Proceedings of the 38th
Conference on Decision & Control
Phoenix, Arizona USA • December 1999

On the numerical solution of LMIs derived from the KYP lemma

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

Abstract

An efficient algorithm for the numerical solution of LMIs arising from the Kalman-Yakubovich-Popov lemma is presented. The procedure is an outer approximation method based on the algorithms used in the computation of $\mathcal{H}_\infty$ norms for LTI systems. The result is especially useful for systems with large state dimension.

1 Introduction

The KYP lemma (Kalman-Yakubovich-Popov [10], see [9] for an elementary proof) establishes the equivalence between a frequency domain inequality and the feasibility of a particular kind of LMI (linear matrix inequality). It is an important generalization of classical linear control results, such as the bounded real and positive real lemma.

It is also a fundamental tool in the practical application of the IQC (integral quadratic constraints) framework [7] to the analysis of uncertain systems. The theorem replaces an infinite family of LMIs, parameterized by ω, by a finite dimensional problem. This is extremely useful from a practical viewpoint, since it allows for the use of standard finite dimensional LMI solvers.

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

In this paper, we present an efficient algorithm for the solution of this type of inequalities. The approach is an *outer approximation* method [8], and is based on the algorithms used in the computation of $\mathcal{H}_\infty$ system norms. The idea is to impose the frequency domain inequality (1) only at a discrete number of frequencies. These frequencies are then updated by a mechanism similar to the $\mathcal{H}_\infty$ norm computation case.

Previous related work includes of course the literature on the computation of $\mathcal{H}_\infty$ system norms. In particular, references [2, 4, 1] developed quadratically convergent algorithms, based explicitly on the Hamiltonian approach. Also, a somewhat related approach in [6] implements a cutting-plane based algorithm, where linear constraints are imposed on the optimization variables.

The paper is organized as follows: in Section 2 the notation is presented and some basic facts are reviewed. In Section 3 the algorithm is presented and analyzed, and its convergence properties are considered. In the following section, the use of the dual problem is discussed, and in Section 5 some examples are presented. Finally, in the last section we present conclusions and outlines for future research.

2 Preliminaries

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

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

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

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

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

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

0-7803-5250-5/99/$10.00 © 1999 IEEE

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

Lemma 1 *Let Q be a partitioned matrix*

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

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

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

A fairly general version of the KYP lemma, as presented in [9] is the following:

Theorem 1 *Let $A \in \mathbb{R}^{n\times n}, B \in \mathbb{R}^{n\times m}, M = M^T \in \mathbb{R}^{(n+m)\times(n+m)}$, with A having no purely imaginary eigenvalues. Let*

$$F(j\omega) \triangleq \begin{bmatrix} (j\omega I - A)^{-1}B \\ I \end{bmatrix}^* M \begin{bmatrix} (j\omega I - A)^{-1}B \\ I \end{bmatrix}.$$

Then, the two following statements are equivalent:

1. $$F(j\omega) < 0, \quad \forall \omega \in \mathbb{R} \cup \{\infty\} \tag{1}$$

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

$$\begin{bmatrix} A^TP + PA & PB \\ B^TP & 0 \end{bmatrix} + M < 0 \tag{2}$$

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

In this paper, we will deal only with the strict version of the KYP lemma, i.e. with a strict inequality in (1), (2). The reason is twofold: in the first place, no controllability/stabilizability assumptions are necessary, simplifying the proofs. Secondly, since the resulting LMIs will in general be solved using interior-point methods, the existence of a strictly feasible solution is usually guaranteed. Note also that the strict inequality implies (in both (1) and (2)) that $M_{22} < 0$.

3 The Algorithm

The basic idea is to replace the semi-infinite optimization problem (1) by a finite dimensional relaxation. We choose to impose the constraint only at a finite number of frequencies $\omega_k \in \Omega$ (see [5] for a related approach). For a given ω, equation (1) is an LMI in M.

A high-level description of the algorithm follows:

Algorithm 1

1. *Initialize the set of frequencies $\Omega \triangleq \{0\}$.*
2. *Solve (1) with the current Ω set.*
3. *Find a frequency ω_k where the constraint (1) is violated (up to an ϵ tolerance). If no such frequency exists, exit.*
4. *Add that frequency to the set Ω, and go to step 2.*

As we can see, the underlying idea of an outer approximation algorithm is a generalization of a *cutting plane* method [8]. We replace the description of the feasible set by a convenient relaxation. If the resulting solution does not satisfy the original constraints, a cutting plane (in this case, a possibly curved hypersurface) that separates that solution from the true feasible set is added. The process is repeated until the desired tolerance is reached.

