$. This is the size of the resulting LMI, assuming no simplifications occur, which is not usually the case. The number of constraints (additional variables in the LMIs) can be large, especially when using many variables and high degree polynomials. For a fixed degree, however, that number is always a polynomial expression in n (it is always bounded by N_z^2, for instance).

4.4 Algebraic geometry

At its most basic level, algebraic geometry deals with the study of the solution set of a system of polynomial equations. From a more general viewpoint, it focuses on the close relationship between geometric objects and the associated abstract 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 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, some critical elements from algebraic geometry theory are presented. The usual name for the specific class of theorems we introduce 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 $\mathbb{C}$. When we are interested only in real roots, we need to introduce the Artin-Schreier theory of formally real fields, that was developed along the search for a solution of Hilbert's 17th problem.

4.4.1 Hilbert's Nullstellensatz

The theorem below establishes a correspondence between a geometric object (an *affine variety*) and an algebraic concept, a *polynomial ideal.* Hilbert's Nullstellensatz basically establishes that in an algebraically closed field, the only ideal that represents the empty variety is the entire polynomial ring.

Theorem 4.2 (Hilbert's Nullstellensatz) *Let k be an algebraically closed field (such as $\mathbb{C}$) and let $I \subset k[x_1, \ldots, x_n]$ be an ideal satisfying $\mathbf{V}(I) = \emptyset$. Then $1 \in I$, or equivalently, $I = k[x_1, \ldots, x_n]$.*

The Nullstellensatz can be applied to show the nonexistence of solutions for a given system of polynomial equations

$$f_i(x) = 0, \qquad i = 1, \ldots, m, \qquad x \in \mathbb{C}^n. \tag{4.6}$$

To do this, we need to show that the corresponding variety is empty. Using the Nullstellensatz, this is equivalent to verifying that the identity polynomial is in the ideal generated by the given equations. In other words, we need to find polynomials g_i such that

$$f_1(x)g_1(x) + \cdots + f_m(x)g_m(x) = 1. \tag{4.7}$$

The sufficiency of the condition should be obvious. If the equality above is actually satisfied for some polynomials g_i, and assuming there exists a point x_0 in the variety, after evaluating (4.7) at x_0 we immediately reach the contradiction 0=1.

Remark 4.1 *In the case when $i = 2$, and the f_i are elements of the ring of stable and proper transfer functions, the Nullstellensatz reduces to the* Bezout identity *used in the coprime factorization approach to feedback stabilization. In that case, the nonexistence of common zeros can be interpreted as the lack of unstable pole/zero cancellations.*

The polynomials g_i provide a *certificate* (usually called a Nullstellensatz refutation) that the variety described by (4.6) is empty. Given the g_i, the equality (4.7) can be checked in a number of operations that is polynomial in their length (if we only count real operations, in the number of nonzero coefficients).

There are at least two possible approaches to effectively find polynomials g_i. 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 [19, 53, 7]. 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.

Given a upper bound on the degree, and a parameterization of the unknown polynomials g_i, then a solution can be obtained by solving a system of linear equations.

The other procedure is based on Gröbner basis methods [26, 64]. 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. For example, for the case of degree one (only linear terms), and a lexicographic term ordering, a Gröbner basis is basically equivalent to Gaussian elimination. If the variety is empty, the corresponding basis has only one element, the identity polynomial. As a byproduct of the computation of a Gröbner basis, it is possible to explicitly obtain the polynomials g_i.

Example 4.6 *As an example of a Nullstellensatz refutation, we will 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

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

After simple algebraic manipulations, we can verify that

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

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

Example 4.7 (Modus ponens) *The* modus ponens *is a basic rule of inference for propositional logic. It establishes that from the truth of the statements* A *and* $A \to B$, *we can conclude that* B *is also true.*

It is possible to give an algebraic "translation" of the modus ponens rule. For example, for the case of three propositions, we have that the statements

$$P_1, \qquad P_1 \to P_2, \qquad P_2 \to P_3, \qquad \neg P_3,$$

cannot all be true simultaneously. Let the variables $x_i \in \{0, 1\}$ *(this can be dealt with by adding the constraints* $x_i(1 - x_i) = 0$*). Associating the "truth" of* P_i *to the variable* x_i *taking the value one, then the modus ponens is equivalent to the nonexistence of a common solution of the equations* $f_i(x) = 0$, *where*

$$\begin{aligned}
f_1(x) &:= 1 - x_1 && (4.8)\\
f_2(x) &:= x_1(1 - x_2) && (4.9)\\
f_3(x) &:= x_2(1 - x_3) && (4.10)\\
f_4(x) &:= x_3 && (4.11)
\end{aligned}$$

An algebraic proof of the validity of the modus ponens in this case is given by the *identity:*

$$(1-x_2)f_1 + f_2 + f_3 + x_2 f_4 = 1.$$

It can be shown [22] that in the general case of n variables, the minimum degree of the required Nullstellensatz refutation is approximately $\log n$, with the result being tight. □

An equivalent statement of the Nullstellensatz is in an ideal-theoretic formulation. Let $\mathbf{I}(V)$ be the polynomial ideal associated with a given algebraic variety, and $\mathbf{V}(I)$ the variety generated by the ideal I (see Appendix A). In this case, given a polynomial ideal I, a concise statement of the (strong) Nullstellensatz is

$$\mathbf{I}(\mathbf{V}(I)) = \sqrt{I},$$

where $\sqrt{I}$ is the *radical* of I.

In many cases, we need to verify if a polynomial f vanishes on a given algebraic variety V. In algebraic terms, this turns out to be equivalent to the *radical membership* problem: does f belong to the radical of the ideal associated with the variety V? By introducing a slack variable and applying Theorem 4.2, the characterization below can be obtained [26].

Theorem 4.3 *Let k be an algebraically closed field. If $f, f_1, \ldots, f_m \in k[x_1, \ldots, x_n]$ are such that*

$$f_1(x) = \ldots = f_m(x) = 0 \Longrightarrow f(x) = 0,$$

then there exists an integer $k \geq 1$ and $g_i \in k[x_1, \ldots, x_n]$ such that

$$f^k = g_1 f_1 + \cdots + g_n f_n.$$

Example 4.8 *Let $I := \langle x^2+y^2, x-y\rangle$. The corresponding variety V consists of just an isolated point, the origin, i.e., $V = \{(0,0)\}$. Therefore, the polynomial $f := x+y$ vanishes in V. However, it can be verified that f does not belong to the ideal I, that* *is, we cannot write f as a linear combination*

$$f = g_1(x,y)(x^2+y^2) + g_2(x,y)(x-y),$$

where $g_1, g_2 \in \mathbb{R}[x,y]$. To see this, just let $y = x$, obtaining $2x = g_1(x,x)2x^2$, from where a contradiction follows.

However, f is in the radical $\sqrt{I}$, as the following equality shows:

$$f^2 = 2(x^2+y^2) + (y-x)(x-y).$$

□

4.4.2 Positivstellensatz

The conditions in the Nullstellensatz are only necessary and sufficient in the case when the field is algebraically closed (as in the case of $\mathbb{C}$). In the case 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 corresponding variety is empty). However, the corresponding polynomial ideal does not include the element 1.

In the case where we are mainly interested in real solutions, the lack of algebraic closure 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 [12, 76] and the references therein.

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

$$\sum_{i=1}^{n} x_i^2 = 0 \Longrightarrow x_1 = \ldots = x_n = 0. \tag{4.12}$$

A field will be called *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 this chapter. For example, an alternative (but equivalent) statement of (4.12) 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 and Positivstellensatz [85, 12].

Before proceeding further, we need to introduce a few concepts. Given a set of polynomials $p_i \in k[x_1, \ldots, x_n]$, denote by $M(p_i)$ the *multiplicative monoid* generated by the p_i, i.e., the set of finite products of the elements p_i (include the empty product, the identity). The following definition introduces the ring-theoretic concept of *cone*. Despite having the same name, it should not be confused with the geometric cones described in Chapter 2.

Definition 4.3 *Let R be a commutative ring. A* cone *P of R is a subset of R 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. $a \in R \Rightarrow a^2 \in P$

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

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

With these elements, a Positivstellensatz for the reals can be formulated [85]. For concreteness it is stated for $\mathbb{R}$, instead of the general case of a real closed field.

Theorem 4.4 ([12, 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*

$$\{x \in \mathbb{R}^n | f_j(x) \geq 0, j = 1,\ldots,s, g_k(x) \neq 0, k = 1,\ldots,t, h_\ell(x) = 0, j = 1,\ldots,\ell\}$$

   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 [12] for the necessity of condition 1.*

Assume that the set is not empty, and consider an element x_0 *from the set. 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* $f + g^2 + h = 0$. □

As we can see, the Positivstellensatz guarantees the existence of *infeasibility certificates*, given by the polynomials f, g and h. Again, for complexity reasons these certificates cannot be polynomial time checkable for every possible instance, unless NP=co-NP.

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 over a semialgebraic set. Consider for simplicity the problem of verifying if the implication

$$a(x) = 0 \Rightarrow b(x) \geq 0 \tag{4.13}$$

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 particularly simple solution 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:

$$b(x) + r(x)a(x) \qquad \text{is a sum of squares,} \tag{4.14}$$

which clearly implies that (4.13) holds.

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 [81] 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 4.4 only guarantees in this case the existence of g, h in the cone such that $fg = 1 + h$.

Example 4.9 *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. □

Theorem 4.4 provides the basis for a whole class 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 we can express the sum of squares condition as a set of LMIs. Restricting the degree of the possible multipliers, we obtain semidefinite programs, that can be efficiently solved.

4.4.3 The S-procedure

The well-known *S-procedure* [17] can be interpreted (in the finite dimensional case) as a specialization of the Positivstellensatz, in the case when the polynomials are quadratic forms and the associated "multipliers" are essentially constants. To see this, consider the usual problem of establishing that a set $\mathcal{A}$ described by quadratic equations is empty:

$$\mathcal{A} := \left\{ \mathbf{x} \in \mathbb{R}^n \,|\, \mathbf{x} \neq 0,\ A_i(\mathbf{x}) := \mathbf{x}^T A_i \mathbf{x} \geq 0, i = 1, \ldots, m \right\} \tag{4.15}$$

Define now the following polynomials:

$$\begin{aligned} f &= \mathbf{x}^T\mathbf{x}\,(s + \sum_{i=1}^{m} \lambda_i\, A_i(\mathbf{x})), \qquad \lambda_i \geq 0, s \text{ is a sum of squares} \\ g &= \mathbf{x}^T\mathbf{x} \\ h &= 0 \end{aligned}$$

Notice that f is in the cone generated by the $A_i(\mathbf{x})$, and g in the monoid corresponding to the inequality $\mathbf{x} \neq 0$. In this case, the equality $f + g^2 + h = 0$ reduces to:

$$(\mathbf{x}^T\mathbf{x})(s + \sum_{i=1}^{m} \lambda_i A_i(\mathbf{x}) + \mathbf{x}^T\mathbf{x}) = 0,$$

or equivalently, since $\mathbb{R}[x]$ is an integral domain and quadratic forms are sum of squares if and only if they are positive definite:

$$\sum_{i=1}^{m} \lambda_i A_i \leq -I, \qquad \lambda_i \geq 0$$

which is a standard formulation of the S-procedure.

In Chapter 6, we present stronger versions of the S-procedure, based on the sum of squares approach.

4.5 A simple interpretation

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 (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 simultaneosly searching over all these possibilities, and combining the results, we can obtain a proof of the infeasibility of the original system. These operations are carried over by the optimization procedure.

It would be interesting to expand the connections with related ideas that have been explored in the context of "lift-and-project" methods [59, 58, 82] for deriving valid inequalities in zero-one combinatorial optimization problems. In those papers, the authors 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 programs, a known NP-hard problem. 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 differences in our work, however, are the extensions to the general continuous case via the sum of squares decomposition, and the use of the Positivstellensatz to formulate the corresponding sufficient conditions.

4.6 Application example

In the following example, we use the Positivstellensatz to compute a lower bound on the distance between a point and an algebraic curve.

