Nonsmooth H^∞ Synthesis

Pierre Apkarian, Dominikus Noll

► To cite this version:

Pierre Apkarian, Dominikus Noll. Nonsmooth H^∞ Synthesis. IEEE Transactions on Automatic Control, 2006, 51 (1), pp.71-86. 10.1109/TAC.2005.860290 . hal-01868781

HAL Id: hal-01868781

<https://hal.science/hal-01868781>

Submitted on 5 Sep 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Nonsmooth H_∞ Synthesis

Pierre Apkarian, *Member, IEEE*, Dominikus Noll, *Member, SIAM*,

Abstract—We develop nonsmooth optimization techniques to solve H_∞ synthesis problems under additional structural constraints on the controller. Our approach avoids the use of Lyapunov variables and therefore leads to moderate size optimization programs even for very large systems. The proposed framework is versatile and can accommodate a number of challenging design problems including static, fixed-order, fixed-structure, decentralized control, design of PID controllers and simultaneous design and stabilization problems.

Our algorithmic strategy uses generalized gradients and bundling techniques suited for the H_∞ norm and other nonsmooth performance criteria. We compute descent directions by solving quadratic programs and generate steps via line search. Convergence to a critical point from an arbitrary starting point is proved and numerical tests are included to validate our methods. The proposed approach proves to be efficient even for systems with several hundreds of states.

Index Terms— H_∞ -synthesis, static output feedback, fixed-order synthesis, simultaneous stabilization, NP-hard problems, nonsmooth optimization, bundle methods, Clarke subdifferential, bilinear matrix inequality (BMI), linear matrix inequality (LMI).

I. INTRODUCTION

In this paper we consider H_∞ synthesis problems with additional structural constraints on the controller. This includes static and reduced-order H_∞ output feedback control, structured, sparse or decentralized synthesis, simultaneous stabilization problems, multiple performance channels, and much else. We propose to solve these problems with a nonsmooth optimization method exploiting the structure of the H_∞ norm.

In nominal H_∞ synthesis, feedback controllers are computed via semidefinite programming (SDP) [1] or algebraic Riccati equations [2]. When structural constraints on the controller are added, the H_∞ synthesis problem is no longer convex. Some of the problems above have even been recognized as NP-hard [3] or as rationally undecidable [4]. These mathematical concepts indicate the inherent difficulty of H_∞ synthesis under constraints on the controller.

Even with structural constraints, the bounded real lemma may still be brought into play. The difference with customary H_∞ synthesis is that it no longer produces LMIs, but bilinear matrix inequalities, BMIs, which are genuinely non-convex. Optimization code for BMI problems is currently developed by several groups, see e.g. [5]–[9], but it appears that the BMI approach runs into numerical difficulties even for problems of moderate size. This is mainly due to the presence of Lyapunov variables, whose number grows quadratically with the number of states.

Our present approach does *not* use the bounded real lemma and thereby avoids Lyapunov variables. This leads to moderate size optimization programs even for very large systems. In exchange, our cost functions are nonsmooth and require special optimization techniques, which we develop here. We evaluate the H_∞ norm via the Hamiltonian bisection algorithm [10]–[12] and exploit it further to compute subgradients, which are then used to compute descent steps. Notice, however, that our method is not a pure frequency domain method. In fact, it allows *both* frequency domain *and* state space domain parameterizations of the unknown controller. This makes it a very flexible tool in a number of situations of practical interest.

Several iterative methods for reduced-order control have been proposed over recent years, see for instance [13]–[15]. In [13], a comparison among four of these methods on a large set of test problems is arranged, with the result that successive linearization [14], also known as the Frank and Wolfe (FW) algorithm [16], performed best. Whenever possible, we have therefore compared our new nonsmooth methods and the augmented Lagrangian algorithm in [17], [18] with the Frank and Wolfe method. The results are presented in the experimental section.

