Semidefinite Programming Relaxations and Algebraic Optimization in Control

Pablo A. Parrilo[1] and Sanjay Lall[2*]

[1]Automatic Control Laboratory, Swiss Federal Institute of Technology, CH-8092 Zürich, Switzerland; [2]Department of Aeronautics and Astronautics, Stanford University, Stanford CA 94305-4035, USA

We present an overview of the essential elements of semidefinite programming as a computational tool for the analysis of systems and control problems. We make particular emphasis on general duality properties as providing suboptimality or infeasibility certificates. Our focus is on the exciting developments which have occured in the last few years, including robust optimization, combinatorial optimization, and algebraic methods such as sum-of-squares. These developments are illustrated with examples of applications to control systems.

1. Introduction

This paper, prepared as part of a mini-course to be given at the 2003 European Control Conference, presents a self-contained treatment of the basics of semidefinite programming, as well as an overview of some recent developments in convex optimization and their application to control systems. There have been significant advances in theory, numerical approaches, and applications in the past few years, with developments including robust optimization, combinatorial optimization and integer programming, and algebraic methods. Our objective in the course, and in this paper, is to highlight some of these tremendously exciting developments, with an emphasis on control applications and examples.

The roots of semidefinite programming can be traced back to both control theory and combinatorial optimization, as well as the more classical research on optimization of matrix eigenvalues. We are fortunate that many excellent works dealing with the development and applications of semidefinite programming (SDP) are available. In particular, we mention the well-known work of Vandenberghe and Boyd [60] as a wonderful survey of the basic theory and initial applications, and the handbook [64] for a comprehensive treatment of the many aspects of the subject. Other survey works, covering different complementary aspects are the early work by Alizadeh [1], Goemans [22], as well as the more recent ones due to Todd [58], and Laurent and Rendl [33]. The upcoming book [11] presents a beautiful and clear exposition of the theory, numerical approaches and broad applications of convex programming. Other works dealing more specifically with semidefinite programming for a control audience are the classical research monograph [10] and the subsequent book [21]. Excellent available

Correspondence and offprint requests to: P.A. Parrilo, Automatic Control Laboratory, Swiss Federal Institute of Technology, CH-8092 Zürich, Switzerland.
E-mail: parrilo@control.ee.ethz.ch
*E-mail: lall@stanford.edu

references that collect results at the different stages of the development of SDPs from a control perspective are the review paper of Doyle et al. [16] and the lecture notes by Scherer and Weiland [53].

1.1. Notation

We make use of the following notation. The *non-negative orthant* in $\mathbb{R}^n$ is denoted by $\mathbb{R}^n_+$, defined by

$$\mathbb{R}^n_+ = \{x \in \mathbb{R}^n \mid x_i \geq 0 \quad \text{for all } i = 1, \ldots, n\},$$

and the *positive orthant* $\mathbb{R}^n_{++}$ is

$$\mathbb{R}^n_{++} = \{x \in \mathbb{R}^n \mid x_i > 0 \quad \text{for all } i = 1, \ldots, n\}.$$

The nonnegative orthant is a closed convex cone. It defines a partial ordering on $\mathbb{R}^n$ as follows. For x, $y \in \mathbb{R}^n$ we write $x \geq y$ to mean $x - y \in \mathbb{R}^n_+$. Similarly $x > y$ means $x - y \in \mathbb{R}^n_{++}$. The set of real symmetric $n \times n$ matrices is denoted $\mathbb{S}^n$. A matrix $A \in \mathbb{S}^n$ is called *positive semidefinite* (PSD) if $x^T A x \geq 0$ for all $x \in \mathbb{R}^n$, and is called *positive definite* if $x^T A x > 0$ for all non-zero $x \in \mathbb{R}^n$. The set of PSD matrices is denoted $\mathbb{S}^n_+$ and the set of positive definite matrices is denoted by $\mathbb{S}^n_{++}$. Both $\mathbb{S}^n_+$ and $\mathbb{S}^n_{++}$ are convex cones. The closed convex cone $\mathbb{S}^n_+$ induces a partial order on $\mathbb{S}^n$, and we write $A \succeq B$ to mean $A - B \in \mathbb{S}^n_+$. Similarly $A \succ B$ means $A - B \in \mathbb{S}^n_{++}$. Both the inequality $\geq$ defined on $\mathbb{R}^n$ and the inequality $\succeq$ defined on $\mathbb{S}^n$ are called *generalized inequalities*. We also use the notation $A \leq 0$ to mean $-A \geq 0$, and similarly for $<$, $\prec$, and $\preceq$. A set $S \subset \mathbb{R}^n$ is called *affine* if $\lambda x + (1 - \lambda)y \in S$ for all x, $y \in S$ and all $\lambda \in \mathbb{R}$. It is called *convex* if $\lambda x + (1 - \lambda)y \in S$ for all x, $y \in S$ and all $\lambda \in [0, 1]$. The ring of polynomials in n variables with real coefficients is denoted $\mathbb{R}[x_1, \ldots, x_n]$.

2. Formulation of Optimization Problems

2.1. Optimization and Feasibility Problems

In this paper we consider two main types of problems; *feasibility problems* and *optimization problems*. In a feasibility problem, we typically have a system of inequalities defined by functions $f_i : \mathbb{R}^n \to \mathbb{R}$ for $i = 1, \ldots, m$. We are then interested in the set

$$P = \{x \in \mathbb{R}^n \mid f_i(x) \leq 0 \text{ for all } i = 1, \ldots, m\}. \tag{1}$$

Typically we would like to know if the set P is non-empty; i.e., if there exists $x \in \mathbb{R}^n$ which simultaneously satisfies all of the inequalities. Such a point is called a

feasible point and the set P is then called *feasible*. If P is not feasible, it is called *infeasible*. If P is known to be feasible, it is a further and typically more computationally difficult question to find a feasible x. Feasibility problems may also be defined using generalized inequalities, where the functions f_i map into $\mathbb{R}^n$ or $\mathbb{S}^n$, and the inequalities are induced by arbitrary proper cones.

In an optimization problem, we have an additional function $f_0 : \mathbb{R}^n \to \mathbb{R}$, and would like to solve the following problem.

$$
\begin{aligned}
\text{minimize} \quad & f_0(x) \\
\text{subject to} \quad & x \in P.
\end{aligned}
\tag{2}
$$

There may also be additional linear equality constraints of the form $Ax = b$. These can be incorporated in the following development using minor additions, and we omit them to simplify the presentation.

2.2. Semidefinite Programming

An SDP is a generalization of a *linear program* (LP), where the inequality constraints in the latter are replaced by generalized inequalities corresponding to the cone of positive semidefinite matrices.

Concretely, an SDP in the *pure primal form* is defined as the optimization problem

$$
\begin{aligned}
\text{minimize} \quad & \text{trace}(CX) \\
\text{subject to} \quad & \text{trace}(A_i X) = b_i \quad \text{for all } i = 1, \ldots, m, \\
& X \succeq 0
\end{aligned}
\tag{3}
$$

where $X \in \mathbb{S}^n$ is the decision variable, $b \in \mathbb{R}^m$ and C, $A_1, \ldots, A_m \in \mathbb{S}^n$ are given symmetric matrices. There are m affine constraints among the entries of the PSD matrix X given by the equality relations. The crucial feature of SDPs is that the feasible set defined by the constraints above is always convex. This is therefore a convex optimization problem, since the objective function is linear.

There are several alternative ways to write an SDP. We choose the above form since it exactly parallels the standard form used for linear programming, and is currently the most commonly used in the optimization community. Much of the literature on the use of SDPs for control has used a slightly different form with free variables, corresponding either to a parametrization of the solution set (see below) or to the standard dual SDP discussed in Section 3.3. We join here the strong trend towards this unification of terminology.