Example 4.10 *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

$$\min_{C(x,y)=0} (x - x_0)^2 + (y - y_0)^2 \tag{4.16}$$

A lower bound on the optimal value can be obtained as in equation (4.14). Restricting the degree of the auxiliary Positivstellensatz polynomials to a simple linear expression in x, *we can compute the maximum value of* γ *that satisfies*

$$(x-1)^2 + (y-1)^2 - \gamma + (\alpha + \beta x)(x^3 - 8x - 2y) \quad \text{is a sum of squares.} \tag{4.17}$$

It should be clear that if condition (4.17) holds, then every pair of points (x, y) *in the curve are at a distance at least equal to* $\gamma^{1/2}$ *from* (x_0, y_0). *To see this, note that if* (x, y) *are in the curve* $C(x, y) = 0$, *then the last term in (4.17) vanishes, and therefore* $(x-1)^2 + (y-1)^2 \geq \gamma$. *Since the expression is affine in* α, β, *and* γ, *the problem can be solved by LMI methods.*

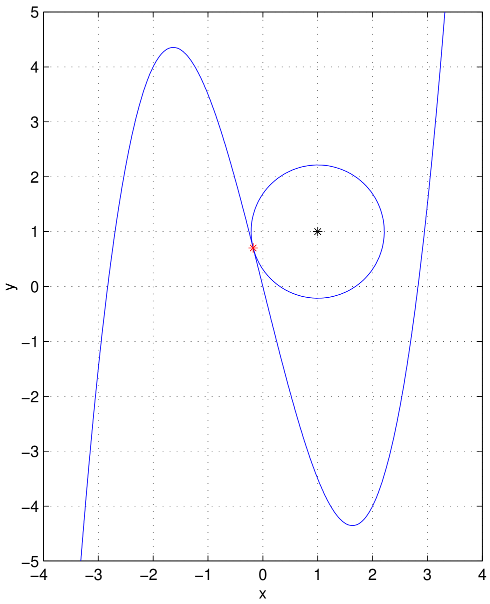

Figure 4.1: The curve $C(x, y) = 0$ and the minimum distance circle.

The optimal solution of the LMIs is:

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

In this case it can be shown that the obtained bound γ is actually optimal, since it is achieved by the values

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

In Figure 4.1 a plot of $C(x)$ and the optimal solution is presented.

Notice that the original optimization formulation (4.16) 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.* □

In the upcoming chapters, we present some concrete applications of the general approach developed so far.

Chapter 5

Copositive matrices

The verification of matrix copositivity is a well-known computationally hard problem, with many applications in continuous and combinatorial optimization. In this chapter, we present a hierarchy of semidefinite programming based sufficient conditions for a real matrix to be copositive. These conditions are obtained through the use of the sum of squares decomposition for multivariable forms, presented in Chapter 4. As can be expected, there is a tradeoff between conservativeness of the tests and the corresponding computational requirements. The proposed tests will be shown to be exact for a certain family of extreme copositive matrices.

5.1 Copositivity

A real matrix M is said to be *copositive* if the quadratic form $x^T M x$ takes only positive values in the nonnegative orthant (except the origin). Without loss of generality, we can take M to be symmetric. As opposed to positive definiteness, which can be efficiently verified (for example, using Gaussian elimination), there seems to be no polynomial time algorithms for checking copositiveness.

Copositive matrices have numerous applications in diverse fields of applied mathematics, especially in optimization. It is a critical ingredient in the characterization of local solutions of constrained optimization problems [65], such as the linear complementarity problem [25]. Also, it has been recently shown that its use can notably improve certain convex relaxation bounds in quadratic programming problems with linear constraints [75]. As we have seen in the past chapters, these convex relaxations are the underlying basis of many important results in robustness analysis. A recent example of an application of copositive matrices in a control setting is in the stability analysis using piecewise quadratic Lyapunov functions [48].

From a computational complexity viewpoint, the recognition problem for copositive matrices is hard, in general. It has been shown that checking if a given matrix is not copositive is an NP-complete problem [65]. Equivalently, checking copositivity is in co-NPC (see [36, 70] for background material on computational complexity). This implies that, unless co-NP=NP (a consequence of P=NP), in general it is not possible to construct polynomial time certificates of copositivity (i.e., copositivity is not in NP).

In many cases, however, it is possible to efficiently construct such certificates. For example, assume that the matrix M has a decomposition $M = P + N$, with P positive semidefinite and N componentwise nonnegative. It is clear that this implies that M is copositive, with the matrices P and N providing a polynomial time verifiable certificate.

In a similar way, the results presented in this chapter provide a unified methodology of constructing sufficient conditions for copositivity. The procedure uses as a basic tool a sum of squares decomposition for multivariable forms, that can be obtained by semidefinite programming methods, as we have seen before. As in the other examples analyzed, one of the main advantages of the proposed procedure is that it can also be applied to the case when the coefficients of M are variable (or uncertain).

5.2 Background and notation

The notation is mostly standard. A matrix $M \in \mathbb{R}^{n \times n}$ is *copositive* if $x^T M x \geq 0 \quad \forall x \in \mathbb{R}^n, x_i \geq 0$. Equivalently, the quadratic form is nonnegative on the closed nonnegative orthant. If $x^T M x$ takes only positive values on the closed orthant (except the origin, of course), then M will be *strictly copositive*.

Recall the definition of geometric cones in page 11 of Chapter 2. A point x of a convex cone C is an *extreme point* if $x = x_1 + x_2$, $x_i \in C$ implies $x_1 = \lambda x$, $x_2 = (1-\lambda)x$, $0 \leq \lambda \leq 1$. It can be shown that the set of copositive matrices $\mathcal{C}$ is a closed convex cone [39]. We also denote as $\mathcal{P}, \mathcal{N}$ the self-dual cones of positive semidefinite and elementwise nonnegative matrices, respectively.

As in Chapter 4, denote by $\mathcal{F}_{n,m}$ the set of homogeneous polynomials (forms) of degree m in n variables $\{x_1, \ldots, x_n\}$, with real coefficients. Every such form can be written as a sum of $\binom{n+m-1}{m}$ monomials, each one of the form $c_\alpha \prod_{i=1}^n x_i^{\alpha_i}$, with $\sum_{i=1}^n \alpha_i = m$.

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 [90, 25] and the references therein). However, since checking copositivity of a matrix is a co-NP-complete problem [65], this implies 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.

We describe next two applications of copositive matrices, mentioned in the introduction. Consider first the problem of obtaining a lower bound on the optimal solution of a linearly constrained quadratic optimization problem [75].

Theorem 5.1 *Let f^* be the solution of the constrained minimization problem:*

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

If the linear matrix inequality in C, γ:

$$Q - A^T C A \geq \gamma I, \tag{5.1}$$

is feasible, with a copositive *C, then the inequality $f^* \geq \gamma$ holds.*

Proof: *Multiply (5.1) left and right by any feasible x of the original problem. Since Ax is componentwise nonnegative and C is copositive, we obtain $x^T Q x \geq \gamma$.*

□

A difficulty in the direct application of Theorem 5.1 is the fact that the set of copositive matrices, though convex, does not have a "nice" description, since even the problem of checking membership is provably hard. For this reason, having semidefinite programming conditions that guarantee copositivity would allow for enhanced bounds for this type of problems.

The other application, presented in [48, 49], deals with the analysis of piecewise linear systems using piecewise quadratic Lyapunov functions. One of the basic issues in that problem is checking nonnegativity of the Lyapunov function, in a region (or "cell") defined by linear inequalities. To this end, an LMI-based condition sufficient condition is usually employed. By using the improved copositivity tests presented in this paper, less conservative answers can be obtained, especially in the case of systems of large state dimension. The conditions in [48] basically correspond to the sufficient condition (5.3) below.

5.3 SDP conditions

In order to apply the sum of squares decomposition to the matrix copositivity problem, we need a way of dealing with the constraints in the variables, since each x_i has to be nonnegative. While we could directly impose the conditions $x_i \geq 0$ and deal with them in the standard way suggested by the Positivstellensatz, we choose here a perhaps more natural, though equivalent, procedure.

The alternative way of addressing the positivity constraint on the x variables is the following: to check copositivity of M, we can consider the change of variables $x_i = z_i^2$, 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, an obvious sufficient condition for M to be copositive is that $P(\mathbf{z})$ can be written as a sum of squares.

In order to do that, as explained in previous chapters, we have to express $P(\mathbf{z})$ as a quadratic form in the variables z_i^2 and $z_i z_j$, for $i \neq j$. In principle, the dimension of the new matrix Q is now $n + \binom{n}{2}$. The nonuniqueness of the representation follows from the identities

$$\begin{aligned}
(z_i z_j)^2 &= (z_i^2)(z_j^2) \\
(z_i z_j)(z_i z_l) &= (z_i^2)(z_j z_l) \\
(z_i z_j)(z_k z_l) &= (z_i z_k)(z_j z_l) = (z_i z_l)(z_j z_k).
\end{aligned}$$

Denote the associated free multipliers by the variables λ_{ij}, ν_{ijl}, and μ_{ijkl}, μ'_{ijkl} respectively. By grouping the variables in a vector $\mathbf{Z}$ (first the z_i^2, then the $z_i z_j$), and writing

$$P(\mathbf{z}) = \mathbf{Z}^T Q \mathbf{Z},$$

the matrix Q can be shown to have the structure

$$Q = \begin{bmatrix}
m_{11} & m_{12} - \lambda_{12} & \dots & m_{1n} - \lambda_{1n} & * & * & * & * \\
m_{12} - \lambda_{12} & m_{22} & \dots & m_{2n} - \lambda_{2n} & * & * & * & * \\
\vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots \\
m_{1n} - \lambda_{1n} & m_{2n} - \lambda_{2n} & \dots & m_{nn} & * & * & \vdots & * \\
* & * & \cdots & * & 2\lambda_{12} & * & \dots & * \\
* & * & \cdots & * & * & 2\lambda_{13} & \dots & * \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\
* & * & \cdots & * & * & * & \cdots & 2\lambda_{(n-1)n}
\end{bmatrix},$$

where the places with asterisks are either zero or a linear combination of the ν and μ variables.

Therefore, P will have a sum of squares decomposition if and only if there exists variables λ, μ, ν such that the matrix Q above is positive semidefinite. Without loss of generality, it is always possible to choose the μ, ν to be zero, since they appear only in the off-diagonal subblocks. Consequently, all the λ_{ij} should be nonnegative, and the LMI can be reduced to:

$$\begin{bmatrix} m_{11} & m_{12}-\lambda_{12} & \dots & m_{1n}-\lambda_{1n} \\ m_{12}-\lambda_{12} & m_{22} & \dots & m_{2n}-\lambda_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ m_{1n}-\lambda_{1n} & m_{2n}-\lambda_{2n} & \dots & m_{nn} \end{bmatrix} \geq 0, \qquad \lambda_{ij} \geq 0. \tag{5.2}$$

It is easy to verify that existence of such λ_{ij} turns out to be equivalent to the condition that the original matrix M can be written as the sum of a positive semidefinite and an elementwise nonnegative matrix, i.e.,

$$M = P + N, \qquad P \geq 0, \quad n_{ij} \geq 0. \tag{5.3}$$

As mentioned earlier, this is a well-known sufficient condition for copositivity (see for example [29]). The equivalence between these two tests has also been noticed in [23, Lemma 3.5]. Note that condition (5.3) can be obtained by considering the enhanced Shor relaxation, where new quadratic constraints are obtained by considering the pairwise products of linear constraints [75]. These products of constraints are exactly what can be obtained via a Positivstellensatz construction, when restricting the constraints to have degree at most equal to two.

From what we have seen so far, we are able to derive a standard sufficient test for copositivity, based on the sum of squares framework. The advantage of the approach presented here is that even *stronger* conditions can be obtained. By considering higher order forms, a hierarchy of increasingly powerful tests is obtained. Of course, the computational requirements increase accordingly.

Consider 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 LMI methods, as described in Chapter 4), we can guarantee the nonnegativity of $P(\mathbf{z})$, and as a consequence, copositivity of M.

For concreteness, we will analyze in some detail the case $r = 1$. We will see that as in the case for $r = 0$ described above, some variables automatically drop out from the optimization due to the particular structure of the resulting LMIs.