As in the case of $\mathcal{H}_\infty$ norm computation [2, 4] the effectiveness of the algorithm hinges on the possibility of detecting in an efficient manner the frequencies at which the inequality is violated. To this end, define the $2n \times 2n$ Hamiltonian matrix:

$$H = \begin{bmatrix} A - BM_{22}^{-1}M_{21} & -BM_{22}^{-1}B^T \\ -M_{11} + M_{12}M_{22}^{-1}M_{21} & -A^T + M_{12}M_{22}^{-1}B^T \end{bmatrix} \tag{3}$$

It can be shown (see for example [10, 11]) that the conditions (1), (2) are satisfied if and only if $M_{22} < 0$ and H has no imaginary eigenvalues. In this case, it is possible to obtain a solution P of the LMI (1) by computing a solution of the Riccati equation associated with the Hamiltonian (or a suitable perturbation, if the subspace complementarity condition is not satisfied). If the eigenvalue condition is violated, then there is a relationship between the critical frequencies, as the following theorem shows.

Theorem 2 *Assume $M_{22} < 0$. Then, $F(j\omega_0)$ is singular, if and only if $j\omega_0$ is an imaginary eigenvalue of H.*

Proof: *Consider the partitioned matrix*

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

The diagonal submatrices are invertible, since A has no imaginary eigenvalues and $M_{22} < 0$. Applying Lemma 1, we immediately have the identity

$$\begin{aligned} \det Q &= \det(j\omega I - H) \det M_{22} \\ &= \det(j\omega I + A^T) \det F(j\omega) \det(j\omega I - A). \end{aligned}$$

from where the result follows. □

Special cases of this theorem are the ones used in [2] to compute the $\mathcal{H}_\infty$ norm or the minimum dissipation of a transfer function.

Several options are available for the choice of the frequency to add to the set Ω. A particularly good one is to choose ω_k as the frequency at which $F(j\omega)$ is maximally positive (i.e., where its first singular value achieves its maximum over frequency). This can be obtained at a computational cost similar to that of an $\mathcal{H}_\infty$ norm. In the following section we present a convergence argument for the procedure resulting from this choice. A cheaper alternative is to pick a criterion similar to the one used in [4]. Given the imaginary eigenvalues of H, consider the midpoint frequencies, and choose the one where the constraint is most violated. The computational requirements of this step are minimal, compared to the one required to solve the LMIs.

An important difference of the LMI case discussed here with the simpler $\mathcal{H}_\infty$ norm case (where the *only* LMI variable is the KYP one) is that at optimality more than one constraint can be active. In fact, the results in [5] show that at most $n + 1$ frequencies are active, where n is the number of IQCs.

In the algorithm as described, no constraint dropping occurs. That is, we keep adding constraints, until convergence. Since we know *a priori* a bound on the number of active constraints, dropping old, nonbinding constraints seems a natural idea. This issue will be explored in more detail in future versions of this paper.

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

A somewhat related approach is used in [6], where the eigenvectors of the Hamiltonian are used to construct linear constraints for the elements of M. In our approach, the constraints are matrix valued (not linear) and we do not impose the restrictions directly at the critical frequencies, but at other points where they are more violated. This way, convergence should be improved (in the $\mathcal{H}_\infty$ case, it is even quadratic). Further numerical experiments are needed to formulate accurate comparisons.

3.1 Convergence

It is possible to prove convergence of the first version of the algorithm. This corresponds to the choice of ω_k as the point at which the frequency domain inequality is maximally violated. In fact, for this variation we can apply the results on the convergence of more abstract version of the outer approximation method (Conceptual Algorithm 3.5.19 in [8]).

It is possible to show (see [8]) that if the algorithm produces a infinite sequence of solutions, then any accumulation point of this sequence is a global solution of the original problem. The infinite set of frequency constraints can be "compactified" either by considering the extended real line or by a standard bilinear transformation.

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

4 Using the dual

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

This can be solved by focusing instead on the dual optimization problem, as is well known from the linear programming literature, for instance. In this case, new *variables* are added to the problem at each iteration. Note that this can also be interpreted as having a dual feasible starting point, which is useful in case we are using a primal-dual LMI solver (such as SDPSOL [3]).