As far as comparison with existing methods is concerned, let us mention that for specific classes of plants, it is possible to compute reduced-order controllers without the use of optimization techniques. This has for instance been investigated in [19]–[21]. These approaches usually make strong additional assumptions like singularity, or hypotheses about unstable invariant zeros. In such cases it may then even be possible to assure global optimality of the computed controllers. Unfortunately, in these approaches, the order of the controller is not a priori known, and in particular, it is not possible to compute static controllers with this type of technique. In the absence of these additional assumptions, and in particular when structural constraints are imposed, synthesis via nonlinear optimization appears to be the most general approach to H_∞ synthesis.

The structure of the paper is as follows. In section II we present the H_∞ synthesis problem and give several motivating examples. Section III computes subgradients of the H_∞ norm, which are then applied to closed-loop scenarios in section IV. In section V we start to develop our first-order descent method, which is completed in section VI. Section VI-G discusses practical aspects of the method, and the final section VII presents a number of experiments to validate our approach.

NOTATION

Let $\mathbb{M}_{n,m}$ be the space of $n \times m$ matrices, equipped with the corresponding scalar product $\langle X, Y \rangle = \text{Tr}(X^H Y)$, where X^H is the transconjugate of the matrix X , $\text{Tr } X$ its trace. The space of $m \times m$ Hermitian matrices is denoted $\mathbb{S}_m$. For Hermitian

or symmetric matrices, $X \succ Y$ means that $X - Y$ is positive definite, $X \succeq Y$ that $X - Y$ is positive semi-definite. For ease of notations, we define the following sets of Hermitian matrices: $\mathbb{B}_m := \{X \in \mathbb{S}_m : X \succeq 0, \text{Tr}(X) = 1\}$. Consider q -tuples of Hermitian matrices $(Y_1, \dots, Y_q)$, we define the set $\mathbb{B}_m^q := \{(Y_1, \dots, Y_q) : Y_i \in \mathbb{S}_m, Y_i \succeq 0, \sum_{i=1}^q \text{Tr}(Y_i) = 1\}$. For short, we shall use $\mathbb{B}$ or $\mathbb{B}^q$ when the dimension needs not be specified.

We use the symbol λ_1 to denote the maximum eigenvalue of a symmetric or Hermitian matrix. We shall use general notions from nonsmooth analysis covered by [22]. Notions on ϵ -subdifferentials and ϵ -enlarged subdifferentials for spectral functions and their relationships are discussed at length in [23]–[25]. Unless stated otherwise, the symbol x designates a vector gathering (controller) decision variables and must not be confused with the plant state in Section II. In the notation x_k , the subscript k refers to the iteration index.

II. H_∞ SYNTHESIS

The general setting of the H_∞ synthesis problem is as follows. We consider a linear time-invariant plant described in standard form by the state-space equations:

$$P(s) : \begin{bmatrix} \dot{x} \\ z \\ y \end{bmatrix} = \begin{bmatrix} A & B_1 & B_2 \\ C_1 & D_{11} & D_{12} \\ C_2 & D_{21} & D_{22} \end{bmatrix} \begin{bmatrix} x \\ w \\ u \end{bmatrix}, \quad (1)$$

where $x \in \mathbb{R}^n$ is the state vector, $u \in \mathbb{R}^{m_2}$ the vector of control inputs, $w \in \mathbb{R}^{m_1}$ the vector of exogenous inputs, $y \in \mathbb{R}^{p_2}$ the vector of measurements and $z \in \mathbb{R}^{p_1}$ the controlled or performance vector. Without loss of generality, it is assumed throughout that $D_{22} = 0$.

Let $u = K(s)y$ be a dynamic output feedback control law for the open-loop plant (1), and let $T_{w \rightarrow z}(K)$ denote the closed-loop transfer function of the performance channel mapping w into z . Our aim is to compute $K(s)$ such that the following design requirements are met:

- *Internal stability*: For $w = 0$ the state vector of the closed-loop system (1) and (2) tends to zero as time goes to infinity.
- *Performance*: The H_∞ norm $\|T_{w \rightarrow z}(K)\|_\infty$ is minimized among all stabilizing K .