As explained, we consider now the sixth order form:

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

To express it as a quadratic form, in principle we need to define the new variables z_i^3, $z_i^2 z_j$ $(i \neq j)$ and $z_i z_j z_k$ $(i \neq j \neq k)$. There are n, $n(n-1)$, and $\binom{n}{3}$ different variables corresponding to each type. A particularly convenient ordering for the variables is the following:

$$\mathbf{Z} = [z_1 z_1^2, \ldots, z_1 z_n^2, z_2 z_1^2, \ldots, z_2 z_n^2, \ldots, z_n z_1^2, \ldots, z_n z_n^2, z_1 z_2 z_3, \ldots, z_{n-2} z_{n-1} z_n]$$

As in the case of the quartic form described above, without loss of generality it is always possible to choose some multipliers to be identically zero. This induces a block diagonal structure in the matrix Q, simplifying the final conditions.

Theorem 5.2 *Consider the system of LMIs given by:*

$$\begin{aligned} M - \Lambda^i &\geq 0, & i &= 1, \ldots, n \\ \Lambda^i_{ii} &= 0, & i &= 1, \ldots, n \\ \Lambda^i_{jj} + \Lambda^j_{ji} + \Lambda^j_{ij} &= 0, & i &\neq j \\ \Lambda^i_{jk} + \Lambda^j_{ki} + \Lambda^k_{ij} &\geq 0, & i &\neq j \neq k \end{aligned} \tag{5.4}$$

where the n matrices $\Lambda^i \in \mathbb{R}^{n \times n}$ are symmetric ($\Lambda^i_{jk} = \Lambda^i_{kj}$). If there exists a feasible solution, then M is copositive. Furthermore, this test is at least as powerful as condition (5.3).

Proof: The nonnegativity of $P_1(\mathbf{z})$ follows immediately from the LMIs above, since

$$\sum_{j,k} m_{jk} x_j x_k \geq \sum_{j,k} \Lambda^i_{jk} x_j x_k \Longrightarrow \sum_{i,j,k} m_{jk} x_i x_j x_k \geq \sum_{i,j,k} \Lambda^i_{jk} x_i x_j x_k,$$

and the coefficients of this last form are nonnegative.

It is also possible to verify directly that if the LMIs (5.2) have a solution, then so does the system (5.4). Just let

$$\Lambda^k_{ij} = \lambda_{ij}(1 - \delta_{ik} - \delta_{jk})$$

where δ is the usual Kronecker symbol, and $\lambda_{ii} = 0$ for all i. This is a consequence of the "nested" properties of the P_r-based tests. □

As we have shown, this class of tests is at least as powerful as the standard condition (5.3). A question naturally arises: how conservative is this procedure? To this end, consider the following theorem of Pólya:

Theorem 5.3 *[40, Section 2.24] Given a form $F(x_1, x_2, \ldots, x_n)$ strictly positive for $x_i \geq 0$, $\sum_i x_i > 0$, then F can be expressed as*

$$F = \frac{G}{H},$$

where G and H are forms with positive coefficients. In particular, we can choose

$$H = (x_1 + x_2 + \cdots + x_n)^r$$

for a suitable r.

Applying the theorem to a strictly copositive M (i.e., to the associated positive definite form $P(\mathbf{z})$), it is clear then that there is a finite r for which the condition based on P_r is exact. However, the minimum r in Theorem 5.3 cannot always be chosen as a polynomial expression of n (uniformly over the psd forms). The known lower bounds for r usually involve a "condition number" for the form P:

the minimum r grows as the form tends to degeneracy (nontrivial solutions). Some of these effective bounds are presented in [28, 27, 73]. However, these bounds can also be conservative: even if P has nontrivial zeros, it might be possible to prove copositivity with a small value of r, as the examples we present shows.

Some interesting questions, yet unanswered, relate to the conservativeness of the proposed tests, for fixed values of r. For example, it is known [29] that the test in equation (5.3) (i.e., the case of $r = 0$) is exact if the dimension n of the matrix M is less than or equal to four. Do similar results hold for every r? In particular, what is the minimum n for which the $r = 1$ test is not exact? In the examples, we show that a particular family of extreme copositive forms for which the $r = 0$ test fails, can be exactly recognized with the $r = 1$ criterion.

5.4 Examples

As a confirmation that the proposed technique can be strictly stronger than the standard relaxations, we will consider some particular examples from the literature.

Consider the quadratic form associated with the matrix J below.

$$J = \begin{bmatrix} 1 & -1 & 1 & 1 & -1 \\ -1 & 1 & -1 & 1 & 1 \\ 1 & -1 & 1 & -1 & 1 \\ 1 & 1 & -1 & 1 & -1 \\ -1 & 1 & 1 & -1 & 1 \end{bmatrix}$$

This form, originally introduced by A. Horn, appeared previously in [29, 75]. It has been noted in [29, note added in proof] that it is copositive, even though *it does not* satisfy the condition (5.3).

Nevertheless, it is still possible to prove its copositiveness by the method presented in this paper. For the numerical implementation of the presented procedure, we used the semidefinite programming solver SeDuMi [86]. Let $\mathbf{x} := [x_1, x_2, x_3, x_4, x_5]^T$.

Taking $r = 1$, after solving the corresponding LMIs we obtain the decomposition:

$$\begin{aligned}
(\mathbf{x}^T J\mathbf{x})\,(x_1+x_2+x_3+x_4+x_5) \;=\;& x_1(x_1-x_2+x_3+x_4-x_5)^2+\\
& x_2(x_2-x_3+x_4+x_5-x_1)^2+\\
& x_3(x_3-x_4+x_5+x_1-x_2)^2+\\
& x_4(x_4-x_5+x_1+x_2-x_3)^2+\\
& x_5(x_5-x_1+x_2+x_3-x_4)^2+\\
& 4\,(x_1x_2x_4+x_2x_3x_5+x_3x_4x_1+x_4x_5x_2+x_5x_1x_3)
\end{aligned}$$

from where copositivity of J follows immediately.

This example can be generalized to a family of copositive forms, with interesting theoretical properties. Consider the following cyclic quadratic form in $n = 3m + 2$ variables ($m \geq 1$), analyzed in [6]:

$$B(x) := \left(\sum_{i=1}^{3m+2} x_i\right)^2 - 2\sum_{i=1}^{3m+2} x_i \sum_{j=0}^{m} x_{i+3j+1} \tag{5.5}$$

where $x_{r+n} = x_r$. It is clear that the Horn form presented above corresponds to the special case $m = 1$. It has been shown in [6] that this is a extreme copositive form. Therefore, since $B(x)$ is neither componentwise nonnegative or positive semidefinite, it cannot satisfy condition (5.3). Generalizing the decomposition above, we have the following theorem:

Theorem 5.4 *Let $B(x)$ be as in equation (5.5). Then, it has the decomposition:*

$$B(x)\sum_{i=1}^{n} x_i = \sum_{i=1}^{n} x_i\left(\sum_{j=1}^{n} x_j - 2\sum_{j=0}^{m} x_{i+3j+1}\right)^2 + 4\sum_{i=1}^{n} x_i \sum_{k=1}^{m} x_{i+3k-2}\sum_{j=k}^{m} x_{i+3j} \tag{5.6}$$

Proof: *For notational simplicity, let $s_i(\mathbf{x}) := x_{i+1} + x_{i+4} + \cdots + x_{i+3m+1}$. Let $L(\mathbf{x})$*

be the left-hand side of (5.6). Then,

$$L(\mathbf{x}) = \sum_{i=1}^{n}\sum_{j=1}^{n}\sum_{k=1}^{n} x_i x_j x_k - 2\sum_{i=1}^{n}\sum_{k=1}^{n} x_i x_k s_i(\mathbf{x})$$

The first term in the right-hand size of (5.6) can be written as:

$$\begin{aligned} R_1(\mathbf{x}) &= \sum_{i=1}^{n} x_i \left(\sum_{j=1}^{n} x_j - 2s_i(\mathbf{x})\right)^2 \\ &= \sum_{i=1}^{n}\sum_{j=1}^{n}\sum_{k=1}^{n} x_i x_j x_k - 4\sum_{i=1}^{n}\sum_{j=1}^{n} x_i x_j s_i(\mathbf{x}) + 4\sum_{i=1}^{n} x_i s_i^2(\mathbf{x})^2 \end{aligned}$$

Subtracting, we obtain:

$$\begin{aligned} L(\mathbf{x}) - R_1(\mathbf{x}) &= 2\sum_{i=1}^{n}\sum_{j=1}^{n} x_i x_j s_i(\mathbf{x}) - 4\sum_{i=1}^{n} x_i s_i^2(\mathbf{x})^2 \\ &= 2\sum_{i=1}^{n} x_i \left(\sum_{j=1}^{n} x_j - 2s_i(\mathbf{x})\right) s_i(\mathbf{x}) \end{aligned}$$

Expanding inside the sum, and cancelling identical terms corresponding to different values of i, after some manipulations we obtain the expression:

$$R_1(\mathbf{x}) = 4\sum_{i=1}^{n} x_i \sum_{k=1}^{m} x_{i+3k-2}(x_{i+3k} + x_{i+3(k+1)} + \cdots + x_{i+3m}),$$

from where the result follows. $\square$

Chapter 6

Higher order semidefinite relaxations

In this chapter, we specialize the general machinery presented earlier in order to formulate improved versions of the standard semidefinite relaxation for quadratic programming. This framework 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 case, 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.

6.1 Introduction

Many problems in systems and control theory, especially in robustness analysis and synthesis, have intrinsically "bad" computational complexity properties. As mentioned in the introduction, 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 are increasingly being used for a variety of problems in diverse fields of applied mathematics. These SDP relaxations are convex optimization problems, that can be solved in polynomial time. The procedure by which we obtain a relaxed problem and its dual is known in the literature under several different names, i.e., S-procedure, Shor relaxation, covariance relaxation, lifting, etc. [91]. 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 [32]) are significantly harder, since even the approximation problem within an arbitrary constant factor is NP-hard.

In this chapter, we present a novel convex relaxation of quadratic programming problems, that runs in polynomial time. The idea can be interpreted as finding a separating functional (not necessarily linear) that proves that the intersection of two sets is empty. As in the previous chapter, we employ as a basic technique the existence of a sum of squares decomposition as a sufficient condition for nonnegativity of a multivariable form.

6.2 The standard SDP relaxation

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

$$\mathcal{A} := \left\{\mathbf{z} \in \mathbb{R}^m \,|\, z_i \geq 0,\ z_i = \mathbf{x}^T A_i \mathbf{x},\, \mathbf{x} \in \mathbb{R}^n/\{0\}, i = 1, \ldots, m\right\} \tag{6.1}$$

For future reference, let $\mathbf{a}(\mathbf{x}) := [\mathbf{x}^T A_1 \mathbf{x}, \ldots, \mathbf{x}^T A_m \mathbf{x}]^T$. Both logical implications and constrained optimization problems can be put in the form (6.1), by checking for the existence of a counterexample, or a feasible point that achieves a given level

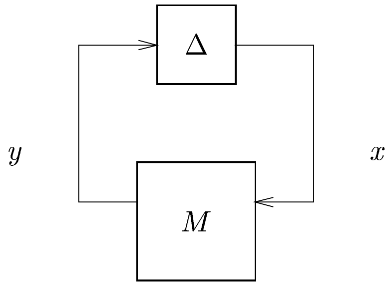

Figure 6.1: Plant M and uncertainty diagonal structure Δ.

of optimality, respectively.

A simple sufficient condition for the set $\mathcal{A}$ defined in (6.1) to be empty is given by the existence of numbers λ_i that satisfy the condition:

$$\sum_{i=1}^{m} \lambda_i A_i < 0, \qquad \lambda_i \geq 0. \tag{6.2}$$

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 should be nonnegative, since both vectors are componentwise nonnegative. However, equation (6.2) makes that inner product negative. As a consequence, $\mathcal{A}$ is empty.