For the frequency domain inequalities arising from IQC optimization, the dual problem has been extensively analyzed in [5]. It has been shown there that upper bounds, or even the optimal value, of the quantities of interest (for example, $\mathcal{L}_2$-induced norms) can be obtained from a finite number of frequencies. However, no procedure to compute or approximate these frequencies was available, other than a standard gridding.

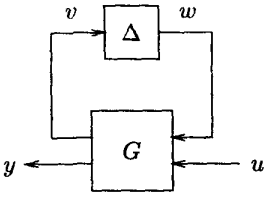

Figure 1: Standard block diagram.

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

Table 1: Numerical values for Example 1.

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

5 Example

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

5.1 First example

Consider the standard block diagram in Fig. 1. We will use the proposed algorithm to compute the worst case $\mathcal{L}_2$ induced norm between u and y, for the plant given by

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

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

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

The results of the sequence of subproblems are shown in Table 1 and Fig. 2.

As we can see, on the third and last iteration we obtain a value of the parameters that makes the frequency domain inequality to be satisfied. That makes possible, if desired, to recover the value of the optimal KYP variable P, by solving a Riccati equation. In this case, we obtain

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

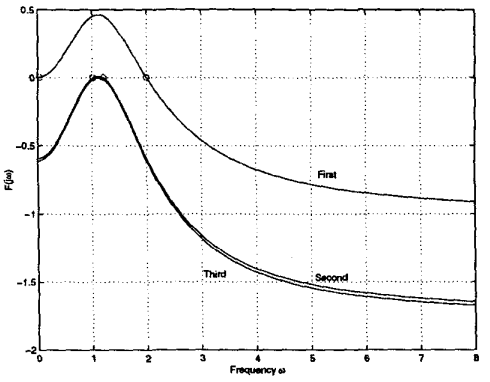

Figure 2: Frequency domain plots corresponding to Example 1.

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

Table 2: Numerical values for Example 2.

This is within numerical error of the solution obtained by directly solving the LMI (1).

5.2 Second example

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

The system is again in the standard form of Fig. 1. The plant G, chosen randomly, has 50 states, and the signals u, y, v, w are vector-valued, with each having 10 components. The uncertainty Δ corresponds to a diagonal gain bounded LTV operator, and therefore there are 10 IQCs associated with it.

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

Note that in this example, as opposed to the previous

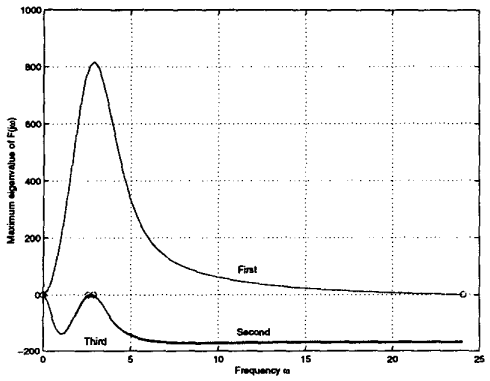

Figure 3: Frequency domain plots corresponding to Example 2.

one, more than one constraint is active at optimality. A result from [5] is that at most $n+1$ frequencies are active, so this is consistent with the expected behavior.

Finally, we remark that even though we are currently using a relatively inefficient implementation (since we are not using the information obtained in earlier stages in the solution of the subproblems), the algorithm still outperforms the standard approach. A future version of the paper will include numerical results for the dual approach.

6 Conclusions

The paper presents a new and efficient algorithm for the solution of KYP-derived LMIs, based on ideas from $\mathcal{H}_\infty$ norm computation methods.

The results are particularly interesting in the case where the dimension of the KYP variable is large when compared to the number of available IQCs. This corresponds to systems with large state dimension, where the standard KYP approach can be practically useless. Any advance in methods for LMI solving can be applied immediately, providing improved performance, since the subproblems are themselves standard LMIs.

Future work will address a robust numerical implementation of the outlined algorithm, where the partial solutions of previous stages are used in the solution of the subsequent problems. The important practical issues of constraints dropping and the necessity of a strictly feasible initial starting point should be addressed. It would also be interesting to have further insight on the exact conditions under which this approach is superior to solving directly for the KYP variable.

References

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

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

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

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

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

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

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

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

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

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

[11] K. Zhou, K. Glover, and J. C. Doyle. *Robust and Optimal Control.* Prentice Hall, 1995.