We can understand the geometry of SDPs in the following way. For a square matrix X, let X_k be the

$k \times k$ submatrix consisting of the first k rows and columns of X. Then

$$X > 0 \quad \text{if and only if } \det(X_k) > 0$$
$$\text{for all } k = 1, \ldots, n.$$

That is, in principle we can replace the positivity condition in Eq. (3) by a set of polynomial inequalities in the entries of X. Note that, in the case of nonstrict inequalities, the condition must be modified to include all minors, not just the principal ones. Sets defined by polynomial inequalities and equations are called *semialgebraic sets*. We therefore see that the feasible set defined by an SDP is a semialgebraic set. This construction is not too useful practically; however we shall see later that the reverse construction (converting a semialgebraic set to an SDP) is very useful. Unlike the LP case where the feasible sets are polyhedra, the feasible sets of SDPs (called *spectrahedra* in [48]) can have *curved* faces, together with sharp corners where faces intersect. For this reason, or equivalently, because of nondifferentiability of the eigenvalues, problem (3), while convex, is not *smooth*.

An example is as follows. Consider the feasible set defined by $X \succ 0$ and the equality constraints

$$X_{11} - X_{12} + X_{22} = 7, \qquad X_{33} - X_{12} + X_{22} = 4,$$
$$X_{13} = 1, \qquad X_{23} = 0.$$

These four affine constraints define a two-dimensional matrix subspace, since $\dim(\mathbb{S}^3) = 6$, and $6 - 4 = 2$. For easier visualization, we find a parametrization of this two-dimensional subspace in terms of the free variables (x, y), obtaining that the feasible set is

$$\left\{ X \in \mathbb{S}^n \mid X = \begin{bmatrix} 3 - x & -(x + y) & 1 \\ -(x + y) & 4 - y & 0 \\ 1 & 0 & -x \end{bmatrix}, X \succ 0, x, y \in \mathbb{R} \right\}. \tag{4}$$

This set may readily be expressed as an SDP of the form (3). By looking at the determinants of its principal submatrices, we have X satisfies Eq. (4) if and only if the parameters x and y simultaneously satisfy the polynomial inequalities:

$$3 - x > 0, \quad \text{(C)}$$
$$(3 - x)(4 - y) - (x + y)^2 > 0, \quad \text{(A)}$$
$$-x((3 - x)(4 - y) - (x + y)^2) - (4 - y) > 0. \quad \text{(B)}$$

The feasible set of x and y is shown in Fig. 1.

We have seen that the feasible sets of SDPs are convex and semialgebraic. A natural question is therefore whether every convex semialgebraic set is


Fig. 1. Example feasible set for an SDP.

representable by an SDP. Without using additional variables this is not always possible, as there are nontrivial obstructions that must be avoided (for instance, exposed faces [48]). We refer the reader to the recent result by Helton and Vinnikov [27] for a complete answer in the two-dimensional case. The results in Section 7 can be used to *constructively* build arbitrarily close approximations.

2.3. Other Classes of Optimization Problems

Semidefinite programs define a fairly wide class of optimization problems. However, it is sometimes convenient to further refine the classification. This leads to significant computational benefits, since some subclasses of SDP possess highly efficient specialized algorithms. This also provides analytical benefits, since in this case further algebraic or geometric properties may provide additional insight (for instance, the polyhedrality of LP).

The language of *conic programming* [39] provides a convenient framework to analyze these issues. By posing an abstract optimization problem of the form

$$\begin{aligned} \text{minimize} \quad & \langle C, X \rangle \\ \text{subject to} \quad & AX = b, \\ & X \in K, \end{aligned}$$

where X is an element of a vector space V, and $\mathcal{A} : V \to \mathbb{R}^m$ is a linear operator and $\mathcal{K} \subset V$ is a given proper cone, a whole collection of different problems can be simultaneously analyzed. Among the classes of problems that can be interpreted as particular cases of the general conic formulation we have *linear programs, second-order cone programs* (SOCP), and SDP, when we take the cone $\mathcal{K}$ to be the nonnegative

orthant $\mathbb{R}^n_+$, the second order cone in n variables, or the PSD cone $\mathbb{S}^n_+$, respectively. We have then the following natural inclusion relationship among the different optimization classes.

$$\text{LP} \subseteq \text{SOCP} \subseteq \text{SDP}.$$

For numerical efficiency reasons, we want to formulate and solve a given problem in as simplified a setting as possible.

Second-order cone programs partially enjoy the expressive modeling power that non-polyhedral cones such as the PSD cone have, but at the same time share with LP the scalability properties necessary for solving large-scale instances (of the order of tens of thousands of variables), that are currently out of reach for SDP. More details on the theory and applications of second-order cone programming can be found in [39,34,2].

We are confident that the application areas where SOCP techniques are used will increase dramatically in the future, either as the natural modeling framework, or as a computational efficient device to approximate an underlying more complicated set, such as the SDP cone.

The case of optimization problems on the cone of sum of squares (SOS) polynomials, to be discussed in Section 7, has attracted lately the attention of several researchers. From a strictly mathematical viewpoint it is equivalent to SDP; SDP corresponds to non-negativity conditions on quadratic forms, which coincide with the SOS forms, and conversely, we can reformulate SOS programs as SDPs. Even though the possibilities of numerically exploiting the structure of SOS programs is not yet fully understood, there are at least a couple of reasons why it is convenient to think of them as a distinguished category of problems; the richness of the underlying algebraic structure, as well their convenient modeling flexibility.

Yet another interesting class, that only in the last few years has begun to be explored, is that of optimization problems involving *hyperbolic polynomials* and their associated cones [24,5], which are of great theoretical interest. Research in this area is still at its early stage, with some open fundamental mathematical questions, and practical applications still under investigation.

3. Certificates and Duality

3.1. Duality for Optimization Problems

If we have an optimization in the form (2), there is a useful and important duality theory. Define the

Lagrangian function $L : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ by

$$L(x, \lambda) = f_0(x) + \sum_{i=1}^{m} \lambda_i f_i(x)$$
$$\text{for all } x \in \mathbb{R}^n, \ \lambda \in \mathbb{R}^m,$$

and define the set D by

$$D = \left\{ \lambda \in \mathbb{R}^m \,|\, \lambda \geq 0, \ \inf_{x \in \mathbb{R}^n} L(x, \lambda) \text{ is finite} \right\}.$$

Then define the *Lagrange dual function* $g : D \to \mathbb{R}$ by

$$g(\lambda) = \inf_{x \in \mathbb{R}^n} L(x, \lambda) \quad \text{for all } \lambda \in D,$$

and the *dual problem* as

$$\begin{aligned} \text{maximize} \quad & g(\lambda) \\ \text{subject to} \quad & \lambda \in D. \end{aligned} \tag{5}$$

A point $\lambda \in D$ is called *dual feasible*. The importance of the dual problem is that, given any dual feasible λ, we have a lower bound on the set of achievable values of the primal problem. That is, if x is primal feasible, and λ is dual feasible, we have

$$g(\lambda) \leq f_0(x).$$

This is known as *weak duality*, and follows because, for any $\lambda \in D$ and $x \in P$, we have

$$\begin{aligned} g(\lambda) &\leq L(x, \lambda) \\ &= f_0(x) + \sum_{i=1}^{m} \lambda_i f_i(x) \,. \\ &\leq f_0(x) \end{aligned}$$

There are also important cases under which the optimal value of the dual problem (5) is equal to the optimal value of the primal problem (2), called *strong duality*. This holds if the functions $f_0, \ldots, f_m$ are convex, and further standard conditions, called *constraint qualifications* hold. An example of a typical constraint qualification is that if there exists $x \in \mathbb{R}^n$ which is strictly feasible, i.e., $f_i(x) < 0$ for all $i = 1, \ldots, m$, then strong duality holds. Note that we assume that the functions f_i are defined on all $\mathbb{R}^n$.

3.2. Duality for Feasibility Problems

For feasibility problems, the Lagrange dual function is

$$g(\lambda) = \inf_{x \in \mathbb{R}^n} \sum_{i=1}^{m} \lambda_i f_i(x) \quad \text{for all } \lambda \in \mathbb{R}^m,$$

and the function g is assigned the value $-\infty$ if the minimization above is unbounded below. We define the set

$$D = \left\{\lambda \in \mathbb{R}^m \mid \lambda \geq 0, \; g(\lambda) > 0\right\}. \tag{6}$$

This gives the *theorem of alternatives*, which states that at most one of P and D is feasible. To show this, suppose P is nonempty, and $x \in P$. Then we have, for all $\lambda \in \mathbb{R}^m$,

$$g(\lambda) \leq \sum_{i=1}^m \lambda_i f_i(x) \leq 0,$$

and so D is empty. The primal and dual feasibility problems defined by P and D are called *weak alternatives*. Note that it may happen that both P and D are infeasible. Again, if the functions $f_1, \ldots, f_m$ are convex, then under suitable constraint qualifications the sets P and D are *strong alternatives*; exactly one of P or D is nonempty.

