



ELSEVIER

Contents lists available at ScienceDirect

C. R. Acad. Sci. Paris, Ser. I

www.sciencedirect.com



Numerical analysis

A successive constraint approach to solving parameter-dependent linear matrix inequalities



Une méthode de contraintes successives pour résoudre les inégalités matricielles linéaires paramétriques

Robert O'Connor

RWTH Aachen University, Aachen, Germany

ARTICLE INFO

Article history:

Received 26 January 2017

Accepted after revision 15 May 2017

Available online 30 May 2017

Presented by Olivier Pironneau

ABSTRACT

We present a successive constraint approach that makes it possible to cheaply solve large-scale linear matrix inequalities for a large number of parameter values. The efficiency of our method is made possible by an offline/online decomposition of the workload. Expensive computations are performed beforehand, in the offline stage, so that the problem can be solved very cheaply in the online stage. We also extend the method to approximate solutions to semidefinite programming problems.

© 2017 Académie des sciences. Published by Elsevier Masson SAS. All rights reserved.

R É S U M É

Nous présentons une méthode de contraintes successives qui réduit le travail nécessaire pour résoudre les inégalités matricielles linéaires paramétriques de grande dimension. Une caractéristique importante de notre méthode est la décomposition hors ligne/en ligne du travail. Les calculs coûteux sont effectués à l'avance, hors ligne, pour nous permettre de résoudre le problème de manière très économique en ligne. La même méthode est aussi appliquée à l'approximation des solutions des problèmes d'optimisation SDP.

© 2017 Académie des sciences. Published by Elsevier Masson SAS. All rights reserved.

1. Introduction

Linear matrix inequalities (LMIs) are a general type of convex constraint that includes linear as well as quadratic constraints [11,12] and lead to very natural formulations of a large number of problems in control and systems theory [2,4]. They can be solved using a wide range of existing methods [1,5,6,10], but that can be expensive for large-scale problems. In particular, problems resulting from the discretization of partial differential equations can be extremely expensive due to their high dimensionality. The computations are even more expensive if parameter-dependent problems are considered and

E-mail address: oc Connor@aices.rwth-aachen.de.

solutions are needed for a large number of parameter values. In that case traditional solution methods are extremely inefficient. We propose the construction of reduced-order models that take advantage of the parametric nature of the problem and allow us to very cheaply produce solutions for a large number of parameter values.

Let us introduce a finite-dimensional, bounded parameter domain $\mathcal{D} \in \mathbb{R}^p$ and the parameter-dependent LMI

$$F(x; \mu) := \sum_{q=1}^{Q_F} \left[\theta_q^0(\mu) + \theta_q^L(\mu)x \right] F_q \succeq 0. \quad (\text{L})$$

Here $F(x; \mu) \in \mathbb{R}^{\mathcal{N} \times \mathcal{N}}$ is a symmetric matrix that depends on both the parameter $\mu \in \mathcal{D}$ and the decision variable $x \in \mathbb{R}^n$ and is composed of Q_F parameter-independent matrices $F_q \in \mathbb{R}^{\mathcal{N} \times \mathcal{N}}$. The parameter dependencies of $F(x; \mu)$ are given by the functions $\theta^0(\cdot) : \mathcal{D} \rightarrow \mathbb{R}^{Q_F}$ and $\theta^L(\cdot) : \mathcal{D} \rightarrow \mathbb{R}^{Q_F \times n}$. We use subscripts to indicate components of a vector, such that $\theta_q^0(\mu)$ is the q th element of $\theta^0(\mu)$. Similarly, we write $\theta_q^L(\mu)$ to indicate the q th row of $\theta^L(\mu)$. The symbol \succeq will be used in the sense that $P \succeq 0$ indicates that the symmetric matrix P is positive semi-definite.

The goal of this paper is to efficiently solve the following problems for a large number of parameter values $\mu \in \mathcal{D}$:

- (i) the strict feasibility problem: find an $x \in \mathbb{R}^n$ such that $F(x; \mu) \succ 0$;
- (ii) the semidefinite program (SDP):

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad c(\mu)^T x \quad \text{subject to} \quad F(x; \mu) \succeq 0. \quad (\text{S})$$

Rather than directly solving these problems for each new parameter value, we will solve them for only a small set of intelligently chosen parameter values. We will then use the resulting solutions to build a reduced-order model that can approximate the solution anywhere in \mathcal{D} .