We assume that the controller K has the following frequency domain representation:

$$K(s) = C_K(sI - A_K)^{-1}B_K + D_K, \quad A_K \in \mathbb{R}^{k \times k}, \quad (2)$$

where k is the order of the controller, and where the case $k = 0$ of a static controller $K(s) = D_K$ is included. Often practical considerations dictate additional challenging structural constraints. For instance it may be desired to design low-order controllers ($0 \leq k \ll n$) or controllers with prescribed-pattern, sparse controllers, decentralized controllers, observed-based controllers, PID control structures, synthesis on a finite set of transfer functions, and much else. Formally, the synthesis problem may then be represented as

$$\begin{aligned} & \text{minimize} && \|T_{w \rightarrow z}(K)\|_\infty \\ & \text{subject to} && K \text{ stabilizes (1)} \\ & && K \in \mathcal{K} \end{aligned} \quad (3)$$

where $K \in \mathcal{K}$ represents a structural constraint on the controller (2) like one of the above.

Without the restriction $K \in \mathcal{K}$, and under standard stabilizability and detectability conditions, it has become customary to synthesize $K(s)$ as follows. After substituting (2) into (1), the H_∞ synthesis problem is transformed into a matrix inequality condition using the bounded real lemma [26]. Then the projection lemma from [1] is used to eliminate the unknown controller data A_K, B_K, C_K, D_K from the cast, leaving an LMI problem, which may be solved by SDP. In a third step the controller state-space representation (2) is recovered.

This scenario changes dramatically as soon as constraints $K \in \mathcal{K}$ are added. Then the problem may no longer be transformed into an LMI or any other convex program, and alternative algorithmic strategies are required. The aim of this paper is to present and analyze one such alternative.

Example 1. Pure stabilization. Often the first important step in controller synthesis (3) is to find a stabilizing controller K . Already at this stage the H_∞ norm plays a prominent role, because of the well-known fact that under stabilizability and detectability, a linear-time invariant system is Lyapunov stable if and only if its H_∞ norm is finite [27]. More specifically, under stabilizability and detectability assumptions, the static control law $u = Ky$ stabilizes the plant

$$G(s) : \begin{cases} \dot{x} &= Ax + B_2u \\ y &= C_2x, \end{cases} \quad (4)$$

if and only if the closed-loop transfer matrix $C_2(sI - (A + B_2KC_2))^{-1}B_2$ has finite H_∞ norm.

In order to construct a static stabilizing controller for an unstable open-loop system (4), the following procedure appears fairly natural. Suppose we are given an initial guess K_0 , which leaves the closed-loop system unstable. Then we pick $a_0 > 0$ such that the a_0 -shifted H_∞ -norm of the closed-loop system is finite:

$$\|C_2(sI - (A + B_2K_0C_2))^{-1}B_2\|_{-a_0, \infty} < +\infty,$$

where the shifted H_∞ norm is given in [28]. The problem of finding a stabilizing controller K may now be addressed by an optimization program

$$\underset{K}{\text{minimize}} \|C_2(sI - (A + B_2KC_2))^{-1}B_2\|_{-a, \infty}, \quad (5)$$

where the shift a is either kept fixed at the initial a_0 , or is gradually decreased after each minimization step to accelerate the procedure. A stabilizing controller K is obviously obtained when the shift reaches $a \leq 0$, but very often this happens already with the initial value a_0 , so that shifting is not even necessary as a rule. Numerical tests for this method will be presented in section VII.

While we stop the optimization (5) as soon as a stabilizing K is reached, it may happen that for a fixed shift a , the method converges to a local minimum K of (5), which fails to stabilize the closed-loop system. This is explained by the fact that (5), just like all the other methods in this paper, are local optimization methods in the sense that they guarantee convergence to a local minimum (or a critical point). If an unsatisfactory local minimum is reached, the only possibility is

to do a restart with a new initial guess K_0 , or switch to another method. Such a local convergence certificate may appear weak at first sight, but experience shows that local methods perform much better than global optimization techniques. Those may have stronger certificates, but run into numerical problems even for small problems. And indeed, our present approach is almost always successful even without restart. A similar comment applies to the pure stabilization method in [29]. $\square$