3.2.1. Certificates of infeasibility

The theorem of alternatives above, and its generalizations discussed below, are of great importance. For many practical problems, one would like to develop an algorithm to find a feasible point x. Once such an x is found, verifying that it satisfies the desired conditions is straightforward, since one may evaluate the functions $f_1, \ldots, f_m$ at x and check that they are all negative. However, if a feasible point x cannot be found, one would like to demonstrate that *no solution exists*, and that the problem specification must be altered. The dual feasibility problem gives precisely this; if instead one can find a dual feasible point, then this provides a *certificate* that no solution to the primal system of inequalities exists. We emphasize that duality is a property by which infeasibility (or suboptimality) can be *concisely certified*.

3.3. Duality in Semidefinite Programming

While we will not delve into the details here, it is possible to extend the duality theory described in the previous subsection to the case of generalized inequalities. Since SDP naturally falls within this framework, it is not surprising then that SDPs, being convex optimization problems, enjoy a rich duality structure.

The *standard dual* associated with the primal program (3) is also an SDP, given by

$$\begin{aligned} \text{maximize} \quad & b^{\mathrm{T}} y \\ \text{subject to} \quad & \sum_{i=1}^m y_i A_i \preceq C, \end{aligned} \tag{7}$$

where the dual variable $y \in \mathbb{R}^m$. As above we have weak duality; any feasible solution of the dual provides a lower bound on the achievable values of the primal; conversely, feasible primal solutions give upper bounds on dual solutions. For SDPs, this may be shown directly, since for any primal feasible X and dual feasible y we have that the difference between the two objective functions satisfies

$$\begin{aligned} \text{trace}(CX) - b^{\mathrm{T}} y &= \text{trace}(CX) - \sum_{i=1}^m y_i \, \text{trace}(A_i X) \\ &= \text{trace}\left(\left(C - \sum_{i=1}^m y_i A_i\right) X\right) \\ &\geq 0, \end{aligned}$$

with the last inequality being true because of self-duality of the PSD cone. The following result expresses a strong duality property possessed by SDPs.

Theorem 1. Consider the primal-dual SDP pairs (3) and (7). If either feasible set has has a nonempty interior, then for every $\epsilon > 0$, there exist feasible X, y such that

$$\text{trace}(C^{\mathrm{T}} X) - b^{\mathrm{T}} y < \epsilon.$$

Furthermore, if both feasible sets have nonempty interiors, then the optimal solutions are achieved by some $X_\star$, $y_\star$.

4. Semidefinite Programming and Control

There has been extensive application of SDP techniques for analysis and synthesis for control systems, and the field is still experiencing a continuous growth. Significant changes have resulted from this, one of the most important being a radical change in what it means to have *solved* a control problem. Historically, control problems have been considered solved when they have been reduced to methods computable with the technology of the time; at various points in history, this has meant a reduction to equations soluble via graphical plots (such as root-locus), and a reduction to linear algebraic systems, such as Riccati equations. Nowadays it is almost universally accepted that the reduction of a question to an easily tractable convex optimization problem (in particular, an SDP) constitutes a solution.

It is also widely accepted that matrix inequalities and duality are inextricably linked with the study of linear control systems; this is witnessed by the fundamental role played by the Lyapunov and Riccati inequalities [4,63]. Very significant work by many

researchers has continuously expanded these boundaries; robustness analysis techniques based on quadratic stability, multipliers, structured singular values, or integral quadratic constraints (IQCs), as well as the corresponding synthesis methodologies, have dealt with uncertainties and nonlinearities, and can all be cast under the convenient computational machinery of matrix inequalities. Given our purposes, we will not go here into a detailed survey of the many available results, but instead refer the reader to the extensive available literature, a good start being the general reference works mentioned in the introduction. Also, as a representative example that illustrates several of the now standard techniques available, we present next a simple case of multiobjective design.

Consider the linear system

$$\dot{x}(t) = Ax(t) + B_z z(t) + B_u u(t),$$
$$w(t) = Cx(t).$$

We are interested in designing a state-feedback controller $u = Kx$ that minimizes the $\mathcal{H}_2$ norm of the transfer function T_{zw} from the input z to the output w. Additionally, we would like the poles of the closed-loop system to be located in a specific region of the complex plane, described by the feasible set of the linear matrix inequality

$$\mathcal{P} = \{Z \in \mathbb{C} \mid P + Rz + R^{\mathrm{T}}\bar{z} \succ 0\},$$
$$P \in \mathbb{S}^k, \quad R \in \mathbb{R}^{k \times k},$$

where the inequality should be interpreted as constraining the Hermitian matrix on the left-hand side to be positive semidefinite. It is well known [10,15,53] that the two desired design specifications $\|T_{zw}\|_2 \le \gamma$ and $\sigma(A + B_u K) \subset \mathcal{P}$ can be expressed as the SDP conditions

$$\mathrm{trace}(Z) \le \gamma^2,$$

$$\begin{bmatrix} A_{\mathrm{cl}}X_1 + X_1 A_{\mathrm{cl}}^{\mathrm{T}} & X_1 C^{\mathrm{T}} \\ C X_1 & -I \end{bmatrix} \preceq 0,$$

$$\begin{bmatrix} X_1 & B_z \\ B_z^{\mathrm{T}} & Z \end{bmatrix} \succeq 0,$$

and

$$P \otimes X_2 + R \otimes A_{\mathrm{cl}}X_2 + R^{\mathrm{T}} \otimes X_2 A_{\mathrm{cl}}^{\mathrm{T}} \succ 0, \quad X_2 \succ 0,$$

respectively, where $A_{\mathrm{cl}} = A + B_u K$ is the closed-loop state matrix, and $\otimes$ is the Kronecker product.

The conditions here are not yet affine in the decision variables, since there are bilinear terms such as KX_2. To convert this into a tractable synthesis problem, a (possibly conservative) approach is to assume $X_1 = X_2 = X$, and then introduce the variable $Y = KX$, so

$A_{\mathrm{cl}}X_1 = A_{\mathrm{cl}}X_2 = AX + B_u Y$. The system to be solved reduces now to

minimize $\mathrm{trace}(Z)$

subject to $0 \succeq \begin{bmatrix} (AX + B_u Y) + (AX + B_u Y)^{\mathrm{T}} & XC^{\mathrm{T}} \\ CX & -I \end{bmatrix}$,

$$0 \preceq \begin{bmatrix} X & B_z \\ B_z^{\mathrm{T}} & Z \end{bmatrix},$$
$$0 \prec P \otimes X + R \otimes (AX + B_u Y)$$
$$\qquad + R^{\mathrm{T}} \otimes (AX + B_u Y)^{\mathrm{T}},$$
$$0 \prec X,$$

which is a *bona-fide* SDP. After solving it numerically, the controller can be recovered from the solution using $K = YX^{-1}$.

As a simple illustration, assume the numerical data

$$A = \begin{bmatrix} 2 & 1 \\ -1 & 4 \end{bmatrix},$$
$$B_u = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad B_z = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 1 \end{bmatrix},$$

and the region in the complex plane defined by $|z + 2| < 1$, i.e.,

$$P = \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix}, \quad R = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}.$$

Ignoring the constraint on the closed-loop pole locations, the optimal controller achieves a value of $(\mathrm{trace}\, Z)^{1/2} = \|T_{zw}\|_2$ arbitrarily close to $\sqrt{10} \approx 3.16$, but the gains K grow unboundedly, as does the magnitude of one of the poles. Imposing now the left-over constraint, the optimal solution achieves a value of $(\mathrm{trace}\, Z)^{1/2} \approx 21.26$. This is only an upper bound on the $\mathcal{H}_2$ norm; the actual achieved value is $\|T_{zw}\|_2 = 7.57$, with the final controller being $K = [-9.92, 35]$, and closed-loop poles at $-1.96 \pm 0.66j$.