The method that we propose can be viewed as a generalization of the successive constraint method (SCM) [3,7], which is often used in the field of reduced basis methods to evaluate stability constants [9]. This method will allow us to very cheaply determine feasible solutions for any $\mu \in \mathcal{D}$. That is made possible by decomposing the computational workload into offline and online stages. All expensive computations will be performed in advance, during the offline stage. During the online stage the cost to solve the problem for a new parameter value will be independent of the size of the original constraint, \mathcal{N} . In that way the computational cost of each new solution will remain cheap even if the original constraint has very large dimensions.

Our methods are applicable to a wide range of LMIs and can also be used to extend the applicability of SCM. In the context of reduced basis methods, applications could involve bounding stability constants with respect to parameter-dependent norms or the selection of Lyapunov functions for the computation of error bounds [8]. In Section 4 we present an example in which we optimize a system while ensuring that it remains stable.

2. Reduced-order modeling for strict feasibility

SCM was originally designed to approximate coercivity constants. We will apply a modified version of SCM to the coercivity constant

$$\alpha(x; \mu) := \inf_{v \in \mathbb{R}^{\mathcal{N}}} \frac{v^T F(x; \mu) v}{v^T F_S v}, \quad (\text{1})$$

where $F_S \in \mathbb{R}^{\mathcal{N} \times \mathcal{N}}$ is a fixed symmetric positive-definite matrix. From the definition it is clear that $\alpha(x; \mu) \geq 0$ is equivalent to $F(x; \mu) \succeq 0$ for all symmetric positive definite matrices F_S . Nevertheless, an appropriate choice of F_S could be beneficial from a numerical point of view. If we are dealing with PDE discretizations, it can be advantageous to choose a matrix associated with an energy norm.

The first step in applying SCM is reformulating the coercivity constant as follows:

$$\alpha(x; \mu) = \inf_{y \in \mathcal{Y}} \left[\theta^0(\mu) + \theta^L(\mu)x \right]^T y, \quad \text{where} \quad \mathcal{Y} := \left\{ y \in \mathbb{R}^{Q_F} \mid y_q = \frac{v^T F_q v}{v^T F_S v}, v \in \mathbb{R}^{\mathcal{N}} \right\}. \quad (\text{2})$$

This formulation has the advantage that the complexity of the problem has been shifted to the definition of the set \mathcal{Y} . That allows us to compute lower and upper bounds for $\alpha(x; \mu)$ by approximating \mathcal{Y} .

A lower bound for $\alpha(x; \mu)$ can be derived by approximating \mathcal{Y} from the outside. A bounded but primitive approximation for \mathcal{Y} is given by

$$\mathcal{B}_Q := \prod_{i=1}^{Q_F} \left[\inf_{y \in \mathcal{Y}} y_q, \sup_{y \in \mathcal{Y}} y_q \right] \subset \mathbb{R}^{Q_F}. \quad (\text{3})$$

\mathcal{B}_Q will generally be much larger than \mathcal{Y} so we will add constraints to restrict it and better approximate \mathcal{Y} . Let us assume that we have computed lower bounds $\bar{\alpha}(\bar{x}; \bar{\mu}) \leq \alpha(\bar{x}; \bar{\mu})$ for each pair $(\bar{x}, \bar{\mu})$ in a predetermined set $\mathcal{R} \subset \mathbb{R}^n \times \mathcal{D}$. The set \mathcal{Y} is then contained in the set

$$\mathcal{Y}_{\text{out}} := \left\{ y \in \mathcal{B}_Q \mid \left[\theta^0(\bar{\mu}) + \theta^L(\bar{\mu})\bar{x} \right]^T y \geq \bar{\alpha}(\bar{x}; \bar{\mu}), \forall (\bar{x}, \bar{\mu}) \in \mathcal{R} \right\}. \tag{4}$$

Using \mathcal{Y}_{out} , we can define

$$\alpha_{\text{out}}(x; \mu) := \inf_{y \in \mathcal{Y}_{\text{out}}} \left[\theta^0(\mu) + \theta^L(\mu)x \right]^T y, \tag{5}$$

such that $\alpha_{\text{out}}(x; \mu) \leq \alpha(x; \mu)$ for any $x \in \mathbb{R}^n$ and any $\mu \in \mathcal{D}$.