Note that condition (6.2), also known as the S-procedure, is a linear matrix inequality (LMI), also known as an instance of a *semidefinite program* [91]. As is widely recognized today, this class of convex optimization problems can be efficiently solved, both in theory and practice.

Example 6.1 *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 [69]. For simplicity, let the matrix $M \in \mathbb{R}^{n\times n}$, $\Delta = \mathrm{diag}(\delta_1, \ldots, \delta_n)$ and the scalar uncertainties δ_i be real. In the notation of Figure 6.1, the condition that the absolute value of the uncertainties δ_i is bounded by $1/\gamma$, is equivalent to the quadratic inequalities:*

$$\delta_i^2 \leq 1/\gamma^2 \Longleftrightarrow y_i^2 - \gamma^2 x_i^2 = \mathbf{x}^T(M_i^T M_i - \gamma^2 E_{ii})\mathbf{x} \geq 0, \tag{6.3}$$

where E_{ii} is the matrix with zero elements, except for a one in the (i,i) position, and

M_i is the ith row of the matrix M. Therefore, for this particular case, the matrices A_i are given by $A_i = M_i^T M_i - \gamma^2 E_{ii}$. In this case, the nonexistence of nontrivial solutions can be interpreted as the robust stability of the system under uncertainty.

When we apply the SDP relaxation to the system of inequalities (6.3), we obtain the usual μ upper bound LMI, with D being a diagonal matrix:

$$M^T D M - \gamma^2 D < 0, \qquad D > 0. \tag{6.4}$$

□

It is also interesting to study the dual problem of (6.2). It consists of checking for the existence of a symmetric matrix $Z \neq 0$, that satisfies

$$\mathrm{trace} A_i Z \geq 0, \qquad Z \geq 0. \tag{6.5}$$

This dual problem can also be obtained directly from (6.1), by using the cyclic properties of the trace function, and dropping the rank one condition on the matrix $Z := \mathbf{x}\mathbf{x}^T$ [91]. If this dual problem does not have a solution, then neither does the original one. But at least in principle, an affirmative answer to the feasibility of (6.5) does not necessarily say anything about the set $\mathcal{A}$ (in some special cases, it is possible to extract useful information from the matrix Z).

6.3 Separating functionals and a new SDP relaxation

In order to extend the standard condition, we will be considering the well-known interpretation of the multipliers λ_i in (6.2) 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 (6.2) 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 [54, Section 13.5]. For concreteness, we will 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 computation requirements can be demanding.

Extending the definitions from the previous chapter, a functional $\phi : \mathbb{R}^n \to \mathbb{R}$ 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:

$$\phi(\mathbf{a}(\mathbf{x})) < 0, \qquad \forall \mathbf{x} \in \mathbb{R}^n / \{0\} \tag{6.6}$$

The reasons are exactly as above: the existence of a possible $\mathbf{x}$ 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 (6.6)? From a complexity viewpoint, these two questions are as intractable as the original problem. It turns out that for the case of polynomial functionals ϕ, a partial answer to both questions can be obtained by using the sum of squares decomposition presented in Chapter 4.

For the exact copositivity problem, the results mentioned in the previous chapter show that checking if a quadratic form is not copositive is an NP-complete problem [65]. As we have seen, a simple sufficient condition for copositivity of a matrix Φ (see Chapter 5 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.:

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

Notice 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. For reasons that will be clear later, we would like the LHS of (6.6) 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 6.1 *Assume there exists solutions $Q_i, T_i \in \mathbb{R}^{n\times n}, r_{ij} \in \mathbb{R}$ to the equation*

$$\sum_{i=1}^{n_a} Q_i(x) A_i(x) + \sum_{1\le i<j\le n_a} r_{ij} A_i(x) A_j(x) + \sum_{j=1}^{n_b} T_j(x) B_j(x) < 0, \qquad \forall x \in \mathbb{R}^n/\{0\}. \tag{6.7}$$

where $Q_i(x) := x^T Q_i x, T_j(x) := x^T T_j x, Q_i \ge 0$ and $r_{ij} \ge 0$. Then, the only solution of

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

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.* □

Note that the left-hand size 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 LHS of (6.7) can be written (except for the sign change) as a sum of squares. As we have seen before in Chapter 4, 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 (6.2) is feasible, then the left-hand side of (6.7) is obviously a sum of squares (recall that positive definite quadratic forms are always sums of squares).

Remark 6.1 *It is interesting to compare this condition with the Nullstellensatz and Positivstellensatz in Chapter 4. The first two terms in (6.7) belong to the* cone *generated by the $A_i(x)$, and the remaining one to the ideal corresponding to the $B_i(x)$. The degree of the multipliers is restricted, so the whole expression is an homogeneous form of fixed degree.*

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 (6.7) 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 (6.7) 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 size of (6.7) linear in γ. In principle, this last technique can be conservative, when comparing to the case where all the variables are free.

6.3.1 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 dimensions $\binom{n+1}{2}$. However, the main difficulty is really caused by the large number of variables, since the ones arising from the redundant constraints, as explained before, are $O(n^4)$. 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 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.

6.4 Relaxations and moments

In the case where the relaxation is not exact, we do not obtain a feasible point in the primal problem, and end up only with lower bounds on the optimal value. Naturally, we would also like to have some upper bounds, so it would be interesting to have some approximate procedure or guidelines to construct a primal feasible point. In this case, a sensible approach, very successful in some specific problems, is a *randomized* procedure.

In the standard SDP relaxation, the dual variables can be interpreted as providing the matrix of second moments for a particular probability distribution. In the case of the MAX CUT problem discussed in the examples below, for instance, primal points can be constructed by randomly generating points consistent with this probability density (given by the matrix Y), and rounding them to the values ± 1. In this specific case, good bounds can be obtained on the expected value of the resulting cut [37].

In principle, in certain instances we can do so in our case too. However, there are some important differences. In the quadratic case, *any* positive semidefinite matrix is a valid candidate for a set of second moments; for example, we can construct a multivariate normal distribution with that preassigned covariance. However, for higher order moments, not every set of numbers obtained from the relaxation necessarily correspond to the moments of a measure [1, 8]. The root of this problem, it turns out, is again the distinction between the conditions of nonnegativity of a polynomial and being a sum of squares.

A notable exception is the one dimensional case, since given a sequence of moments, positive semidefiniteness of the corresponding Hankel matrix is enough to guarantee the existence of a measure with exactly those moments [1]. Interestingly enough, this problem is very related to the Nevanlinna-Pick interpolation questions studied in $\mathcal{H}_\infty$ control.

6.5 Examples

In this section, we present a few applications of the new relaxations to some practically important problems.

6.5.1 Structured singular value upper bound

As mentioned in Example 6.1, the standard upper bound of the structured singular value μ [69] 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 presented techniques.

Given a matrix $M \in \mathbb{C}^{n\times n}$, and an uncertainty structure $\mathbf{\Delta}$, define the structured singular value μ as:

$$\mu_{\mathbf{\Delta}}(M) := \frac{1}{\min\{\|\Delta\| : \Delta \in \mathbf{\Delta}, \det(I - M\Delta) = 0\}}, \tag{6.8}$$

unless no Δ makes $I - M\Delta$ singular, in which case $\mu_{\mathbf{\Delta}}(M) := 0$. An upper bound for μ can be obtained by solving the LMI presented in Example 6.1.

We consider next the counterexample, due to Morton and Doyle, to the proposition that μ equals to its standard upper bound in the case with four scalar uncertainties [69, Section 9.2]. This corresponds to a certain rank two matrix $M \in \mathbb{C}^{4\times 4}$, given by:

$$M := UV^*, \qquad U = \begin{bmatrix} a & 0 \\ b & b \\ c & jc \\ d & f \end{bmatrix}, \qquad V = \begin{bmatrix} 0 & a \\ b & -b \\ c & -jc \\ -jf & -d \end{bmatrix},$$

with $a = \sqrt{2/\gamma}, b = 1/\sqrt{\gamma}, c = 1/\sqrt{\gamma}, d = -\sqrt{\beta/\gamma}, f = (1+j)\sqrt{1/\gamma\beta}, \gamma = 3+\sqrt{3}$, and $\beta = \sqrt{3}-1$. This matrix has a value of $\mu \approx 0.8723$. However, the standard μ upper bound, given by equation (6.4), 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.

6.5.2 The MAX CUT problem

The maximum cut (MAX CUT) problem consists in finding a partition of the nodes of a graph in two disjoint sets 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 [36].

By casting the problem as a boolean maximization, we can write the MAX CUT problem as an equality constrained quadratic program. One standard formulation is the following:

$$\max_{y_i \in \{-1,1\}} \frac{1}{2} \sum_{i,j} w_{ij}(1 - y_i y_j), \tag{6.9}$$

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 (6.9) using semidefinite programming. Removing the constant term, and changing the sign, the original problem is clearly equivalent to:

$$\min_{y_i^2=1} \sum_{i,j} w_{ij} y_i y_j. \tag{6.10}$$

The corresponding semidefinite relaxation is given by:

$$\min_{Y \geq 0, Y_{ii}=1} \operatorname{trace} WY, \tag{6.11}$$

and its dual

$$\max_{D \leq W} \operatorname{trace} D, \tag{6.12}$$

where D is a diagonal matrix. Any feasible solution of the dual (6.12) provides a lower bound on the optimal value of (6.11), and therefore on that of (6.10).

It has been recently shown by Goemans and Williamson [37] that by randomly truncating in an appropriate manner the solution Y of this relaxation, a cut with an expected value greater than 87% of the optimal MAX CUT solution is obtained. In this sense, for the MAX CUT problem the semidefinite relaxation is provably "good." Note however that for other NP-complete problems, such as MAX CLIQUE, no such approximation results hold, unless P=NP.

The enhanced relaxations developed earlier in this chapter can be directly applied, by testing if the set of solutions y_i of (6.9) that achieve a value greater than or equal to γ is empty. Since the constraints defining the problem are quadratic, this problem formulation corresponds exactly to the setting of Theorem 6.1. The variable γ can be included in the optimization problem, as described in page 77.

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

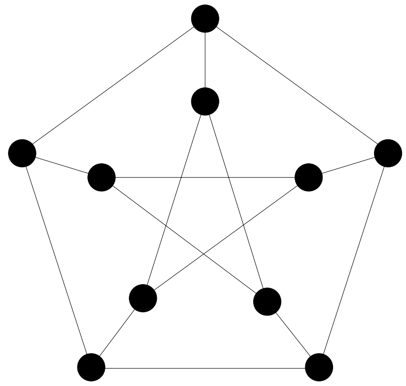

Figure 6.2: The Petersen graph.

(i.e., equal to n). For the odd n case, we have the upper bound

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

For this class of graphs, the gap is maximal in the case of the 5-cycle ($k = 2$). 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 developed procedure 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 6.2. This nonplanar graph has ten nodes and fifteen edges, and has very interesting theoretical properties [43]. 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 [4], a different strengthened SDP relaxation for 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[1]. Of course, a fair comparison should also take into

[1]In a very recent work [5], the same authors present yet another relaxation, which attains exact bounds for these cases. The possible connections between this new relaxation and the one proposed here certainly deserve more analysis.

account the computational requirements, which are higher in our proposed method than in that of [4]. We also note that a usual technique to decrease the possible conservativeness of the MAX CUT relaxation is to add linear *odd cycle* constraints. The complexity of doing this (for the three point case) is lower than the one of our proposed relaxation. For this case, in the small problems we have tested, the results seem to be equivalent. However, more numerical experience and theoretical insight is needed in order to formulate accurate comparisons.

6.6 Final overview

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 idea is to use a 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, especially in terms of practical performance.

Chapter 7

Applications in systems and control

In this chapter, we show how the methods developed in the preceding sections can be profitably applied to systems and control related problems. Some of the presented applications correspond to well-studied problems, such as Lyapunov function computation, while others, such as robust bifurcation analysis, are relatively new.

The main insight underlying the results in this chapter is that under certain assumptions, many conditions (for example, existence of a Lyapunov function) can be equivalently formulated in terms of polynomial equalities and inequalities. In other words, the set of feasible parameters is a *semialgebraic set.* In this case, operations such as testing for emptyness, obtaining bounds on the distance to a given point, etc., can all be formulated and solved within the framework described in Chapter 4. The main advantages are the resulting computational tractability (since it reduces to semidefinite programs), as well as the algorithmic character of the solution procedure.