Several improved modifications of these methods are also available, including better ways of handling the requirement on a common Lyapunov function, and extensions to multiobjective output feedback; once again, we refer the reader to the cited bibliography for the details.

5. Combinatorial Optimization

Many well-known problems in linear control theory have the property that they can be *exactly* reformulated in terms of SDPs, although in some cases the required transformations can be far from obvious. When attempting to extend convexity-based methods to larger classes of problems, a complete

characterization of the solution set via a simple SDP may no longer be possible. The difficulties arise in two different (though related) fronts. On the one hand, it may be computationally difficult to find feasible points, that show that the problem indeed has a solution; on the other hand, there may not exist a *concise* way of demonstrating infeasibility. There is a clear interdependence of these two issues, since if the first question were always easy, then the second would follow, but they are not equivalent. For the case of feasibility problems defined by convex functions, the first property is satisfied since we can efficiently search for solutions. As we have seen, convex duality gives us a useful way of certifying the nonexistence of solutions. After introducing the formal machinery of a computational model and polynomial-time algorithms to make their meaning more precise, these questions play an essential role in some of the central issues in computational complexity; whether $P = NP$ and $NP = co\text{-}NP$, respectively [20].

Note also that convexity properties by themselves do not automatically imply polynomial-time solvability. The specific *representation* of the feasible set plays a crucial role here, and the existing results require the availability of either subgradients [55], or a self-concordant barrier function [39]. There are many examples of optimization problems over (possibly implicitly given) convex sets, where all these operations, or even checking membership, are computationally hard.

Computational hardness arises in many cases as a direct consequence of the combinatorial nature of the problem. When this happens, a class of approximate methods, usually called *convex relaxations*, are typically used to either bound the achievable optimal values, or obtain reasonably good candidate solutions. While relaxations based on linear programming have been and still are extensively used in many application areas edited such as integer programming, in the last few years a vast collection of new and powerful SDP-based relaxations have attracted the attention of many researchers.

An important and well-known example where these SDP relaxations are extensively used is boolean programming. Consider the NP-hard problem of minimizing a quadratic form over the discrete set given by the vertices of a hypercube, i.e.,

$$\begin{aligned} \text{minimize} \quad & x^T Q x \\ \text{subject to} \quad & x_i \in \{+1, -1\} \quad \text{for all } i = 1, \ldots, n, \end{aligned}$$
(8)

where $Q \in \mathbb{S}^n$. Notice that the feasible set is discrete and has 2^n points, so convexity-based techniques

would seem not to be applicable. However, a lower bound on the optimal solution can be found by solving the primal-dual pair of SDPs given by

$$\begin{aligned} \text{minimize} \quad & \text{trace}(QX) \\ \text{subject to} \quad & X \succeq 0, \\ & X_{ii} = 1 \quad \text{for all } i = 1, \ldots, n, \end{aligned}$$
(9)

and its dual

$$\begin{aligned} \text{maximize} \quad & \text{trace}(\Lambda) \\ \text{subject to} \quad & \Lambda \text{ diagonal}, \\ & Q - \Lambda \succeq 0. \end{aligned}$$
(10)

This pair of SDPs have several possible interpretations, extensively discussed in the literature. First notice, that, on the primal side, letting $X = xx^T$, we can equivalently rewrite the original objective function as $x^T Q x = \text{trace}(Q x x^T) = \text{trace } Q X$. The feasible set is given now by $X \succeq 0$, $X_{ii} = 1$, $\text{rank}(X) = 1$. Dropping the rank constraint, we obtain directly the SDP (9).

Another interpretation is given from the dual SDP; an argument based on Lagrangian duality applied to the constraints $x_i^2 - 1 = 0$ is enough. We have

$$\begin{aligned} L(x, \lambda) &= x^T Q x - \sum_i \lambda_i (x_i^2 - 1) \\ &= x^T (Q - \Lambda) x + \text{trace } \Lambda, \end{aligned}$$

where: $\Lambda = \text{diag } \lambda_i$.

For L to be bounded below, the condition $Q - \Lambda \succeq 0$ is required, and therefore the SDP (10) directly follows.

This relaxation can also be interpreted as a specific case of the so-called S-procedure; see [10] and the references therein. Several important results are available regarding the performance of these relaxations; the celebrated Goemans–Williamson approximation scheme for MAXCUT [23] relies exactly on this relaxation, followed by a randomized rounding step. For many other related problems, it has been possible to prove *a priori* approximation guarantees.

5.1. Combinatorial Optimization and Control

An example of combinatorial optimization techniques applied to control is as follows. Consider the problem of designing an optimal open-loop input for the single-input discrete-time system

$$\begin{aligned} x(t + 1) &= A x(t) + B u(t), \\ y(t) &= C x(t) \end{aligned}$$

for $t = 0, \ldots, N$, where one would like to minimize the ℓ_2 criterion $q(u) = \|y(t) - y_r(t)\|^2$, with y_r being the

desired reference output trajectory, and the input $u(t)$ constrained by $|u(t)| = 1$ for all $t = 0, \ldots, N$. In other words, we have an open-loop LQR-type problem, but where the input is *bang-bang*, taking only two possible values.

It is clear that a simple lower bound on the optimal value can be obtained from the solution of an associated unconstrained LQR problem. We are interested in computing better, less conservative lower bounds on the achievable value of $q(u)$. We can formulate the search for such a lower bound as an SDP. For notational simplicity, let $q(u) = u^T Q u + 2r^T u + s$, where the expressions of Q, r, s can be easily obtained from those of A, B, C and y_r. We are thus interested in bounding the value of the following optimization problem.

$$\text{minimize} \quad \begin{bmatrix} u \\ 1 \end{bmatrix}^T \begin{bmatrix} Q & r \\ r^T & s \end{bmatrix} \begin{bmatrix} u \\ 1 \end{bmatrix}$$
$$\text{subject to} \quad u_i \in \{+1, -1\} \quad \text{for all } i.$$

Let q^* denote the optimal value of this problem. Similar to the one described earlier, a simple SDP bound on the optimal value can be obtained as follows. Let q_* be the optimal value of

$$\text{minimize} \quad \text{trace}(D)$$
$$\text{subject to} \quad \begin{bmatrix} Q - D & r \\ r^T & s \end{bmatrix} \succeq 0,$$

where D is a diagonal matrix. Clearly $q_* \leq q^*$, as can easily be seen by multiplying the matrix inequality left and right by $[u^T \ 1]$ and its transpose. Notice also that $q_* \geq 0$, since $D = 0$ is always a feasible point.

6. Robust Optimization

An important development in recent years in the field of robust optimization has been *robust semidefinite programming*. Here one has a parametrized family of linear matrix inequalities, and one would like to find a point which is simultaneously feasible for the whole family.

We give here an example from finite-horizon control, which may be used, for example, in receding horizon control. We have a linear dynamical system

$$x(t + 1) = Ax(t) + B_w w(t) + B_u u(t),$$
$$y(t) = Cx(t),$$
$$(11)$$

where $x(t) \in \mathbb{R}^n$, $u(t) \in \mathbb{R}^{n_u}$ and $w(t) \in \mathbb{R}^{n_w}$ for all $t = 0, \ldots, N$, and the system evolves on the discrete-time interval $[0, N]$. We would like to design an input signal $u(0), \ldots, u(N-1)$ so that the system output y tracks

a desired signal y_{des}. Define

$$u = \begin{bmatrix} u(0) \\ \vdots \\ u(N-1) \end{bmatrix}, \quad w = \begin{bmatrix} w(0) \\ \vdots \\ w(N-1) \end{bmatrix}$$

$$y = \begin{bmatrix} y(0) \\ \vdots \\ y(N) \end{bmatrix}, \quad y_{\text{des}} = \begin{bmatrix} y_{\text{des}}(0) \\ \vdots \\ y_{\text{des}}(N) \end{bmatrix}.$$

Suppose the set of disturbances w is given by

$$W = \Big\{ w \in \mathbb{R}^{Nn_w} \mid \|w(t)\|_2 \leq w_{\max}$$
$$\text{for all } t = 0, \ldots, N - 1 \Big\}.$$

In this way the disturbances are specified so that at each time t the vector $w(t)$ lies within a Euclidean ball. Such sets, and generalizations thereof, are called *ellipsoidal uncertainty* in [6].