The lower bound $\alpha_{\text{out}}(x; \mu)$ can be formulated as the solution to a linear programming problem. To see that we note that \mathcal{Y}_{out} is a linearly constrained subset of \mathbb{R}^{Q_F} . As a result we can find a matrix $A_{\text{out}} \in \mathbb{R}^{\ell \times Q_F}$ and a vector $b_{\text{out}} \in \mathbb{R}^{\ell}$ such that $\mathcal{Y}_{\text{out}} = \{y \in \mathbb{R}^{Q_F} \mid A_{\text{out}}y \geq b_{\text{out}}\}$. Here the symbol \geq indicates a componentwise comparison of vectors and $\ell \ll \mathcal{N}$ is the number of constraints. The lower bound $\alpha_{\text{out}}(x; \mu)$ can be calculated, for any $x \in \mathbb{R}^n$ and any $\mu \in \mathcal{D}$, as the solution to the linear program (P) or its dual (D):

| | |
|---|---|
| $\begin{aligned} & \text{minimize}_{y \in \mathbb{R}^{Q_F}} \quad \left[\theta^0(\mu) + \theta^L(\mu)x \right]^T y \\ & \text{subject to} \quad A_{\text{out}}y \geq b_{\text{out}} \end{aligned} \tag{P}$ | $\begin{aligned} & \text{maximize}_{p \in \mathbb{R}^{\ell}} \quad b_{\text{out}}^T p \\ & \text{subject to} \quad A_{\text{out}}^T p = \theta^0(\mu) + \theta^L(\mu)x \\ & \quad \quad \quad p \geq 0 \end{aligned} \tag{D}$ |
|---|---|

To find an optimal value of x we will maximize (D) over all $x \in \mathbb{R}^n$ using the problem

$$\text{maximize}_{p \in \mathbb{R}^{\ell}, x \in \mathbb{R}^n} b_{\text{out}}^T p \quad \text{subject to} \quad A_{\text{out}}^T p - \theta^L(\mu)x = \theta^0(\mu) \quad \text{and} \quad p \geq 0. \tag{RF}$$

If $b_{\text{out}}^T p$ is strictly greater than 0, then the associated x is guaranteed to strictly satisfy (L). Even if that is not the case, (RF) is always feasible. To show that we note that the boundedness of $\mathcal{Y}_{\text{out}} \subset \mathcal{B}_Q$ implies the boundedness of (P). By duality (D) is then feasible for any $x \in \mathbb{R}^n$ and any $\mu \in \mathcal{D}$ which implies that (RF) is feasible for any $\mu \in \mathcal{D}$. (RF) is also bounded if $\max_{x \in \mathbb{R}^n} \alpha(x; \mu)$ is bounded; otherwise, we will be content with any x associated with a large, positive value of $b_{\text{out}}^T p$. The main reason to use (RF) is that it is cheap to solve: It is a small linear program with $n + \ell$ variables, Q_F equality constraints, and ℓ inequality constraints.

Since it is not possible to work directly with the infinite set \mathcal{D} , we introduce a large set $\Xi \subset \mathcal{D}$ of discrete points that are representative of \mathcal{D} and a much smaller set $\mathcal{C}_k \subset \Xi$ with cardinality k . The set \mathcal{R} will be made up of parameter values $\bar{\mu} \in \Xi$ and associated approximate solutions $\bar{x} \in \mathbb{R}^n$. For parameter values $\bar{\mu}$ that are in \mathcal{C}_k we will use solutions \bar{x} to the original problem (L) and set $\bar{\alpha}(\bar{x}; \bar{\mu}) = \alpha(\bar{x}; \bar{\mu})$. For parameter values in the much larger set $\Xi \setminus \mathcal{C}_k$ we will use solutions \bar{x} to (RF), and set $\bar{\alpha}(\bar{x}; \bar{\mu}) = \alpha_{\text{out}}(\bar{x}; \bar{\mu})$. The pairs $(\bar{x}, \bar{\mu})$ that will be included in \mathcal{R} will depend on the parameter μ for which we would like a feasible solution. For two natural numbers $M_{\mathcal{C}}$ and M_{Ξ} we define \mathcal{R} to be a set of $M_{\mathcal{C}}$ pairs $(\bar{x}, \bar{\mu})$ with $\bar{\mu} \in \mathcal{C}_k$ and M_{Ξ} pairs $(\bar{x}, \bar{\mu})$ with $\bar{\mu} \in \Xi \setminus \mathcal{C}_k$. In both cases we will choose the pairs $(\bar{x}, \bar{\mu})$ with $\bar{\mu}$ closest to μ in a predetermined metric. If $M_{\mathcal{C}} > k$, we simply include all of \mathcal{C}_k in \mathcal{R} .