Notice that pure stabilization is just a special case of the more general H_∞ synthesis problem in section II when we specialize the standard form to

$$P(s) : \begin{bmatrix} \dot{x} \\ z \\ y \end{bmatrix} = \begin{bmatrix} A & B_2 & B_2 \\ C_2 & 0 & 0 \\ C_2 & 0 & 0 \end{bmatrix} \begin{bmatrix} x \\ w \\ u \end{bmatrix}. \quad (6)$$

Example 2. Spectral abscissa. In [30], [31] Burke *et al.* present an alternative approach to computing static stabilizing controllers K . The authors propose to solve the nonsmooth optimization program

$$\underset{K}{\text{minimize}} \quad \alpha(A + B_2 K C_2), \quad (7)$$

where α is the spectral abscissa of a matrix $M \in \mathbb{M}_{n,n}$ defined as

$$\alpha(M) = \max\{\text{Re } \lambda : \lambda \text{ eigenvalue of } M\}.$$

Unfortunately, this function is not even locally Lipschitz, which renders application of existing nonsmooth algorithms impossible. The authors have therefore developed a probabilistic algorithm which allows to treat problems like (7); see [31], [32]. Since optimality tests and approximate subgradients for α are difficult to compute (see [29]), we prefer the use of (5) over (7). Numerical tests for (5) are presented in section VII and show that this method is successful as a rule. Our own numerical tests with program (7) are published in [29]. $\square$

The above scenario covers the case of a static stabilizing controller K , but it is clear that stabilization problems including structural constraints $K \in \mathcal{K}$ can be treated in exactly the same way. Several examples of classes $\mathcal{K}$ will be presented in the sequel.

Example 3. Simultaneous stabilization. Another instance of interest is the simultaneous stabilization problem, which can be cast as minimizing a finite family of closed-loop transfer functions. Formally, given the open-loop plants, G_i , $i = 1, \dots, r$, we consider the problem

$$\underset{K}{\text{minimize}} \quad \max_{i=1, \dots, r} \gamma_i \|C_i(sI - (A_i + B_i K C_i) - a_i I)^{-1} B_i\|_\infty$$

where $\gamma_i > 0$ are appropriate weights, and the a_i are chosen so that the initial guess K renders the i th system stable after shifting by a_i . A lower bound for a_i is therefore the spectral abscissa (7) of the i th system. As before, the shifts are decreased after each H_∞ norm minimization step and a simultaneously stabilizing K is obtained e.g. if $a_i \leq 0$ for all $i = 1, \dots, r$. But even when some $a_i > 0$, the solution may produce a simultaneously stabilizing K . $\square$

Example 4. System reduction. A technique of considerable importance is system reduction. It is used by practitioners whenever an open-loop system $G(s)$ of large order N is difficult to control. Let $G(s)$ denote such a large size open-loop plant, and suppose a decomposition $G(s) = G_{\text{instab}}(s) + G_{\text{stab}}(s)$ into an unstable and a stable part is available. Then we may consider the problem

$$\underset{\tilde{G}_{\text{stab}} \in \mathcal{K}}{\text{minimize}} \quad \|G_{\text{stab}}(s) - \tilde{G}_{\text{stab}}(s)\|, \quad (8)$$

where $\tilde{G}_{\text{stab}}(s)$ ranges over a prespecified class $\mathcal{K}$ of stable system of reduced order, $n \ll N$, and where some norm criterion is used to evaluate the mismatch between nominal and reduced systems. If $\|\cdot\|$ represents the Hankel norm, an explicit expression for $\tilde{G}_{\text{stab}}$ is available [33]. But it may be preferable to use other criteria like the H_∞ norm, a problem which then falls within the class of problems considered in this work. Once a solution $\tilde{G}_{\text{stab}}(s)$ to (8) is obtained, the new system $\tilde{G}(s) = G_{\text{instab}}(s) + \tilde{G}_{\text{stab}}(s)$, while easier to control, may be expected to have characteristics similar to those of the original system. $\square$

III. SUBDIFFERENTIAL OF THE H_∞ NORM

In this section, we start characterizing the subdifferential of the H_∞ norm, and derive expressions for the Clarke subdifferential of several nonconvex composite functions $f(x) = \|\mathcal{G}(x)\|_\infty$, where $\mathcal{G}$ is a smooth operator defined on some $\mathbb{R}^n$ with values in the space of stable matrix transfer functions $\mathbf{H}_\infty$.