We would like to find an input sequence u which solves

$$\min_u \max_{w \in W} \|y - y_{\text{des}}\|_2.$$

As standard, one may construct block-Toeplitz matrices T and S such that $y = Tu + Sw$ for all u, y satisfying the dynamics (11), and so we can write this problem as

$$\min_u \max_{w \in W} \| Tu + Sw - y_{\text{des}} \|_2 . \quad (12)$$

This problem is therefore a *robust least-squares* problem; recent work [17] in this area has produced solutions using semidefinite programming for this and more general versions, and certain classes of this problem may be solved using second-order cone programming [34].

Let $s_1, \ldots, s_{Nn_w}$ be the columns of S, and define

$$F_0(u) = \begin{bmatrix} t & (Tu - y_{\text{des}})^T \\ Tu - y_{\text{des}} & I \end{bmatrix},$$

$$F_i(u) = \begin{bmatrix} 0 & s_i^T \\ s_i & 0 \end{bmatrix} \quad \text{for all } i = 1, \ldots, Nn_w.$$

Then problem (12) is equivalent to the robust SDP

$$\text{minimize} \quad t$$
$$\text{subject to} \quad F_0(u) + \sum_{i=1}^{Nn_w} w_i F_i(u) \succeq 0 \quad \text{for all } w \in W,$$
$$(13)$$

where w_i is the ith (scalar) element of w. An upper bound on the solution of this robust SDP, along with an input u that achieves it, may be found by solving the following SDP in the variables $t, u, P_1, \ldots, P_N$, $Q_1, \ldots, Q_N$.

minimize $\quad t$

subject to
$$\begin{bmatrix} P_i & F_{(i-1)n_w+1} & \cdots & F_{in_w-1} \\ F_{(i-1)n_w+1} & Q_i & & \\ \vdots & & \ddots & \\ F_{in_w-1} & & & Q_i \end{bmatrix} \succeq 0$$
for all $i = 1, \ldots, N$, $\qquad(14)$
$$\begin{bmatrix} t & (Tu - y_{\text{des}})^{\text{T}} \\ Tu - y_{\text{des}} & I \end{bmatrix}$$
$$- \frac{1}{2} \sum_{i=1}^{N} (P_i + Q_i) \succeq 0.$$

Given a feasible solution, we have

$$\max_{w \in W} \| Tu + Sw - y_{\text{des}} \|_2 \le t^{1/2}.$$

The particular problem shown above is of course just an example of the application of the robust SDP approach to control. Many other problems for which one has synthesis conditions in terms of linear matrix inequalities may be similarly generalized to construct robust synthesis conditions.

7. Sum of Squares Methods

Many optimization problems have feasible sets specified by polynomial equations and inequalities. This includes the feasible set of a SDP, as well as many more general classes. Finite integer constraints can be easily expressed as polynomial equations; e.g., $x \in \{0, 1\}$ is equivalent to the constraint equation $x(x - 1) = 0$.

Specifically, a (basic) *semialgebraic set* is a subset of $\mathbb{R}^n$ defined by a finite number of polynomial equations and inequalities. For example, the subset of $\mathbb{R}^2$ defined by

$$\{(x_1, x_2) \in \mathbb{R}^2 \mid x_1^2 + x_2^2 \le 1, x_1^3 - x_2 \le 0\}$$

is a semialgebraic set. Given a set of polynomials specifying such a set, one would like to find either a feasible point or a certificate of infeasibility. Clearly, a semialgebraic set need not be convex, and in general, testing feasibility of such sets is intractable.

Recent developments have led to an approach using semidefinite programming to test feasibility of semialgebraic sets [43–45] (see also [32] for a dual

approach, and [14]). By solving a semidefinite program, one may obtain a *certificate of infeasibility* for an infeasible semialgebraic set. The size of this certificate (i.e., the size of the SDP to be solved) may depend on the specific set. One of the appealing features of this approach is that there is a hierarchy of such certificates, allowing one to solve a sequence of SDPs in an attempt to prove infeasibility of a semialgebraic set.

The simplest case is when we are concerned with a set defined by a single polynomial inequality, as follows. Given $f \in \mathbb{R}[x_1, \ldots, x_n]$, we would like to know if there exists $x \in \mathbb{R}^n$ such that $f(x) < 0$. If there does not exist such an x, we have $f(x) \ge 0$ for all x and f is called positive semidefinite, or PSD.

An example is as follows. Suppose f is given by

$$f(x) = 4x_1^4 + 4x_1^3 x_2 - 7x_1^2 x_2^2 - 2x_1 x_2^3 + 10x_2^4,$$

and we would like to determine the feasibility of the set

$$P = \{x \in \mathbb{R}^2 \mid f(x) < 0\}.$$

Assume now that there exists a matrix $Q \in \mathbb{S}^3$ such that

$$f(x) = \begin{bmatrix} x_1^2 \\ x_2^2 \\ x_1 x_2 \end{bmatrix}^{\text{T}} \begin{bmatrix} q_{11} & q_{12} & q_{13} \\ q_{12} & q_{22} & q_{23} \\ q_{13} & q_{23} & q_{33} \end{bmatrix} \begin{bmatrix} x_1^2 \\ x_2^2 \\ x_1 x_2 \end{bmatrix}$$
$$= z^{\text{T}}(x) Q z(x). \qquad(15)$$

In general, there are many such Q, since this is just a system of linear equations in the entries of Q, as can be seen by matching coefficients on the right- and the left-hand sides. However, an important consequence is the following; *if we can find a PSD Q that satisfies Eq. (15), then $f(x)$ would have to take only non-negative values*, so one immediately has a proof that there cannot exist a real $x \in P$ and the matrix Q is then a certificate of infeasibility of the semialgebraic set P.

In the example above, we can parametrize the set of possible Q by

$$f(x) = \begin{bmatrix} x_1^2 \\ x_2^2 \\ x_1 x_2 \end{bmatrix}^{\text{T}} \begin{bmatrix} 4 & -\lambda & 2 \\ -\lambda & 10 & -1 \\ 2 & -1 & -7 + 2\lambda \end{bmatrix} \begin{bmatrix} x_1^2 \\ x_2^2 \\ x_1 x_2 \end{bmatrix}$$

for all $\lambda \in \mathbb{R}$, and therefore if there exists $\lambda \in \mathbb{R}$ such that

$$\begin{bmatrix} 4 & -\lambda & 2 \\ -\lambda & 10 & -1 \\ 2 & -1 & -7 + 2\lambda \end{bmatrix} \succeq 0,$$

then P is infeasible. This is an SDP. In this case, picking $\lambda = 6$ gives a positive semidefinite matrix, and therefore a valid proof of the infeasibility of the set P.

This approach is quite general. For every polynomial of degree d in n variables $f \in \mathbb{R}[x_1, \ldots, x_n]$, we form the vector of monomials $z(x)$ which has as elements all monomials of degree $d/2$ or less. In fact, fewer monomials may be included, which gives a reduced-size SDP, depending on the particular polynomial f. The set of matrices Q for which $f(x) = z(x)^{\mathrm{T}} Qz(x)$ for all x is an affine set, and is simple to construct.

The above condition can be interpreted as follows. If there exists a positive semidefinite matrix Q such that $f(x) = z(x)^{\mathrm{T}} Qz(x)$, then we may factorize Q. Every PSD matrix Q may be factorized as $Q = VV^{\mathrm{T}}$, where $V \in \mathbb{R}^{n \times r}$ is a real matrix, and r is the rank of Q. With $v_1, \ldots, v_r$ the columns of V, we have

$$f(x) = z(x)^{\mathrm{T}} Qz(x)$$
$$= \sum_{i=1}^{r} (v_i^{\mathrm{T}} z(x))^2$$

and in this case we may choose a factorization with

$$V = \begin{bmatrix} 0 & 2 \\ 1 & -3 \\ 2 & 1 \end{bmatrix},$$

giving

$$f(x) = (x_2^2 + 2x_1 x_2)^2 + (2x_1^2 - 3x_2^2 + x_1 x_2)^2,$$

which is a sum-of squares (SOS) decomposition of f [46,51]. From this decomposition, it is immediately clear that there cannot exist a feasible x, and the decomposition (characterized by Q) is clearly a certificate of infeasibility. Hence we may decide whether a polynomial is SOS by solving an SDP.