To build our model we will make use of the greedy algorithm that was introduced in the context of reduced basis methods [13] and plays a vital role in SCM. The algorithm is initiated by choosing a small initial set \mathcal{C}_k . For each $\bar{\mu} \in \mathcal{C}_k$ the problem (L) is solved once to determine the associated values of \bar{x} and $\bar{\alpha}(\bar{x}; \bar{\mu})$. For $\bar{\mu} \in \Xi \setminus \mathcal{C}_k$ we initially set $\bar{x} = 0$ and $\bar{\alpha}(\bar{x}; \bar{\mu}) = -\infty$. The model is then improved iteratively. In each iteration we will solve (RF) for each $\bar{\mu} \in \Xi \setminus \mathcal{C}_k$ and update the stored values of \bar{x} and $\bar{\alpha}(\bar{x}; \bar{\mu})$. The $\bar{\mu}$ associated with the smallest value of $\bar{\alpha}(\bar{x}; \bar{\mu})$ is then added to \mathcal{C}_k , and (L) is solved to update the stored values of \bar{x} and $\bar{\alpha}(\bar{x}; \bar{\mu})$. The process terminates when the smallest value of $\bar{\alpha}(\bar{x}; \bar{\mu})$ is larger than some positive tolerance.

Solving problems with our method involves two stages. During the offline stage the greedy algorithm builds the model. The expensive operations in the offline stage are k solutions of (L) and nearly $k|\Xi|$ solutions of (RF). Here k is the cardinality of the final \mathcal{C}_k and $|\Xi|$ is the cardinality of Ξ . During the online stage (RF) can be constructed and solved very cheaply for any new parameter value. The online computational cost is independent of \mathcal{N} to ensure that it remains cheap even if \mathcal{N} is very large.

During the greedy algorithm it is necessary to solve (L) for various parameter values. That can be done using a variety of methods including algorithms that solve semidefinite programming problems [1,5,6,10]. In this section we propose a method that is based on SCM and allows us to reuse information that we have already computed.

For a predetermined, finite subset \mathcal{Y}_{in} of \mathcal{Y} we define

$$\alpha_{\text{in}}(x; \mu) := \inf_{\bar{y} \in \mathcal{Y}_{\text{in}}} \left[\theta^0(\mu) + \theta^L(\mu)x \right]^T \bar{y}, \tag{6}$$

such that $\alpha_{\text{in}}(x; \mu) \geq \alpha(x; \mu)$ for all $x \in \mathbb{R}^n$ and $\mu \in \mathcal{D}$. Here we will make use of a variation of the greedy algorithm. In each iteration we solve $\max_{x \in \mathbb{R}^n} \alpha_{\text{in}}(x; \mu)$ to find the optimal x . For that value of x we then compute $\alpha(x; \mu)$ and an element y of \mathcal{Y} for which the infimum in (2) is reached. That y is then added to \mathcal{Y}_{in} to improve the upper bound $\alpha_{\text{in}}(\cdot; \mu)$. The algorithm terminates when $\alpha(x; \mu)$ is sufficiently large or sufficiently close to $\alpha_{\text{in}}(x; \mu)$.

In each iteration of the algorithm both the small linear program associated with $\max_{x \in \mathbb{R}^n} \alpha_{\text{in}}(x; \mu)$ and the \mathcal{N} -dimensional eigenvalue problem associated with $\alpha(x; \mu)$ need to be solved once. The dominant cost is that of the eigenvalue problems. To reduce the number of iterations and eigenvalue solves we will store the last value of \bar{y} that is calculated each time we solve (L) for a new $\mu \in \mathcal{C}_k$. The set \mathcal{Y}_{in} can then be initialized using those values of \bar{y} .