As an motivating example of the methodology, we will deal in the next section mainly with the stability analysis of systems described by polynomial vector fields. Later we will show that the same techniques can be employed to more complicated problems.

7.1 Lyapunov stability

Stability analysis can be reduced, using Lyapunov theory, to the existence of a positive definite function, such that its time derivative along the trajectories of the system is negative. As is well known, to prove asymptotic stability of a fixed point of vector field (the origin, without loss of generality) it is required to find a Lyapunov function $V(x)$ such that:

$$\begin{aligned} \dot{x} &= f(x), \qquad V(x) > 0 \quad x \neq 0, \\ \dot{V}(x) &= \left(\frac{\partial V}{\partial x}\right)^T f(x) < 0, \quad x \neq 0 \end{aligned} \tag{7.1}$$

for all x in a neighborhood of the origin. If we want global results, we need additional conditions such as V being radially unbounded.

In the specific case of linear systems $\dot{x} = Ax$ and quadratic Lyapunov functions $V(x) = x^T P x$, this stability test is equivalent to the well-known LMIs

$$A^T P + PA < 0, \qquad P > 0.$$

The existence of a P satisfying this last condition can be checked efficiently, using for instance interior point methods.

For nonlinear systems, in the general case there are no systematic methodologies for the search for Lyapunov functions [51]. Nevertheless, in the presence of additional structure, such as the case of *mechanical systems*, sometimes it is possible to find natural energy-based Lyapunov functions. Alternative approaches use an embedding (overbounding) of the given nonlinear system in a class of uncertain linear systems. This is the case, to cite a few, of conic sector bounds, Linear Parameter Varying (LPV) and Integral Quadratic Constraints (IQC, [61]) based methods. The methology presented in this section, on the contrary, handles polynomial nonlinearities *exactly*.

In the attempt to extend the algorithmic formulation to more general vector fields (not necessarily linear) or Lyapunov functions (not necessarily quadratic), we are faced with the basic question of how to verify in a systematic fashion the conditions (7.1). If we want to develop an algorithmic approach to nonlinear system analysis, similar to what is available in the linear case, we need some explicit way of *testing the global positivity of a function.* In the case of polynomial functions, a tractable sufficient condition, as presented in Chapter 4, is the existence of a sum of squares decomposition.

Example 7.1 *The system below is from [13, Example 2.5]. Given the nonlinear system*

$$\begin{aligned} \dot{x}_1 &= -x_1^3 - x_2 x_3 - x_1 - x_1 x_3^2 \\ \dot{x}_2 &= -x_1 x_3 + 2x_1^3 - x_2 \\ \dot{x}_3 &= -x_3 + 2x_1^2 \end{aligned}$$

and the (fixed) Lyapunov function $V(x) = \frac{1}{2}(x_1^2 + x_2^2 + x_3^2)$, *test if* $\dot{V}(x)$ *is negative definite.*

After computing $\dot{V}$, *we can test if we can express it as a sum of squares using the methodology described. In this case, the decomposition*

$$-\dot{V}(x) = x_1^2 + x_3^2 + (x_1^2 - x_1 x_3 - x_2)^2$$

is obtained, from where global stability follows. □

7.2 Searching for a Lyapunov function

Given an affine parametrization $V(x,p)$ of the Lyapunov function, the search for a Lyapunov function can be automated, since in this case the polynomial

$$-\dot{V}(x,p) = -\left(\frac{\partial V}{\partial x}\right)^T f(x)$$

is again affine in p. Therefore, by including the parameters p as variables in the LMI, the full problem can be reformulated as a linear matrix inequality.

The following example shows an application of the method to a nonlinear second order system:

Example 7.2 *Consider the system described by:*

$$\begin{aligned} \dot{x}_1 &= -x_1 - 2x_2^2 \\ \dot{x}_2 &= -x_2 - x_1 x_2 - 2x_2^3; \end{aligned}$$

Notice that the vector field is invariant under the symmetry transformation $(x_1, x_2) \to (x_1, -x_2)$. *We could potentially use this information in order to limit the search to symmetric candidate Lyapunov functions. However, we will not do so, to show the method in its full generality. To look for a Lyapunov function, we will use the general expression of a polynomial in* x_1, x_2 *of degree four with no constant or linear terms (because* $V(0) = 0$, *and* V *has to be positive definite). We use a matrix representation for notational clarity.*

$$V(x) = \begin{bmatrix} 1 \\ x_1 \\ x_1^2 \\ x_1^3 \\ x_1^4 \end{bmatrix}^T \begin{bmatrix} 0 & 0 & c_{02} & c_{03} & c_{04} \\ 0 & c_{11} & c_{12} & c_{13} & 0 \\ c_{20} & c_{21} & c_{22} & 0 & 0 \\ c_{30} & c_{31} & 0 & 0 & 0 \\ c_{40} & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ x_2 \\ x_2^2 \\ x_2^3 \\ x_2^4 \end{bmatrix}$$

It is easy to verify that V *can be represented as* $V(x) = \frac{1}{2} z^T Q z$, *where* $z = [x_1, x_1^2, x_1 x_2, x_2^2, x_2]^T$ *and*

$$Q = \begin{bmatrix} 2c_{20} & c_{30} & c_{21} + \lambda_2 & c_{12} + \lambda_1 & c_{11} \\ c_{30} & 2c_{40} & c_{31} & -\lambda_3 & -\lambda_2 \\ c_{21} + \lambda_2 & c_{31} & 2c_{22} + 2\lambda_3 & c_{13} & -\lambda_1 \\ c_{12} + \lambda_1 & -\lambda_3 & c_{13} & 2c_{04} & c_{03} \\ c_{11} & -\lambda_2 & -\lambda_1 & c_{03} & 2c_{02} \end{bmatrix},$$

which λ_i *being arbitrary real numbers. The condition for the existence of a sos*

$$\begin{bmatrix}
2a_{20} & a_{30} & a_{21}+\nu_6 & a_{12}+\nu_2 & a_{11} & -\nu_7+\nu_8 & -\nu_{13} & a_{13}+\nu_3 \\
a_{30} & 2a_{40} & a_{31}+\nu_{13} & -\nu_8+\nu_9 & -\nu_6 & a_{32}+\nu_{14} & a_{41} & -\nu_{10}+\nu_{11} \\
a_{21}+\nu_6 & a_{31}+\nu_{13} & 2a_{22}+2\nu_7 & -\nu_3+\nu_4 & -\nu_2 & a_{23}+\nu_{10} & -\nu_{14} & a_{14}+\nu_5 \\
a_{12}+\nu_2 & -\nu_8+\nu_9 & -\nu_3+\nu_4 & 2a_{04}+2\nu_1 & a_{03} & -\nu_5 & -\nu_{11} & a_{05} \\
a_{11} & -\nu_6 & -\nu_2 & a_{03} & 2a_{02} & -\nu_4 & -\nu_9 & -\nu_1 \\
-\nu_7+\nu_8 & a_{32}+\nu_{14} & a_{23}+\nu_{10} & -\nu_5 & -\nu_4 & 2a_{24}+2\nu_{12} & a_{33} & a_{15} \\
-\nu_{13} & a_{41} & -\nu_{14} & -\nu_{11} & -\nu_9 & a_{33} & 2a_{42} & -\nu_{12} \\
a_{13}+\nu_3 & -\nu_{10}+\nu_{11} & a_{14}+\nu_5 & a_{05} & -\nu_1 & a_{15} & -\nu_{12} & 2a_{06}
\end{bmatrix}$$

Table 7.1: The matrix R.

representation for V, obtained as explained in Chapter 4, is therefore $Q \geq 0$.

For the derivative, we obtain after some algebra that

$$\dot{V}(x) = \begin{bmatrix} 1 \\ x_1 \\ x_1^2 \\ x_1^3 \\ x_1^4 \end{bmatrix}^T \begin{bmatrix} 0 & 0 & a_{02} & a_{03} & a_{04} & a_{05} & a_{06} \\ 0 & a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & 0 \\ a_{20} & a_{21} & a_{22} & a_{23} & a_{24} & 0 & 0 \\ a_{30} & a_{31} & a_{32} & a_{33} & 0 & 0 & 0 \\ a_{40} & a_{41} & a_{42} & 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ x_2 \\ x_2^2 \\ x_2^3 \\ x_2^4 \\ x_2^5 \\ x_2^6 \end{bmatrix}$$

where the a_i are linear functions of the c_i. For example, we have $a_{12} = -4c_{20} - c_{12} - 2c_{12} - 2c_{02}$, and $a_{42} = 0$. The full expressions are omitted for space reasons.

Writing it as a quadratic expression, we have $\dot{V}(x) = -\frac{1}{2}w^T R w$, with the vector $w = [x_1, x_1^2, x_1x_2, x_2^2, x_2, x_1x_2^2, x_1^2x_2, x_2^3]^T$. The expression for the matrix R is given in Table 7.1.

Again, here ν_i are arbitrary real parameters. If $\dot{V}$ has to be negative, then the sos *condition is $R \geq 0$. Notice that having $a_{42} = 0$ immediately implies that the multipliers $\nu_9, \nu_{11}, \nu_{12}, \nu_{13}, \nu_{14}$ and the coefficients a_{41}, a_{33} should also vanish. This way, the LMIs are considerably simplified.*

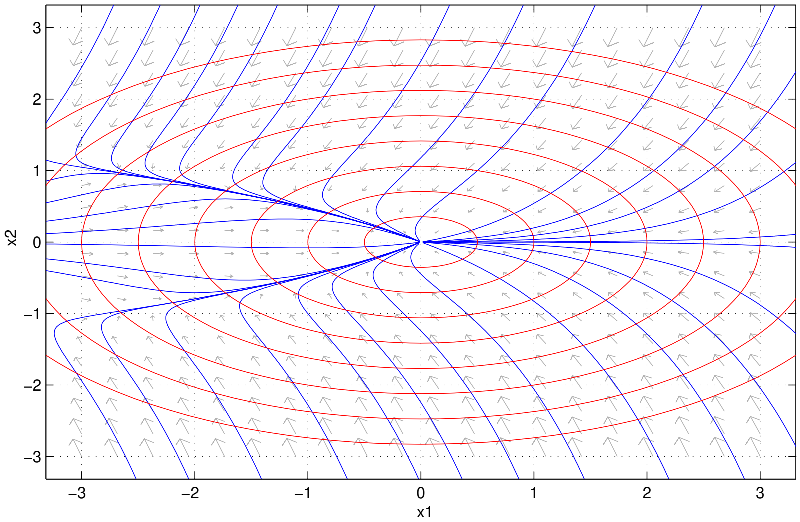

Figure 7.1: Phase plot and Lyapunov function level sets.

After solving the LMIs, it turns out that for this specific example it is even possible to pick a particularly elegant solution, given by a quadratic Lyapunov function. This can be achieved by minimizing the sum of diagonal elements corresponding to the nonquadratic terms, subject to the LMI constraints. In fact, we can choose all the c_i, λ_i *and* ν_i *equal to zero, except* $c_{20} = 1$ *and* $c_{02} = 2$*, i.e.,*

$$V(x_1, x_2) = x_1^2 + 2x_2^2.$$

In this case, we have

$$\begin{aligned} \dot{V}(x) &= (2x_1)(-x_1 - 2x_2^2) + (4x_2)(-x_2 - x_1x_2 - 2x_2^3) \\ &= -4x_2^2 - 2(x_1 + 2x_2^2)^2 \leq 0 \end{aligned}$$

In Figure 7.1 a phase plot of the vector field and the level sets of the obtained Lyapunov function are presented. It should be remarked that the considerable simplification in the final answer is not really necessary. Any feasible solution of the LMIs will provide a stability-proving Lyapunov function. $\square$

It is important to keep in mind that one of the main differences between linear and nonlinear control is that in the latter, *global behavior* is not necessarily the most important consideration: in applications, many successful nonlinear schemes are *not* globally stable. The described techniques can also be employed in this case, by testing nonnegativity only on compact regions. To this end, the Positivstellensatz based conditions, a natural generalization of S-procedure type tests, provide a useful computational approach, as we will see in the next section.