The natural question to ask, when P infeasible, is whether there always exists such a certificate of infeasibility. Expressed another way, is every PSD polynomial a sum of squares? It was shown by Hilbert in 1888 that in general this is not the case; there exist polynomials which are PSD but not SOS. This is to be expected, since it is known that testing feasibility of a semialgebraic set (even if it is defined by just one polynomial) is intractable; but if every PSD polynomial was SOS then feasibility would be testable using a simple SDP. An example of a polynomial which is PSD, but not SOS, is the well-known Motzkin form

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

Besides the more general methods in Section 7.1, there is a further result that gives a method for bridging the gap between SOS and PSD polynomials. It was shown by Reznick [50] that if $f(x) > 0$ for all $x \in \mathbb{R}^n$, then there exists some r such that

$$f(x) \left(\sum_{i=1}^{n} x_i^2 \right)^r$$

is SOS. The coefficients of this polynomial are affine functions of the coefficients of f, and so for each r one may test whether this product is SOS by solving an SDP. This gives a *sequence* of SDPs, of growing dimension, which may be used to test whether a given polynomial is positive definite.

An important remaining question is how to find feasible points x when P is feasible. This is a distinct question from finding a certificate of infeasibility; however the dual problem to the above SDP can be used in certain cases to do this.

7.1. Generalized Duality and the Positivstellensatz

We would now like to consider how to certify infeasibility for semialgebraic sets defined by multiple polynomial inequalities, that is sets P of the form

$$P = \left\{ x \in \mathbb{R}^n \mid f_i(x) \geq 0 \text{ for all } i = 1, \ldots, m \right\},$$

(16)

where the functions f_i are polynomials on $\mathbb{R}^n$. For convex feasibility problems which satisfy a constraint qualification, the theorem of alternatives provides a necessary and sufficient condition for feasibility of the primal. However, for semialgebraic sets, the feasibility problems are defined by nonconvex functions $f_1, \ldots, f_m$. In this case, the following stronger result is very useful. Define $\mathcal{P}_{\mathrm{sos}}^n$ the set of SOS polynomials, as

$$\mathcal{P}_{\mathrm{sos}}^n = \left\{ s \in \mathbb{R}[x_1, \ldots, x_n] \mid s \text{ is an SOS} \right\}.$$

Define the functional $G : \mathcal{P}_{\mathrm{sos}}^n \to \mathbb{R}$ by

$$G(s_1, \ldots, s_m) = \sup_{x \in \mathbb{R}^n} \sum_{i=1}^{n} s_i(x) f_i(x)$$
$$\text{for all } s_1, \ldots, s_m \in \mathcal{P}_{\mathrm{sos}}^n.$$

Then define the dual set of functions

$$D_1 = \left\{ (s_1, \ldots, s_m) \mid s_i \in \mathcal{P}_{\mathrm{sos}}^n, G(s_1, \ldots, s_m) < 0 \right\}.$$

Then P and D_1 are weak alternatives; at most one of P and D_1 is nonempty. To see this, suppose that P is

feasible, and $x \in P$. Then, for all $s_1, \ldots, s_m \in \mathcal{P}_{\mathrm{sos}}^n$, we have

$$G(s_1, \ldots, s_m) \geq \sum_{i=1}^{m} s_i(x) f_i(x) \geq 0$$

and hence D_1 is empty. This refutation is a simple form of weak duality.

7.1.1. The Positivstellensatz

A slight modification of this result is as follows. If there exist SOS polynomials $s_0, s_1, \ldots, s_m$ such that

$$\sum_{i=1}^{n} s_i(x) f_i(x) + s_0(x) + 1 = 0, \qquad (17)$$

then the set P is infeasible. This sufficient condition for infeasibility of P may be tested using semidefinite programming. One picks a fixed degree d over which to search for SOS polynomials $s_1, \ldots, s_m$, satisfying this condition. The decision variables are the coefficients of the polynomials and the constraints that the s_i be SOS are imposed as positive semidefiniteness constraints. Clearly, if we can find a set of functions $s_1, \ldots, s_m$ satisfying Eq. (17) then the set P is infeasible. A similar, but stronger condition is

$$\sum_{i=1}^{n} s_i(x) f_i(x) + \sum_{i=1}^{n} \sum_{j=1}^{n} t_{ij}(x) f_i(x) f_j(x) + s_0(x) + 1 = 0,$$
$$(18)$$

which again may be tested via SDP. Refutations of this form have very strong duality properties. It can be shown that, by allowing for unrestricted degrees of the polynomials s_i and t_{ij}, and arbitrary squarefree products of f_i, one may always construct a refutation for any given infeasible semialgebraic set P using SOS polynomials [8]. No assumptions whatsoever on the polynomials f_i are required. This result is called the *Positivstellensatz*. Software for testing feasibility of semialgebraic sets using the above methods is available in the form of a Matlab toolbox, called SOS-TOOLS [47].

7.1.2. Applications

Testing feasibility of semialgebraic sets has important applications in control and combinatorial optimization. For example, the integer program of Eq. (8) can be formulated as

$$\begin{aligned}
&\text{minimize} && t \\
&\text{subject to} && x^{\mathrm{T}} Q x \leq t, \qquad (19) \\
&&& x_i^2 - 1 = 0 \quad \text{for all } i = 1, \ldots, n,
\end{aligned}$$

hence we check feasibility for a fixed t using the above Positivstellensatz approach. In fact, the relaxation discussed in that section exactly coincides with that obtained using the first level of the general approach just presented.

7.1.3. SOS and Lyapunov functions

The possibility of reformulating conditions for a polynomial to be a sum of squares as an SDP is very useful, since we can use the SOS property in a control context as a convenient sufficient condition for polynomial nonnegativity. Recent work has applied the SOS approach to the problem of finding a Lyapunov function for nonlinear systems [43,41,57]. This approach allows one to search over affinely parametrized polynomial or rational Lyapunov functions for systems with dynamics of the form

$$\dot{x}_i(t) = f_i(x(t)) \quad \text{for all } i = 1, \ldots, n,$$

where the functions f_i are polynomials or rational functions. Then the condition that the Lyapunov function be positive, and that its Lie derivative be negative, are both directly imposed as SOS constraints in terms of the coefficients of the Lyapunov function.

As an example, we consider the following system, suggested by Krstić.

$$\begin{aligned}
\dot{x} &= -x + (1 + x)y, \\
\dot{y} &= -(1 + x)x.
\end{aligned}$$

Using SOSTOOLS [47] we easily find a quartic polynomial Lyapunov function, which after rounding (for purely cosmetic reasons) is given by

$$\begin{aligned}
V(x, y) = {} & 6x^2 - 2xy + 8y^2 - 2y^3 + 3x^4 \\
& + 6x^2 y^2 + 3y^4.
\end{aligned}$$

It can be readily verified that both $V(x,y)$ and $(-\dot{V}(x,y))$ are SOS, since

$$V = \begin{bmatrix} x \\ y \\ x^2 \\ xy \\ y^2 \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} 6 & -1 & 0 & 0 & 0 \\ -1 & 8 & 0 & 0 & -1 \\ 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 6 & 0 \\ 0 & -1 & 0 & 0 & 3 \end{bmatrix} \begin{bmatrix} x \\ y \\ x^2 \\ xy \\ y^2 \end{bmatrix},$$

$$-\dot{V} = \begin{bmatrix} x \\ y \\ x^2 \\ xy \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} 10 & 1 & -1 & 1 \\ 1 & 2 & 1 & -2 \\ -1 & 1 & 12 & 0 \\ 1 & -2 & 0 & 6 \end{bmatrix} \begin{bmatrix} x \\ y \\ x^2 \\ xy \end{bmatrix},$$

and the matrices in the expression above are positive definite. Similar approaches may also be used for finding Lyapunov functionals for certain classes of hybrid systems.