3. The semidefinite programming problem

In this section we will show how our methods can be used to approximate solutions to the semidefinite program (S). The idea is to minimize $c(\mu)^T x$ over all $x \in \mathbb{R}^n$ that satisfy $\alpha_{\text{out}}(x; \mu) \geq 0$. That is done using the following problem:

$$\underset{x \in \mathbb{R}^n, p \in \mathbb{R}^\ell}{\text{minimize}} \quad c(\mu)^T x \quad \text{subject to} \quad \begin{bmatrix} A_{\text{out}}^T & -\theta^L(\mu) \end{bmatrix} \begin{bmatrix} p \\ x \end{bmatrix} = \theta^0(\mu) \quad \text{and} \quad \begin{bmatrix} I \\ b_{\text{out}}^T \end{bmatrix} p \geq \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \quad (\text{RS})$$

Here A_{out} , b_{out}^T and $\alpha_{\text{out}}(x; \mu)$ are all the same as in Section 2 except that the stored values of \bar{x} will be calculated using (S) and (RS) rather than (L) and (RF). The new constraint $b_{\text{out}}^T p \geq 0$ guarantees that $\alpha_{\text{out}}(x; \mu) \geq 0$.

Let us write \mathcal{J}_{out} and \mathcal{J} to denote the optimal values of (RS) and (S), respectively. Based on the fact that the optimal $x \in \mathbb{R}^n$ for (RS) will be feasible for (S) we know that $\mathcal{J}_{\text{out}} \geq \mathcal{J}$. Ideally \mathcal{J}_{out} will also be a good approximation to \mathcal{J} despite being much cheaper to compute: It requires the solution of a linear program with $n + \ell$ variables, Q_F equality constraints, and $\ell + 1$ inequality constraints.

For any $\mu \in \mathcal{D}$, the boundedness of (S) implies the boundedness of (RS), but the question of feasibility is more complicated. To ensure feasibility we first build a model (RF) to find strictly feasible solutions. We then convert that model to the form of (RS). In doing so we consider the stored values of $\bar{x} \in \mathbb{R}^n$, which should be feasible, to be approximate solutions to (S) and we reinitialize \mathcal{C}_k as an empty set. We can then improve the model using another greedy algorithm.

The goal of this greedy algorithm is to reduce the error $\mathcal{J}_{\text{out}} - \mathcal{J}$. Since it is too expensive to compute \mathcal{J} online, we will compute a lower bound for it by considering the problem

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad c(\mu)^T x \quad \text{subject to} \quad \alpha_{\text{in}}(x; \mu) \geq 0. \quad (\text{ER})$$

Here $\alpha_{\text{in}}(x; \mu)$ is defined like in (6) with \mathcal{Y}_{in} initially being the set of $\bar{y} \in \mathcal{Y}$ associated with parameter values in \mathcal{C}_k . The optimal value of (ER), which we will denote \mathcal{J}_{in} , is by definition a lower bound for \mathcal{J} . That allows us to bound the error $\mathcal{J}_{\text{out}} - \mathcal{J} \geq 0$ from above using $\mathcal{J}_{\text{out}} - \mathcal{J}_{\text{in}}$. Like the previous approximations, this bound can be computed online with a number of computations that is independent of \mathcal{N} . Although (ER) is always feasible, it could be unbounded. That can be fixed by increasing \mathcal{C}_k and consequently \mathcal{Y}_{in} .

During each iteration of the greedy algorithm we will solve both (RS) and (ER), compute the difference $\mathcal{J}_{\text{out}} - \mathcal{J}_{\text{in}}$, and update both the estimates $\bar{x} \in \mathbb{R}^n$ and the associated values $\bar{\alpha}(\bar{x}; \bar{\mu})$ for each $\bar{\mu} \in \Xi \setminus \mathcal{C}_k$. The $\bar{\mu}$ that produced the maximum value of $\mathcal{J}_{\text{out}} - \mathcal{J}_{\text{in}}$ is then added to \mathcal{C}_k and a more accurate solution to (S) is computed. That allows us to update both (RS) and (ER). This process terminates when the maximum value of $\mathcal{J}_{\text{out}} - \mathcal{J}_{\text{in}}$ is below a desired tolerance.