7.3 Estimates for the region of attraction

Given a Lyapunov function, consider the problem of estimating the region of attraction of a given asymptotically stable fixed point (the origin, without loss of generality). To compute such an estimate, one possible approach is the *trajectory-reversing* method [51], which uses intensive numerical simulation to propagate outwards estimates of the region of attraction, starting from a neighborhood of the stable fixed point. As an alternative, we can try to find a *positively invariant subset*, on which the time derivative of the Lyapunov function is negative [51].

A way of doing this is to solve the optimization problem

$$\gamma_0 := \inf_{x,y \in \mathbb{R}^n} V(x,y) \qquad \text{subject to} \begin{cases} \dot{V}(x,y) = 0 \\ (x,y) \neq (0,0) \end{cases} \tag{7.2}$$

in which case the invariant subset is given by the connected component of the Lyapunov function sublevel set $\mathcal{S} := \{(x,y) \,|\, V(x,y) < \gamma_0\}$ that includes the origin. Using the machinery introduced in the previous chapters, it is possible to obtain lower bounds on γ_0, which immediately provide estimates for the attracting region.

For concreteness, consider the following system, taken from [88, Example S7]:

$$\begin{aligned} \dot{x} &= -x + y \\ \dot{y} &= 0.1x - 2y - x^2 - 0.1x^3 \end{aligned}$$

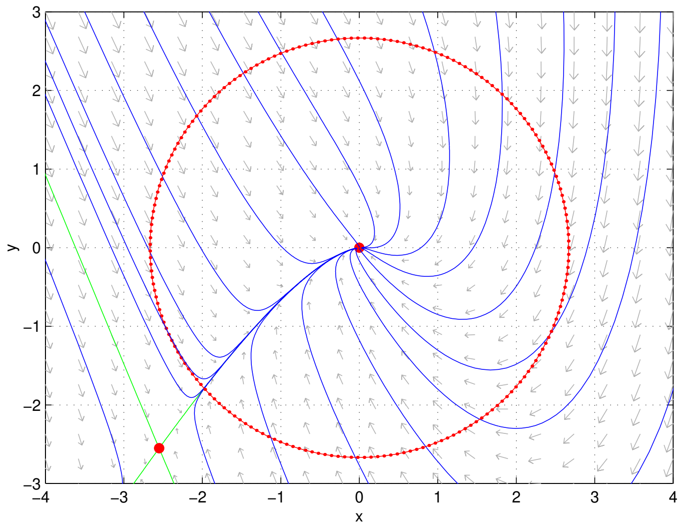

Figure 7.2: Phase plot and region of attraction.

and the Lyapunov function $V(x,y) := x^2 + y^2$. The system has three fixed points, at $(0,0), (-5 \pm \sqrt{6}, -5 \pm \sqrt{6})$.

In [88], a lower bound estimate for $\gamma_0^{1/2}$ equal to 1.911 is presented. In order to use the methods described previously, we can consider the condition:

$$(V(x,y) - \gamma)(x^2 + y^2) + (p_1 + p_2 x + p_3 y + p_4 xy)\dot{V}(x,y) \qquad \text{is a sum of squares.}$$

If this holds for feasible γ and p_i, then obviously the bound $V(x,y) \geq \gamma$ holds for every (x,y) satisfying the constraints in (7.2). Solving the LMIs, the optimal value of $\gamma^{1/2}$ is 2.66673. This actually corresponds to the optimal value, since a feasible solution achieving this bound is given by

$$x \approx -2.26099, \qquad y \approx -1.413999.$$

In Figure 7.2 the vector field and the optimal level set of the Lyapunov function are presented. The fixed points shown are the origin and $(-2.55, -2.55)$. Notice that the result is tight: one of the trajectories is in fact tangent to the curve $V(x,y) = \gamma$.

7.4 Robust bifurcation analysis

Dynamical systems possess some uniquely nonlinear phenomena, such as *local bifurcations* [92]. A vector field $f(x, \mu)$ is said to undergo a *fixed point bifurcation* when the flow around a fixed point x_0 changes qualitatively, when a parameter μ crosses some critical value μ_0.

Local bifurcations are very important in natural and engineered systems. For example, in power systems it has been argued that the significant practical problem of voltage collapse in fact has its origin in a *saddle-node* bifurcation, where the operating equilibrium point suddenly disappears as a consequence of a change in the parameters (for example, reactive load). From a practical viewpoint, is it absolutely critical to recognize such situations, and choose nominal values for the parameters that are far away from the hypersurface where bifurcations occur.

Despite its practical importance, there does not seem to be many systematic approaches to the problem of computing bifurcation margins. In reference [30], Dobson proposed two methods for computing *locally closest* bifurcations to a given set of nominal parameters. These methods (iterative and direct) aim to numerically solve the equations characterizing the closest point in the bifurcation surface. However, the problem with this approach is exactly the same as in standard robustness analysis: what we really need in practice is some way of *guaranteeing a minimum distance* (or safety margin) to a singularity, not just feasible solutions. In other words, if we find a bifurcation "nearby," then we need to be absolutely sure that there are no other points that are even closer. The results in [30, 2] do not fully address this issue: a Monte Carlo approach is employed, where the optimization is restarted from multiple initial conditions.

The techniques developed in previous chapters can be applied to rigorously prove bounds on the distance to the bifurcation surface. The conditions for a vector field $f(x, \mu)$ to have a saddle-node bifurcation at (x_0, μ_0) are [38]:

$$\begin{aligned} f &= 0 & \qquad w^* D_\mu f &\neq 0 \\ w^* D_x f &= 0 & \qquad w^* D_x^2 f(v, v) &\neq 0 \end{aligned}$$

where v, w are the right and left eigenvectors, respectively, of the jacobian $J := D_x f$, corresponding to the simple eigenvalue zero. The two conditions on the left-hand side correspond to the singularity of the jacobian at the fixed point, and the ones on the right-hand side are generic transversality requirements.

As we can see, in the polynomial (or rational) case, the set where bifurcations occur is semialgebraic, since it can be characterized in the form described by Theorem 4.4. Therefore, our methods are immediately applicable to this problem.

The example below also demonstrates another issue: even if the problem contains *nonalgebraic* elements, such as trigonometric functions, it might be possible in certain cases to get around this by changing variables.

The following system, from [30], is a model of a simple power system with a generator slack bus, lossless lines, and a load with real and reactive powers P, Q, respectively. The state variables are (α, V), where $V e^{j\alpha}$ is the load voltage phasor, and the bifurcation parameters μ are (P, Q). The equations that determine the system equilibria are:

$$\begin{aligned} 0 &= -4V \sin\alpha - P \\ 0 &= -4V^2 + 4V \cos\alpha - Q \end{aligned}$$

The system operates at a nominal solution, given by the values $(P_0, Q_0, \alpha_0, V_0) = (0.5, 0.3, -0.1381, 0.9078)$, and shown in Figure 7.3. As the loads P, Q change, the equilibrium moves around, and can eventually disappear. In this problem, we compute "safety margins" for the allowable variations in the loads, that guarantee that a saddle-node bifurcation is not reached.

To handle the trigonometric functions, define $x := \sin\alpha$, $y := \cos\alpha$. The first transversality condition is identically satisfied. If for simplicity we do not consider the second generic transversality condition, the equations we need to solve are:

$$\begin{aligned} f_1 &:= x^2 + y^2 - 1 = 0 \\ f_2 &:= -4Vx - P = 0 \end{aligned}$$

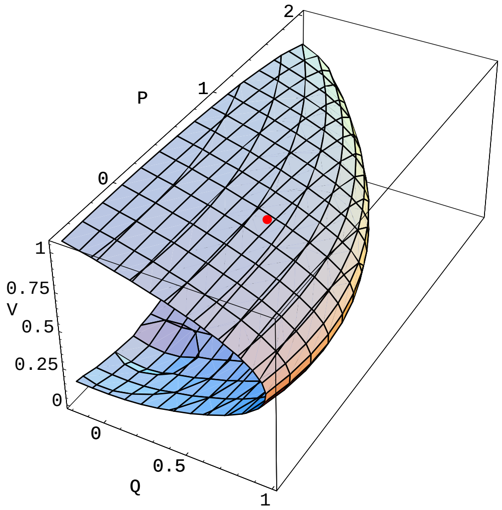

Figure 7.3: Equilibrium points surface and nominal operating point.

$$
\begin{aligned}
f_3 &:= -4V^2 + 4Vy - Q = 0 \\
f_4 &:= \det J = -16V(x^2 + y^2 - 2Vy) = 0
\end{aligned}
$$

Since we are not interested in the case where the voltage is zero, we factor out the first term $-16V$ in the last equation, obtaining:

$$f_4' := (x^2 + y^2 - 2Vy) = 0$$

We would like, therefore, to minimize the function

$$J(P, Q) := (P - 0.5)^2 + (Q - 0.3)^2$$

subject to the equalities above.

Instead of dealing with the problem as a whole, since we have equality constraints in this case it is easier to eliminate the variables that do not appear in the objective. In other words, we will only care about the constraints we can generate that are in the *elimination ideal*, i.e., $\langle f_1, f_2, f_3, f_4' \rangle \cap \mathbb{R}[P, Q]$. The only reason we do this is because of computational efficiency, but is not strictly necessary to do so from a

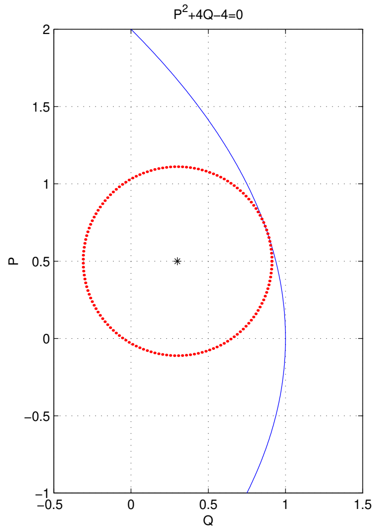

Figure 7.4: Curve where saddle-node bifurcations occur, and computed distance from the nominal equilibrium.

theoretical viewpoint.

An automatic way of generating this ideal is using Gröbner bases, when we use a lexicographic degree monomial ordering [26, 64]. The elimination ideal has only one polynomial, $P^2 + 4Q - 4$. This corresponds to the curve where saddle-node bifurcations occur; see Figure 7.4. Therefore, to compute a lower bound on the distance from the nominal equilibrium to the closest saddle-node bifurcation, we can find the maximum γ^2 that verifies the condition:

$$(P - 0.5)^2 + (Q - 0.3)^2 - \gamma^2 + \lambda(P, Q)(P^2 + 4Q - 4) \qquad \text{is a sum of squares.}$$

In this case, it is sufficient to pick $\lambda(P, Q)$ constant, and we obtain an optimal value of $\gamma^2 \approx 0.3735$, with $\lambda \approx -0.2883$.

To verify that the restriction to the elimination ideal is not crucial, we can easily verify that multiplying the expressions

$$8V^2 + 4 - 2Q, \quad -P + 4Vx, \quad -4Vy - 2 \quad -4 + 8V^2 + 2Q$$

by f_1, f_2, f_3 and f_4' respectively, and adding, we obtain the valid constraint $P^2 + 4Q - 4 = 0$. Therefore, the only difference in that case would be the need of using nonconstant multipliers.

Though not guaranteed *a priori* by the method, in this case again we obtain a bound that is exact. As seen in the figure, there exists a solution of the equations that achieves the computed value of γ^2, corresponding to $P \approx 0.7025$, $Q \approx 0.8766$.

7.5 Zero dynamics stability

When studying the global feedback linearization procedure for nonlinear systems [46], a problem that appears is that of the *zero dynamics* stability. This question, that extends the linear concepts of *minimum phase*, deals with the stability of the system, when the outputs is constrained to be identically zero. At least in certain cases, we can apply the techniques to this kind of problems. A simple example follows.