8. Exploiting Control-Specific Structure

The SDP problems that arise from the analysis and synthesis of linear control systems, usually have a very specific structure, that should be exploited to achieve the best possible computational efficiency. In particular, the ubiquity of Lyapunov-style terms of the form $A^{\mathrm{T}}P + PA$ (or $A^{\mathrm{T}}PA - P$, in the discrete case) suggest that generic implementations using the SDP standard forms described in Eqs (3) and (7) will not be optimal, unless the extra structure is somehow taken into account.

An important subclass for which several customized algorithms are already available is that of optimization problems with an structure induced by the Kalman–Yakubovich–Popov (KYP) lemma (see [63,49] and the references therein). This fundamental result establishes the equivalence between a frequency domain inequality and the feasibility of a particular SDP. It is an important generalization of classical linear control results, such as the bounded real and positive real lemmas, and the cornerstone of several analysis and synthesis results.

The harder direction of the KYP lemma, that the frequency domain inequality implies the existence of a storage function, can be interpreted as a *lossless* property of an associated SDP relaxation, much in the style of the ones presented in Section 5. Yet another related reformulation of the KYP lemma is given by the equivalence of the structured singular value μ and its upper bound in the case where there is a full block and one scalar block [40].

In the case of systems with large state dimension n, the KYP approach is not very efficient, since the matrix variable representing the storage function that appears in the LMI has $(n^2 + n)/2$ components, and therefore the computational requirements are quite large, even for medium sized problems.

Several different methods have been proposed recently for the efficient solution of this and related kind of problems. The approaches by Kao et al. [29,30], Parrilo [42] and Varga and Parrilo [62] all rely on outer approximation methods based on a discretization of the frequency axis, with the former using linear programming and the latter using SDP cuts. The scheme by Hansson and Vandenberghe [25] is based on the interior point machinery, but cleverly exploits the Lyapunov structure at each iteration at the linear algebra level. Rotea and D'Amato exploit several properties of the μ SDP upper bound to obtain significant speed-ups in the computation of the worst-case frequency response [52]. Additionally, methods for taking advantage of autocorrelation structure have also been developed by Alkire and Vandenberghe [3].

9. Available Implementations

Despite the impressive advances on the theoretical and modeling aspects of SDP, much of its impact in applications has undoubtedly been a direct consequence of the efforts of many researchers in producing and making available good quality software. In this section we give pointers to and discuss briefly some of the current computational tools for effectively formulating and solving SDPs.

9.1. Solvers

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

For applications of semidefinite programming to control, the pioneering codes were the MATLAB LMI toolbox [19] and SP [61]. Today there is a wide variety of excellent SDP solvers to choose from. For general-purpose small- and medium-scale problems, interior-point based solvers are probably the best choice, combining good performance and accuracy, primal and dual solutions, as well as reasonable speed-ups depending on the problem sparsity. We mention a few of the best-known ones: SeDuMi [56], SDPT3 [59], SDPA [18], CSDP [9], DSDP [7], among others. Other good pointers to the available SDP solvers are the SDP webpages of C. Helmberg and H. Wolkowicz.

Several approaches other than interior-point methods have also been investigated, a few of them being bundle methods [26,37], or nonlinear approaches based on special factorizations [13]. This research has increased steadily in the last few years. The codes based on these new developments are the only ones achieving satisfactory performance for some of the very large and structured problems arising from combinatorial optimization, with MAXCUT and the Lováasz theta function being prime examples.

A comprehensive benchmarking effort of the performance of several solvers in a representative collection of problems (including some arising from control) has been recently reported by Hans Mittelmann in [38]; an up-to-date version of the results is available at the web page http://plato.la.asu.edu/bench.html.

The consensus and experience among researchers, backed up by the hard data mentioned, seems to be indicate that there is no algorithm or code that uniformly outperforms all others. While most software packages will work satisfactorily on different problems, there may be (sometimes significant) differences in the computation times. It is therefore good advice to try different codes at the initial stages of solving a large-scale problem; not only do we benefit from the possible speed differences, but we can also validate the solutions against each other. The availability of SDP parsers such as the ones described in the following subsection can notably help during this migration process.

9.2. Parsers

The solvers described in the previous subsections usually take as inputs either text files containing a problem description, or directly the matrices (A_i, b, C) corresponding to the standard primal/dual formulation. Needless to say, this is often inconvenient at the initial modeling and solution stages. A natural solution is therefore to formulate the problem in a more neutral description, that can later be automatically translated to fit the requirements of each solver. For generic optimization problems, this has indeed been the approach of much of the operations research community, which has developed some well-known standard file formats, such as MPS, or optimization modeling languages like AMPL and GAMS. An important remark to keep in mind, much more critical in the SDP case than in LP, is the extent to which the problem structure can be signaled to the solver. There is a growing push within the optimization community towards the possibility of adding SDP-oriented extensions to the standard modeling languages mentioned.

In the meantime, however, enterprising researchers in control and related areas have written specific parsers that partially or fully automate the conversion tasks, when used within a problem-solving environment such as MATLAB. Among them we mention the early effort SDPSOL [12], as well as the more recent ones YALMIP [35], SeDuMi Interface [31], and LMILab translator [54], dealing with general SDPs, as well as the more

domain-specific IQCbeta [36], Gloptipoly [28], and SOSTOOLS [47].

Any of these parsers can make the task of posing and solving a specific problem a much simpler and enjoyable procedure than manual, error-prone methods. We strongly encourage the reader to take them for a test drive.

References

1. Alizadeh F. Interior point methods in semidefinite programming with applications to combinatorial optimization. SIAM J Optim 1995; 5(1): 13–51
2. Alizadeh F, Goldfarb D. Second order cone programming. Technical Report 51-2001, RUTCOR, Rutgers University, 2001
3. Alkire B, Vandenberghe L. Convex optimization problems involving finite autocorrelation sequences. Math Programming Series A 2002; 93(3): 331–359
4. Balakrishnan V, Vandenberghe L. Semidefinite programming duality and linear time-invariant systems. IEEE Trans Auto Control January 2003; AC-48(1): 30–41
5. Bauschke HH, Güler O, Lewis AS, Sendov HS. Hyperbolic polynomials and convex analysis. Canad J Math 2001; 53(3): 470–488
6. Ben-Tal A, El Ghaoui L, Nemirovskii A. Robustness. In: Wolkowitz H, Saigal R, Vandenberghe L. (eds). Handbook of semidefinite programming: theory, algorithms and applications, Kluwer, 2000, pp. 138–162
7. Benson S, Ye Y, Zhang X. DSDP, A dual scaling algorithm for positive semidefinite programming. Available from http://www-unix.mcs.anl.gov~benson/dsdp/
8. Bochnak J, Coste M, Roy M-F. Real algebraic geometry. Springer, Berlin 1998
9. Borchers B. CSDP, A C Library for semidefinite programming. Optim methods software 1999; 11(1): 613–623 Available from http://www.nmt.edu/~borchers/csdp.html
10. Boyd S, El Ghaoui L, Feron E, Balakrishnan V. Linear matrix inequalities in system and control theory, vol 15. Studies in applied mathematics. SIAM, Philadelphia, PA, 1994
11. Boyd S, Vandenberghe L. Convex optimization. 2002 Preprint
12. Boyd S, Wu S-P. SDPSOL: a parser/solver for semidefinite programs with matrix structure. User's guide, version beta. Stanford University, March 1996
13. Burer S, Monteiro RDC. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. Math Programming 2003; 95(2): 329–357
14. Chesi G, Garulli A, Tesi A, Vicino A. LMI-based techniques for solving quadratic distance problems. In: Proceedings of the 40th IEEE Conference on Decision and Control, December 2001, pp 3587–3592
15. Chilali M, Gahinet P. H_∞ design with pole placement constraints: an LMI approach. IEEE Trans Auto Control 1996; 41(3): 358–367