The major computational burdens in the offline stage of our method are the construction of the model (RF), k solutions of (S), and nearly $k|\Xi|$ solutions of both (RS) and (ER). During the online stage an approximate solution can be determined for any given $\mu \in \mathcal{D}$ by simply constructing and solving (RS). The error can also be bounded online by solving (ER).

The more accurate solutions that we need can be computed using either general solvers for semidefinite programming problems [1,5,6,10] or a method that is similar to the one presented in Section 2. For the latter we will consider the following problem with a fixed value of $\mu \in \mathcal{D}$:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad c(\mu)^T x \quad \text{subject to} \quad \alpha_{\text{in}}(x; \mu) > \alpha_{\text{min}}, \quad (\text{T})$$

where $\alpha_{\text{min}} > 0$ is some small constant that is used as a tolerance. We again use a sort of greedy algorithm. In each iteration we solve (T) to get the approximate solution x and update \mathcal{Y}_{in} by adding the new value of y . The process is repeated until the stopping condition $\alpha(x; \mu) > 0$ is satisfied. If α_{min} is sufficiently small, the resulting x should be a good approximation to the true solution.

4. Numerical example: system stabilization

To test our method we will consider a reaction–diffusion equation on the unit square, Ω , depicted in Fig. 1(a). To facilitate the explanation we consider a problem with just one parameter $\mu \in \mathcal{D} := [0, 3]$ and one decision variable $x \in \mathbb{R}$. The problem is related to the following system:




Fig. 1. Setup of the numerical example.

Fig. 1. Configuration de l'exemple numérique.




Fig. 2. Results of the numerical experiments.

Fig. 2. Résultats des expériences numériques.

$$\dot{y} = \Delta y + \mu y \mathbf{1}_{\Omega_1} + \mathbf{1}_{\Omega_2} u, \quad \forall z \in \Omega; \quad \frac{\partial y}{\partial \eta} = 0, \quad \forall z \in \partial \Omega; \quad u = -x \zeta; \quad \zeta = \int_{\Omega_2} y. \quad (7)$$

Here $\mathbf{1}_S$ indicates the characteristic function of a set S and η is the outward-pointing unit vector on the boundary $\partial \Omega$. The decision variable x determines the feedback gain from the system's output $\zeta(t)$ to its input $u(t)$ as depicted in Fig. 1(b). We use linear finite elements to discretize the problem in space. That gives us the following 2601-dimensional closed-loop semi-discrete system

$$M \dot{y} + A(x; \mu) y = 0, \quad \text{where } A(x; \mu) = A_0 - \mu A_1 + x A_2, \quad \text{with } M, A_0, A_1, A_2 \in \mathbb{R}^{2601 \times 2601}. \quad (8)$$

Here M is symmetric positive definite and $A(x; \mu)$ is symmetric. Under these assumptions the system (8) is strictly stable for a given $\mu \in \mathcal{D}$ and $x \in \mathbb{R}$ iff $A(x; \mu) > 0$. Our goal will be to minimize the cost of the control, by keeping x small, while ensuring strict stability. Noting that A_0 , which is associated with the operator $-\Delta$, is positive semidefinite and that $A_0 + A_1$ is positive definite, we define $F(x; \mu) := (1 - \rho) A_0 + (-\mu - \rho) A_1 + x A_2$ and $F_S := A_0 + A_1 > 0$ for some small $\rho > 0$. It then holds that

$$F(x; \mu) \geq 0 \quad \iff \quad A(x; \mu) - \rho F_S \geq 0, \quad (9)$$

and hence the system in (8) is strictly stable if $F(x; \mu) \geq 0$. We will search for a minimal stabilizing gain $x \in \mathbb{R}$ as the solution to (S).

For our model we will use values of $M_C = 4$ and $M_\Xi = 3$. C_k is initialized using the smallest and the largest values in \mathcal{D} . For this particular problem that guarantees that (RS) is feasible for all $\mu \in \mathcal{D}$. As a result we can start directly with (RS) and do not need to build the model given in (RF). Another modification of our method is that we will work with the relative error given by $(\mathcal{J}_{\text{out}} - \mathcal{J}_{\text{in}}) / \mathcal{J}_{\text{in}}$ rather than the error $\mathcal{J}_{\text{out}} - \mathcal{J}_{\text{in}}$. That is possible because \mathcal{J}_{in} will always be strictly positive.