Consider the H_∞ -norm of a nonzero transfer matrix function $G(s)$:

$$\|G\|_\infty = \sup_{\omega \in \mathbb{R}} \bar{\sigma}(G(j\omega)),$$

where G is stable and $\bar{\sigma}(X)$ is the maximum singular value of X . Suppose $\|G\|_\infty = \bar{\sigma}(G(j\omega))$ is attained at some frequency ω , where the case $\omega = \infty$ is allowed. Let $G(j\omega) = U\Sigma V^H$ be a singular value decomposition. Pick u the first column of U , v the first column of V , that is, $u = G(j\omega)v/\|G\|_\infty$. Then the linear functional $\phi = \phi_{u,v,\omega}$ defined as

$$\begin{aligned} \phi(H) &= \text{Re}(u^H H(j\omega)v) \\ &= \|G\|_\infty^{-1} \text{Re} \text{Tr } v v^H G(j\omega)^H H(j\omega) \\ &= \|G\|_\infty^{-1} \text{Re} \text{Tr } G(j\omega)^H u u^H H(j\omega) \end{aligned}$$

is continuous on the space $\mathbf{H}_\infty$ of stable transfer functions and is a subgradient of $\|\cdot\|_\infty$ at G [28]. More generally, assume that the columns of Q_u form an orthonormal basis of the eigenspace of $G(j\omega)G(j\omega)^H$ associated with the largest eigenvalue $\lambda_1(G(j\omega)G(j\omega)^H) = \bar{\sigma}(G(j\omega))^2$, and that the columns of Q_v form an orthonormal basis of the eigenspace of $G(j\omega)^H G(j\omega)$ associated with the same eigenvalue. Then for all complex Hermitian matrices $Y_v, Y_u \in \mathbb{B}$,

$$\begin{aligned} \phi(H) &= \|G\|_\infty^{-1} \text{Re} \text{Tr } Q_v Y_v Q_v^H G(j\omega)^H H(j\omega) \\ &= \|G\|_\infty^{-1} \text{Re} \text{Tr } G(j\omega)^H Q_u Y_u Q_u^H H(j\omega) \end{aligned} \quad (9)$$

is a subgradient of $\|\cdot\|_\infty$ at G . Finally, with $G(s)$ rational and assuming that there exist finitely many frequencies $\omega_1, \dots, \omega_p$

where the supremum $\|G\|_\infty = \overline{\sigma}(G(j\omega_\nu))$ is attained, all subgradients of $\|\cdot\|_\infty$ at G are precisely of the form

$$\phi(H) = \|G\|_\infty^{-1} \operatorname{Re} \sum_{\nu=1}^p \operatorname{Tr} G(j\omega_\nu)^H Q_\nu Y_\nu Q_\nu^H H(j\omega_\nu),$$

where the columns of Q_ν form an orthonormal basis of the eigenspace of $G(j\omega_\nu)G(j\omega_\nu)^H$ associated with the leading eigenvalue $\|G\|_\infty^2$, and where $(Y_1, \dots, Y_p) \in \mathbb{B}^p$. See [22, Prop. 2.3.12 and Thm. 2.8.2] and [29] for this.