16. Doyle JC, Packard A, Zhou K. Review of LFTs, LMIs, and μ. In: Proceedings of the 30th IEEE Conference on Decision and Control, 1991

17. El Ghaoui L, Lebret H. Robust solutions to least-squares problems with uncertain data. SIAM J Matrix Analysis Appli 1997; 18(4): 1035–1064

18. Fujisawa K, Kojima M, Nakata K, Yamashita M. SDPA semidefinite programming algorithm – Version 6.00, 2002. Available from http://www.is.titech.ac.jp/~yamashi9/sdpa/

19. Gahinet P, Nemirovski A, Laub AJ, Chilali M. LMI control toolbox. The Math-Works Inc., Natick, Mass., May 1995

20. Garey MR, Johnson DS. Computers and intractability: a guide to the theory of NP-completeness. W. H. Freeman and Company, 1979

21. El Ghaoui L, Niculescu S-I. (eds). Advances on linear matrix inequality methods in control. SIAM, 1999

22. Goemans MX. Semidefinite programming in combinatorial optimization. Math Programming 1997; 79(1–3): 143–161

23. Goemans MX, Williamson DP. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. Journal of the ACM 1995; 42(6): 1115–1145

24. Güler O. Hyperbolic polynomials and interior point methods for convex programming. Math Oper Res 1997; 22(2): 350–377

25. Hansson A, Vandenberghe L. Efficient solution of linear matrix inequalities for integral quadratic constraints. In: Proceedings of the 39th IEEE Conference on Decision and Control, December 2000

26. Helmberg C. SBmethod – A C++ implementation of the spectral bundle method. Available from http://www-user.tu-chemnitz.de/~helmberg/SBmethod/

27. Helton JW, Vinnikov V. Linear matrix inequality representation of sets. Preprint, March 2002. Available from http://math.ucsd.edu/~helton/

28. Henrion D, Lasserre J-B. GloptiPoly – global optimization over polynomials with Matlab and SeDuMi. Available from http://www.laas.fr/~henrion/software/gloptipoly/

29. Kao CY, Jönsson U, Megretski A. A cutting plane algorithm for robustness analysis of periodically time-varying systems. In: Proceedings of the American Control Conference, 1999

30. Kao CY, Megretski A, Jönsson U. An algorithm for solving optimization problems involving special frequency dependent LMIs. In: Proceedings of the American Control Conference, 2000

31. Labit Y, Peaucelle D, Henrion D. SeDuMi Interface 1.02: a tool for solving LMI problems with SeDuMi. In: Proceedings of the CACSD Conference, Glasgow, September 2002

32. Lasserre JB. Global optimization with polynomials and the problem of moments. SIAM Journal of Optim 2001; 11(3): 796–817

33. Laurent M, Rendl F. Semidefinite programming and integer programming. Technical Report PNA-R0210, CWI, Amsterdam, April 2002

34. Lobo M, Vandenberghe L, Boyd S, Lebret H. Applications of second-order cone programming. Linear Algebra and its Applications November 1998; 284: 193–228

35. Löfberg J. YALMIP: a Matlab interface to SP, MAXDET and SOCP. Technical Report LiTH-ISY-R-2328, Department of Electrical Engineering, Linköping University, SE-581 83 Linkping, Sweden, Jan 2001. Available from http://www.control.isy.liu.se/~johanl/yalmip.html

36. Megretski A, Kao C-Y, Jönsson U, Rantzer A. A guide to IQCβ: software for robustness analysis, June 1998

37. Miller SA, Smith RS. Solving large structured semidefinite programs using an inexact spectral bundle method. In: Proceedings of the 39th IEEE Conference on Decision and Control, December 2000, pp 5027–5032

38. Mittelmann HD. An independent benchmarking of SDP and SOCP solvers. Math Progr 2003; 95: 407–430

39. Nesterov YE, Nemirovski A. Interior point polynomial methods in convex programming, vol 13 Studies in Applied Mathematics. SIAM, Philadelphia, PA 1994

40. Packard A, Doyle JC. The complex structured singular value. Automatica 1993; 29(1): 71–109

41. Papachristodoulou A, Prajna S. On the construction of Lyapunov functions using the sum of squares decomposition. In: Proceedings of the 41st IEEE Conference on Decision and Control, December 2002, pp 3482–3487

42. Parrilo PA. On the numerical solution of LMIs derived from the KYP lemma. In: Proceedings of the 38th IEEE Conference on Decision and Control, 1999

43. Parrilo PA. Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization. PhD thesis, California Institute of Technology May 2000. Available at http://www.cds.caltech.edu/~pablo/

44. Parrilo PA. Semidefinite programming relaxations for semialgebraic problems. Math Prog Series B. Preprint available at http://control.ee.ethz.ch/~parrilo/, 2003

45. Parrilo PA, Sturmfels B. Minimizing polynomial functions. In: Basu S, González-Vega L. (eds). Algorithmic and Quantitative Real Algebraic Geometry volume 60 DIMACS Series in Discrete Mathematics and Theoretical Computer Science. AMS, 2003. To appear, available from arXiv:math.OC/0103170

46. Powers V, Wörmann T. An algorithm for sums of squares of real polynomials. Journal of Pure Applied Algebra 1998; 127: 99–104

47. Prajna S, Papachristodoulou A, Parrilo PA. SOSTOOLS: Sum of squares optimization toolbox for MATLAB, 2002. Available from http://www.cds.caltech.edu/sostools and http://control.ee.ethz.ch/~parrilo/sostools

48. Ramana MV. Polyhedra, spectrahedra, semidefinite programming. In: Pardalos PM, Wolkowicz H. (eds). Topics in semidefinite and interior-point methods, Fields Institute Communications, Vol 18. American Mathematical Society, 1998, pp 27–38

49. Rantzer A. On the Kalman-Yakubovich-Popov lemma. Systems and Control Lett 1996; 28: 7–10

50. Reznick B. Uniform denominators in Hilbert 17th problem. Mathematische Zeitschrift 1995; 220(1): 75–97

51. Reznick B. Some concrete aspects of Hilbert's 17th problem. In: Real algebraic geometry and ordered structures, Contemporary Mathematics (American Mathematical Society), Vol 253, American Mathematical Society, 2000, pp 253

52. Rotea MA, D'Amato FJ. LFTB: an optimized algorithm to bound worst-case frequency response functions. In: Proceedings of the American Control Conference, June 2001, pp 3041–3048
53. Scherer C, Weiland S. Linear Matrix Inequalities in Control. DISC course lecture notes, 2000
54. Seiler P. LMILab translator. Available from http://vehicle.me.berkeley.edu/~guiness/lmitrans.html
55. Shor NZ. Minimization methods for nondifferentiable functions, vol 3. Springer Series in Computational Mathematics. Springer-Verlag, Berlin 1985
56. Sturm J. SeDuMi version 1.05, October 2002. Available from http://fewcal.kub.nl/sturm/software/sedumi.html
57. Tibken B, Dilaver KF. Computation of subsets of the domain of attraction for polynomial systems. In: Proceedings of the 41st IEEE Conference on Decision and Control, December 2002, pp 2651–2656
58. Todd M. Semidefinite optimization. Acta Numerica 2001; 10: 515–560
59. Toh KC, Tütüncü RH, Todd MJ. SDPT3 – a MATLAB software package for semidefinite-quadratic-linear programming. Available from http://www.math.cmu.edu/~reha/sdpt3.html
60. Vandenberghe L, Boyd S. Semidefinite programming. SIAM Rev March 1996; 38(1): 49–95
61. Vandenberghe L, Boyd S, Alkire B. SP Version 1.1. Software for Semidefinite Programming. http://www.ee.ucla.edu/~vandenbe/sp.html
62. Varga A, Parrilo PA. Fast algorithms for solving $\mathcal{H}_\infty$ norm minimization problems. In: Proceedings of the 40th IEEE Conference on Decision and Control, 2001
63. Willems JC. Least squares stationary optimal control and the algebraic Riccati equation. IEEE Trans Auto Control 1971; 16(6): 621–634
64. Wolkowicz H, Saigal R, Vandenberghe L. (eds). Handbook of semidefinite programming. Kluwer, 2000