Consider the following system, from [46, Example 4.3.4].

$$\begin{aligned} \dot{x}_1 &= x_3 - x_2^3 \\ \dot{x}_2 &= -x_2 - u \\ \dot{x}_3 &= x_1^2 - x_3 + u \\ y &= x_1 \end{aligned} \tag{7.3}$$

To prove stability of the zero dynamics, it is sufficient to find a positive definite Lyapunov function $V(x)$, that has a negative derivative along the trajectories of (7.3), with the constraint $y \equiv 0$.

Such a V can be obtained by solving the LMIs corresponding to

$$V(x) + \lambda_1(x) y, \quad -\dot{V}(x) + \lambda_2(x) y \qquad \text{are sums of squares}$$

A simple solution is given by $V(x_1, x_2, x_3) := \frac{1}{2}(x_2 + x_3)^2$, since in this case the second expression above reduces to:

$$-\dot{V}(x) + x_1(x_2 + x_3)y = (x_2 + x_3)^2 \geq 0.$$

Since the Lyapunov function and its derivative are not strictly positive, we need to be a bit careful. However, after invoking LaSalle's invariance principle, the stability of the zero dynamics can be established.

7.6 Synthesis

As we have seen, the developed methods can be applied to many nonlinear *analysis* questions. A natural question, therefore, is about the possibility of extending these results to *synthesis* problems, where we try to find stabilizing controllers that satisfy given performance criteria.

In the linear case, the usual LMI solution to stabilization problems goes along the following lines [17]: to find a stabilizing controller K, we need a Lyapunov function $V(x) := x^T P x$, $P > 0$ such that $A + BK$ is stable, i.e.,

$$P(A + BK) + (A + BK)^T P < 0. \tag{7.4}$$

This condition is not affine in both P and K. By multiplying the expression above by $Q := P^{-1}$, and defining a new variable $L := KQ$, we obtain:

$$(A + BK)Q + Q(A + BK)^T = AQ + BL + (AQ + BL)^T < 0, \tag{7.5}$$

which is affine in both Q and L. Since $Q > 0$, we can always find the controller K as $K = LQ^{-1}$.

However, extending this procedure to the nonlinear case does not seem feasible, at least in a reasonably straightforward way. Consider an affine nonlinear system $\dot{x} = f(x) + g(x)u$. While the condition

$$\frac{\partial V}{\partial x}(f(x) + g(x)k(x)) < 0 \tag{7.6}$$

is clearly the nonlinear equivalent of (7.4) above, there does not seem to be an efficient way of searching *simultaneously* over the Lyapunov function V and the controller $k(x)$.

We certainly expect synthesis procedures to be no easier than the corresponding analysis questions. However, the presence of additional properties, such as a triangular structure of the vector field in simple cases of backstepping [56], usually helps in the complexity reduction. The extent to which the presented results can be applied in synthesis procedures still remains to be fully determined.

7.7 Conclusions

The sum of squares decomposition is a very useful sufficient condition for positivity of a multivariable polynomial. It can be obtained at a reasonable computational cost, using LMI methods. We can combine this procedure with Positivstellensatz based tests, in order to extend the class of problems to which the methods are applicable. The obtained results and procedures constitute a sound and natural extension of standard tools in linear systems analysis. One of the big advantages of the proposed procedure is that it is a completely algorithmic procedure. All the computations can be carried through in a deterministic fashion, in polynomial time.

The basic idea of the procedure seems to be relevant in numerous questions in the systems and control area. In this chapter we presented some immediate applications, mainly dealing with the analysis of nonlinear (polynomial or rational) systems.

In conclusion, we have shown that the combination of semialgebraic geometry tools and semidefinite programming optimization is a very powerful general purpose tool for dealing with the analysis of nonlinear dynamical systems.

Chapter 8

Conclusions

To conclude, we briefly summarize our main contributions, and outline some directions of future research.

8.1 Summary

In Chapter 2, an exact characterization of the optimal solution for a class of cone preserving linear matrix inequalities was presented. The results were applied to a variant of the rank minimization problem, and to the computation of the LMI upper bound for the spherical μ problem.

The special structure of the LMIs arising from the Kalman-Yakubovich-Popov lemma was exploited in Chapter 3 in the formulation of numerically efficient algorithms. We introduced an outer approximation based procedure based on the frequency domain description and a semi-infinite programming viewpoint.

From Chapter 4 on, we developed a computational framework for semialgebraic problems, and presented applications to many different problems in systems and control. The proposed methods are extremely general, and are based, on the one hand, on results from real algebraic geometry, and on the other, on semidefinite programming.

The key idea was the use a sum of squares decomposition as a sufficient condition for nonnegativity of a multivariable polynomial. This condition can be tested in polynomial time, using LMI methods. Pairing this computational tool with the

Positivstellensatz in real algebraic geometry, we obtained very powerful generalizations of the successful methods developed in linear robustness analysis during the last decade.

The problem of matrix copositivity was analyzed, and it was shown how improved conditions can be obtained through the presented methodology. The enhanced tests were shown to be exact for a certain family of copositive quadratic forms.

In the important specific case of indefinite quadratic programming, 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 procedures.

8.2 Future research

As future research directions, it would be interesting to analyze the possibility of extending the results in Chapter 2 to more general inequalities, and unifying several results concerning closed forms solutions of LMIs.

Since many LMIs arising in systems and control theory do not possess the cone-invariance property, it is interesting to examine to what extent fast algorithms can be constructed, in the case where only part of the LMI is cone-invariant. An example of this is the computation of μ upper bounds with mixed "norm bounded" and "Frobenius" uncertainty block.

As we have shown, it is perfectly possible to immediately apply the tools developed in Chapter 4 to relatively small problems. However, an important issue is certainly the computational feasibility of applying these relaxations to large scale instances. Though in principle all the relaxations are polynomial time algorithms, we should realize this in only a coarse characterization: in practice, other considerations such as memory usage, or the actual execution time are perhaps more relevant. In this respect, an important factor is the choice of data representation: dense coding of polynomials is clearly unsuitable for large scale problems, and alternatives such as sparse representations or straight-line programs are needed.

For this reasons, more research is needed in the implementation aspects, especially on the issue of exploiting additional problem structure. Some recent interesting approaches, such as the work in reference [21] on the standard MAX CUT relaxation, show that there is lot of room for improvement, especially when working in specific problem classes.

A natural question in the sum of squares decomposition, for instance, is if we really need to introduce additional variables to cast the problem as an LMI, or is it possible to solve the problem directly in the original space. After all, the set of sum of squares polynomials is a "nice" closed convex cone. In this direction, in [66] it has been shown that the natural self-concordant barrier for the cone of positive definite univariate polynomials is essentially optimal.

In the general Positivstellensatz approach, another important practical issue lies in the "customization" of the structure of the polynomial multipliers to that of the original problem. For example, in the enhanced relaxations of Chapter 6, the homogeneous formulation presented in Theorem 6.1 seems natural. Also, as we have seen in Chapter 5, in the copositivity problem some of the multipliers can be chosen without loss of generality to be identically zero. Therefore, for computational reasons it would be interesting to characterize a convenient family of possible multipliers, to a higher level of detail than just degree bounds. In this direction, the Newton polytope ideas used in sparse versions of the Nullstellensatz (see for example [87, 84]) might prove to be useful.

Appendix A

Algebra review

For the convenience of the reader, we present in this appendix some standard background material on abstract algebra. Most of the definitions are from [57, 26, 12].

Definition A.1 *A* group *consists of a set G and a binary operation "$\cdot$" defined on G, for which the following conditions are satisfied:*

1. *Associative: $(a \cdot b) \cdot c = a \cdot (b \cdot c)$, for all $a, b, c \in G$.*
2. *Identity: There exist $1 \in G$ such that $a \cdot 1 = 1 \cdot a = a$, for all $a \in G$.*
3. *Inverse: Given $a \in G$, there exists $b \in G$ such that $a \cdot b = b \cdot a = 1$.*

For example, the integers $\mathbb{Z}$ form a group under addition, but not under multiplication. Another example is the set $GL(n, \mathbb{R})$ of real nonsingular $n \times n$ matrices, under matrix multiplication.

If we drop the condition on the existence of an inverse, we obtain a *monoid*. Note that a monoid always has at least one element, the identity. As an example, given a set S, then the set of all strings of elements of S is a monoid, where the monoid operation is string concatenation and the identity is the empty string λ. Another example is given by $\mathbb{N}_0$, with the operation being addition (in this case, the identity is the zero).

Definition A.2 *A* field *consists of a set k and two binary operations "$\cdot$" and "$+$", defined on k, for which the following conditions are satisfied:*

1. *Associative:* $(a+b)+c = a+(b+c)$ *and* $(a \cdot b) \cdot c = a \cdot (b \cdot c)$*, for all* $a, b, c \in k$.

2. *Commutative:* $a + b = b + a$ *and* $a \cdot b = b \cdot a$*, for all* $a, b \in k$.

3. *Distributive:* $a \cdot (b + c) = a \cdot b + a \cdot c$*, for all* $a, b, c \in k$.

4. *Identities: There exist* $0, 1 \in k$ *such that* $a + 0 = a \cdot 1 = a$*, for all* $a \in k$.

5. *Additive inverse: Given* $a \in k$*, there exists* $b \in k$ *such that* $a + b = 0$.

6. *Multiplicative inverse: Given* $a \in k, a \neq 0$*, there exists* $c \in k$ *such that* $a \cdot c = 1$.

Some commonly used fields are the rationals $\mathbb{Q}$, the reals $\mathbb{R}$ and the complex numbers $\mathbb{C}$. There are also Galois or finite fields (the set k has a finite number of elements), such as $\mathbb{Z}_p$, the set of integers modulo p, where p is a prime. Another important field if given by $k(x_1, \ldots, x_n)$, the set of *rational functions* with coefficients in the field k, with the natural operations.

Dropping the existence of multiplicative inverses, we obtain commutative rings.

Definition A.3 *A* commutative ring *(with identity) consists of a set k and two binary operations "$\cdot$" and "$+$", defined on k, for which the following conditions are satisfied:*

1. *Associative:* $(a+b)+c = a+(b+c)$ *and* $(a \cdot b) \cdot c = a \cdot (b \cdot c)$*, for all* $a, b, c \in k$.

2. *Commutative:* $a + b = b + a$ *and* $a \cdot b = b \cdot a$*, for all* $a, b \in k$.

3. *Distributive:* $a \cdot (b + c) = a \cdot b + a \cdot c$*, for all* $a, b, c \in k$.

4. *Identities: There exist* $0, 1 \in k$ *such that* $a + 0 = a \cdot 1 = a$*, for all* $a \in k$.

5. *Additive inverse: Given* $a \in k$*, there exists* $b \in k$ *such that* $a + b = 0$.

Any field is obviously a commutative ring. Additional examples are the integers $\mathbb{Z}$, and the polynomial ring $k[x_1, \ldots, x_n]$, i.e., the set of polynomials in n variables, with coefficients in k (see Definition 4.1 in page 38).

A commutative ring is called an *integral domain* if it has no *zero divisors*, i.e. $a \neq 0, b \neq 0 \Rightarrow a \cdot b \neq 0$. Any field is also an integral domain. Two examples of rings that are not integral domains are the set of matrices $\mathbb{R}^{n\times n}$, and the set of integers modulo n, when n is a composite number (with the usual operations). If k is an integral domain, then so is $k[x_1, \ldots, x_n]$.

Definition A.4 *A field k is* algebraically closed *if every nonconstant polynomial in $k[x_1, \ldots, x_n]$ has a root in k.*

The Fundamental Theorem of Algebra shows that $\mathbb{C}$ is an algebraically closed field. This is not the case of $\mathbb{R}$, since for example the polynomial $x^2 + 1$ does not have any real root. To deal with the case when the base field is not algebraically closed, the *Artin-Schreier* theory of *formally real fields* was introduced, see equation (4.12) in Chapter 4. A related important notion is that of an *ordered* field:

Definition A.5 *A field k is said to be* ordered *if a relation $>$ is defined on k, that satisfies*

1. *If $a, b \in k$, then either $a > b$ or $a = b$ or $b > a$.*
2. *If $a > b$, $c \in k$, $c > 0$ then $ac > bc$.*
3. *If $a > b$, $c \in k$, then $a + c > b + c$.*

A field can be ordered if and only if it is formally real.