Fig. 2(a) shows the convergence of the solver described in Section 3 for the computation of more accurate solutions to (S). Here the error is measured as $\alpha_{\text{in}}(\bar{x}; \bar{\mu}) - \alpha(\bar{x}; \bar{\mu})$. Two examples are shown as well as the worst case over 30 random parameter values. We recall that the most expensive part of this method is the eigenvalue solves and that one such solve is needed for each iteration. Fig. 2(b) shows the convergence of the greedy algorithm for two different parameter domains \mathcal{D} . The plotted values are the worst relative errors over the respective sets Ξ . For $\mathcal{D} = [0, 3]$ we set Ξ to be 300 uniformly distributed points and for $\mathcal{D} = [0, 1.5]$ we use 150 points. Fig. 2(c) compares the results when C_k is constructed using the greedy algorithm or simply as uniformly distributed points. The plot shows the worst relative errors over 100 random values of $\mu \in \mathcal{D}$. The greedy algorithm performs slightly better than uniform distributions when C_k is large and also has the advantage that it is iterative.

References

- [1] F. Alizadeh, J.-P.A. Haeberly, M.L. Overton, Primal–dual interior–point methods for semidefinite programming: convergence rates, stability and numerical results, *SIAM J. Optim.* 8 (3) (1998) 746–768.
- [2] S. Boyd, V. Balakrishnan, E. Feron, L. ElGhaoui, Control system analysis and synthesis via linear matrix inequalities, in: *American Control Conference*, June 1993, pp. 2147–2154.
- [3] Y. Chen, J.S. Hesthaven, Y. Maday, J. Rodríguez, A monotonic evaluation of lower bounds for inf–sup stability constants in the frame of reduced basis approximations, *C. R. Acad. Sci. Paris, Ser. I* 346 (23) (2008) 1295–1300.
- [4] E. Feron, P. Apkarian, P. Gahinet, Analysis and synthesis of robust control systems via parameter-dependent Lyapunov functions, *IEEE Trans. Autom. Control* 41 (7) (1996) 1041–1046.
- [5] M. Fukuda, M. Kojima, K. Murota, K. Nakata, Exploiting sparsity in semidefinite programming via matrix completion I: general framework, *SIAM J. Optim.* 11 (3) (2001) 647–674.
- [6] C. Helmsberg, F. Rendl, A spectral bundle method for semidefinite programming, *SIAM J. Optim.* 10 (3) (2000) 673–696.
- [7] D.B.P. Huynh, G. Rozza, S. Sen, A.T. Patera, A successive constraint linear optimization method for lower bounds of parametric coercivity and inf–sup stability constants, *C. R. Acad. Sci. Paris, Ser. I* 345 (8) (2007) 473–478.
- [8] R. O'Connor, Lyapunov-based error bounds for the reduced-basis method, *IFAC-PapersOnLine* 49 (8) (2016) 1–6.
- [9] G. Rozza, D.B.P. Huynh, A.T. Patera, Reduced basis approximation and a posteriori error estimation for affinely parametrized elliptic coercive partial differential equations, *Arch. Comput. Methods Eng.* 15 (3) (2008) 229–275.
- [10] K.-C. Toh, Solving large scale semidefinite programs via an iterative solver on the augmented systems, *SIAM J. Optim.* 14 (3) (2004) 670–698.
- [11] J.G. VanAntwerp, R.D. Braatz, A tutorial on linear and bilinear matrix inequalities, *J. Process Control* 10 (4) (2000) 363–385.
- [12] L. Vandenberghe, S. Boyd, Semidefinite programming, *SIAM Rev.* 38 (1) (1996) 49–95.
- [13] K. Veroy, C. Prud'homme, D.V. Rovas, A.T. Patera, A posteriori error bounds for reduced-basis approximation of parametrized noncoercive and nonlinear elliptic partial differential equations, in: *Proceedings of the 16th AIAA Computational Fluid Dynamics Conference*, 2003, AIAA Paper 2003-3847.