We consider next *ideals*, which are subrings with an "absorbent" property:

Definition A.6 *Let R be a commutative ring. A subset $I \subset R$ 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 $b \in R$, then $a \cdot b \in I$.*

A simple example of an ideal is the set of even integers, considered as a subset of the integer ring $\mathbb{Z}$.

To introduce another important example of ideals, we need to define the concept of an algebraic variety as the zero set of a set of polynomial equations:

Definition A.7 *Let k be a field, and let $f_1, \ldots, f_s$ be polynomials in $k[x_1, \ldots, x_n]$. Let the set* $\mathbf{V}$ *be*

$$\mathbf{V}(f_1, \ldots, f_s) = \{(a_1, \ldots, a_n) \in k^n : f_i(a_1, \ldots, a_n) = 0 \qquad \forall 1 \leq i \leq s\}.$$

We call $\mathbf{V}(f_1, \ldots, f_s)$ *the* affine variety *defined by $f_1, \ldots, f_s$.*

Then, the set of polynomials that vanish in a given variety, i.e.,

$$\mathbf{I}(V) = \{f \in k[x_1, \ldots, x_n] : f(a_1, \ldots, a_n) = 0 \qquad \forall (a_1, \ldots, a_n) \in V\},$$

is an ideal, called the *ideal of* V.

By Hilbert's Basis Theorem [26], $k[x_1, \ldots, x_n]$ is a *Noetherian* ring, i.e., every ideal $I \subset k[x_1, \ldots, x_n]$ is finitely generated. In other words, there always exists a finite set $f_1, \ldots, f_s \in k[x_1, \ldots, x_n]$ such that for every $f \in I$, we can find $g_i \in k[x_1, \ldots, x_n]$ that verify $f = \sum_{i=1}^{s} g_i f_i$.

We also define the *radical* of an ideal:

Definition A.8 *Let $I \subset k[x_1, \ldots, x_n]$ be an ideal. The* radical *of I, denoted $\sqrt{I}$, is the set*

$$\{f \mid f^k \in I \text{ for some integer } k \geq 1\}.$$

It is clear that $I \subset \sqrt{I}$, and it can be shown that $\sqrt{I}$ is also a polynomial ideal.

Bibliography

[1] N. I. Akhiezer. *The classical moment problem.* Hafner Publishing Company, New York, 1965.

[2] F. Alvarado, I. Dobson, and Y. Hu. Computation of closest bifurcations in power systems. *IEEE Transactions on Power Systems*, 9(2):918–928, 1994.

[3] B. D. O. Anderson, N. K. Bose, and E. I. Jury. Output feedback stabilization and related problems—solution via decision methods. *IEEE Transactions on Automatic Control*, 20:53–66, 1975.

[4] 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/ABSTRACTS.html`, 1999.

[5] M. F. Anjos and H. Wolkowicz. A tight semidefinite relaxation of the cut polytope. Preprint, available at `http://orion.math.uwaterloo.ca/~hwolkowi/henry/reports/ABSTRACTS.html`, 2000.

[6] V. J. D. Baston. Extreme copositive quadratic forms. *Acta Arithmetica*, XV:319–327, 1969.

[7] 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.

[8] C. Berg. The multidimensional moment problem and semigroups. In H. J. Landau, editor, *Moments in Mathematics*, volume 37 of *Proceedings of Symposia in Applied Mathematics*, pages 110–124. American Mathematical Society, 1987.

[9] A. Berman and R. J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences.* Academic Press, New York, 1979.

[10] G. Birkhoff. Linear transformations with invariant cones. *Amer. Math. Monthly*, 72:274–276, 1967.

[11] V. Blondel and M. Gevers. Simultaneous stabilizability of 3 linear systems is rationally undecidable. *Mathematics of Control, Signals, and Systems*, 6(2):135–145, 1993.

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

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

[14] 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.

[15] S. Boyd and V. Balakrishnan. A regularity result for the singular values of a transfer matrix and a quadratically convergent algorithm for computing its $\mathcal{L}_\infty$ norm. *Systems & Control Letters*, 15(1):1–7, 1990.

[16] S. Boyd, V. Balakrishnan, and P. Kabamba. A bisection method for computing the $\mathcal{H}_\infty$ norm of a transfer matrix and related problems. *Mathematics of Control, Signals, and Systems*, 2:207–219, 1989.

[17] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*, volume 15 of *Studies in Applied Mathematics.* SIAM, Philadelphia, PA, 1994.

[18] S. Boyd and S.-P. Wu. *SDPSOL: A Parser/Solver for Semidefinite Programs with Matrix Structure. User's Guide, Version Beta.* Stanford University, Mar. 1996.

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

[20] N. A. Bruinsma and M. Steinbuch. A fast algorithm to compute the $\mathcal{H}_\infty$-norm of a transfer function. *Systems & Control Letters*, 14:287–293, 1990.

[21] 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.gz`, 1998.

[22] S. R. Buss and T. Pitassi. Good degree bounds on Nullstellensatz refutations of the induction principle. *Journal of computer and system sciences*, 57(2):162–171, 1998.

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

[24] 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.

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

[26] 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.

[27] J. de Loera and F. Santos. Correction of errors to: An effective version of Pólya's theorem on positive definite forms. Available at `http://www.matesco.unican.es/~santos/Articulos/JPAAcorr.ps.gz`, to appear in *J. Pure Appl. Algebra.*

[28] J. de Loera and F. Santos. An effective version of Pólya's theorem on positive definite forms. *J. Pure Appl. Algebra*, 108:231–240, 1996.

[29] 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.

[30] I. Dobson. Computing a closest bifurcation instability in multidimensional parameter space. *Nonlinear Science*, 3:307–327, 1993.

[31] P. Dorato, W. Yang, and C. Abdallah. Robust multi-objective feedback design by quantifier elimination. *J. Symbolic Computation*, 24:153–159, 1997.

[32] M. Fu. The real structured singular value is hardly approximable. *IEEE Transactions on Automatic Control*, 42(9):1286–1288, 1997.

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

[34] P. Gahinet and P. Apkarian. A linear matrix inequality approach to $\mathcal{H}_\infty$ control. *International Journal of Robust and Nonlinear Control*, 4(4):421–448, 1994.

[35] P. Gahinet, A. Nemirovski, A. J. Laub, and M. Chilali. *LMI Control Toolbox*. The MathWorks Inc., Natick, Mass., May 1995.

[36] M. R. Garey and D. S. Johnson. *Computers and Intractability: A guide to the theory of NP-completeness*. W. H. Freeman and Company, 1979.

[37] 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.

[38] J. Guckenheimer and P. Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, 1983.

[39] M. Hall and M. Newman. Copositive and completely copositive quadratic forms. *Proc. Camb. Phil. Soc.*, 59:329–339, 1963.

[40] G. H. Hardy, J. E. Littlewood, and G. Pólya. *Inequalities.* Cambridge University Press, 1967.

[41] 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.

[42] D. S. Hochbaum, editor. *Approximation algorithms for NP-hard problems.* PWS Publishing Company, 1997.

[43] D. A. Holton and J. Sheehan. *The Petersen graph*, volume 7 of *Australian Mathematical Society Lecture Series.* Cambridge University Press, 1993.

[44] R. A. Horn and C. R. Johnson. *Matrix Analysis.* Cambridge University Press, Cambridge, 1995.

[45] Y. Huang. *Nonlinear Optimal Control: An Enhanced Quasi-LPV Approach.* PhD thesis, California Institute of Technology, 1998.

[46] A. Isidori. *Nonlinear Control Systems.* Springer-Verlag, Berlin, third edition, 1997.

[47] M. Jirstrand. Nonlinear control system design by quantifier elimination. *J. Symbolic Computation*, 24:137–152, 1997.

[48] M. Johansson. *Piecewise Linear Control Systems.* PhD thesis, Lund Institute of Technology, 1999.

[49] M. Johansson and A. Rantzer. Computation of piecewise quadratic lyapunov functions for hybrid systems. *IEEE Transactions on Automatic Control*, 43(4):555–559, 1998.

[50] U. Jönsson and A. Rantzer. Duality bounds in robustness analysis. *Automatica*, 33(10):1835–1844, 1997.

[51] H. Khalil. *Nonlinear Systems.* Macmillan Publishing Company, 1992.

[52] S. Khatri and P. A. Parrilo. Spherical μ. In *Proceedings of the American Control Conference*, 1998.

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

[54] A. N. Kolmogorov and S. V. Fomin. *Introductory Real Analysis*. Dover, 1970.

[55] M. G. Krein and M. A. Rutman. Linear operators leaving invariant a cone in a Banach space. *Usp. Mat. Nauk (N.S)*, 3:3–95, 1948. (English transl.: Amer. Math. Soc. Transl. Ser. I, vol. 10, (1962), 199–325).

[56] M. Krstić, I. Kanellakopoulos, and P. V. Kokotović. *Nonlinear and Adaptive Control Design*. John Wiley & Sons, 1995.

[57] S. Lang. *Algebra*. Addison-Wesley, 1971.

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

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

[60] A. Megretski, C.-Y. Kao, U. Jönsson, and A. Rantzer. *A Guide To IQCβ: Software for Robustness Analysis*, June 1998.

[61] A. Megretski and A. Rantzer. System analysis via Integral Quadratic Constraints. *IEEE Transactions on Automatic Control*, 42(6):819–830, 1997.

[62] M. Mesbahi. On the rank minimization problem and its control applications. *Systems & Control Letters*, 33:31–36, 1998.

[63] M. Mesbahi and G. P. Papavassilopoulos. On the rank minimization problem over a positive semidefinite linear matrix inequality. *IEEE Transactions on Automatic Control*, 42(2):239–243, 1997.

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

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

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

[67] Y. E. Nesterov and A. Nemirovski. *Interior point polynomial methods in convex programming*, volume 13 of *Studies in Applied Mathematics*. SIAM, Philadelphia, PA, 1994.

[68] E. E. Osborne. On preconditioning of matrices. *Journal of the Assoc. for Comp. Mach.*, 7:338–349, 1960.

[69] A. Packard and J. C. Doyle. The complex structured singular value. *Automatica*, 29(1):71–109, 1993.

[70] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Dover, 1998.

[71] M. Petkovsěk, H. S. Wilf, and D. Zeilberger. *A=B*. A.K. Peters Ltd., 1996.

[72] E. Polak. *Optimization: Algorithms and Consistent Approximations*, volume 124 of *Applied Mathematics Sciences*. Springer, 1997.

[73] V. Powers and B. Reznick. A new bound for Pólya's theorem with applications to polynomials positive on polyhedra. Preprint, available at `http://www.mathcs.emory.edu/~vicki/preprint/polya.ps`, 2000.

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

[75] 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.

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

[77] A. Rantzer. On the Kalman-Yakubovich-Popov lemma. *Systems & Control Letters*, 28:7–10, 1996.

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

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

[80] B. Reznick. Some concrete aspects of Hilbert's 17th problem. Preprint, available at `http://www.math.uiuc.edu/Reports/reznick/98-002.html`, 1999.

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

[82] 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.

[83] 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).

[84] M. Sombra. A sparse effective Nullstellensatz. *Advances in Applied Mathematics*, 22:271–295, 1999.

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

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

[87] B. Sturmfels. Polynomial equations and convex polytopes. *American Mathematical Monthly*, 105(10):907–922, 1998.

[88] A. Tesi, F. Villoresi, and R. Genesio. On the stability domain estimation via a quadratic Lyapunov function: convexity and optimality properties for polynomial systems. *IEEE Transactions on Automatic Control*, 41(11):1650–1657, 1996.

[89] O. Toker. On the algorithmic unsolvability of some stability problems for hybrid systems. *IEEE Transactions on Automatic Control*, 45(1):165–168, 2000.

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

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

[92] S. Wiggins. *Introduction to applied nonlinear dynamical systems and chaos.* Springer-Verlag, New York, 1990.

[93] J. C. Willems. Least squares stationary optimal control and the algebraic Riccati equation. *IEEE Transactions on Automatic Control*, 16(6):621–634, 1971.