Suppose now we have a smooth operator $\mathcal{G}$, mapping $\mathbb{R}^n$ onto the space $\mathbf{H}_\infty$ of stable transfer functions G . Then the composite function $f(x) = \|\mathcal{G}(x)\|_\infty$ is Clarke subdifferentiable at x with

$$\partial f(x) = \mathcal{G}'(x)^* [\partial \|\cdot\|_\infty(\mathcal{G}(x))], \quad (10)$$

where $\partial \|\cdot\|_\infty$ is the subdifferential of the H_∞ -norm obtained above, and where $\mathcal{G}'(x)^*$ is the adjoint of $\mathcal{G}'(x)$, mapping the dual of $\mathbf{H}_\infty$ into $\mathbb{R}^n$, where $\mathbb{R}^n$ is identified with its dual here. In the sequel, we will compute this adjoint $\mathcal{G}'(x)^*$ for special classes of closed-loop transfer functions. Suitable chain rules covering this case are for instance given in [22, section 2.3].

IV. CLARKE SUBDIFFERENTIALS IN CLOSED-LOOP

Given a stabilizing controller $K(s)$ and a plant with the usual partition

$$P(s) := \begin{bmatrix} P_{11}(s) & P_{12}(s) \\ P_{21}(s) & P_{22}(s) \end{bmatrix},$$

the closed-loop transfer function is obtained as

$$T_{w \rightarrow z}(K) := P_{11} + P_{12}K(I - P_{22}K)^{-1}P_{21},$$

where the state-space data of P_{11} , P_{12} , P_{21} and P_{22} are given in (1) and the dependence on s is omitted for brevity. Our aim is to compute the subdifferential $\partial f(K)$ of $f := \|\cdot\|_\infty \circ T_{w \rightarrow z}$ at K . We first notice that the derivative $T'_{w \rightarrow z}(K)$ of $T_{w \rightarrow z}$ at K is

$$T'_{w \rightarrow z}(K)\delta K := G_{12}\delta K G_{21},$$

where δK is an element of the same matrix space as K and with the definitions

$$\begin{aligned} \begin{bmatrix} T_{w \rightarrow z}(K, s) & G_{12}(K, s) \\ G_{21}(K, s) & \star \end{bmatrix} &:= \\ \begin{bmatrix} P_{11} + P_{12}K(I - P_{22}K)^{-1}P_{21} & P_{12}(I - KP_{22})^{-1} \\ (I - P_{22}K)^{-1}P_{21} & \star \end{bmatrix} &= \\ \begin{bmatrix} \mathcal{C}(K) \\ C_2 \end{bmatrix} (sI - \mathcal{A}(K))^{-1} [\mathcal{B}(K) \quad B_2] + \begin{bmatrix} \mathcal{D}(K) & D_{12} \\ D_{21} & \star \end{bmatrix} \end{aligned}$$

and the closed-loop state-space data

$$\begin{aligned} \mathcal{A}(K) &:= A + B_2 K C_2, \quad \mathcal{B}(K) := B_1 + B_2 K D_{21}, \\ \mathcal{C}(K) &:= C_1 + D_{12} K C_2, \quad \mathcal{D}(K) := D_{11} + D_{12} K D_{21}. \end{aligned}$$

Note that numerical stability requires that transfer functions $T_{w \rightarrow z}$, G_{12} and G_{21} be computed through state-space realizations.

Now let $\phi = \phi_Y$ be a subgradient of $\|\cdot\|_\infty$ at $T_{w \rightarrow z}(K)$ of the form (9), specified by $Y \in \mathbb{B}$ and with $\|T_{w \rightarrow z}(K)\|_\infty$ attained at frequency ω . According to the chain rule, the subgradients Φ_Y of f at K are of the form $\Phi_Y := T'_{w \rightarrow z}(K)^* \phi_Y \in \mathbb{M}_{m_2, p_2}$, where the adjoint $T'_{w \rightarrow z}(K)^*$ acts on ϕ_Y through

$$\begin{aligned} \langle T'_{w \rightarrow z}(K)^* \phi_Y, \delta K \rangle &= \langle T'_{w \rightarrow z}(K) \delta K, \phi_Y \rangle \\ &= \|T_{w \rightarrow z}(K)\|_\infty^{-1} \operatorname{Re} \operatorname{Tr} (T_{w \rightarrow z}(K, j\omega)^H Q Y Q^H T'_{w \rightarrow z}(K) \delta K(j\omega)) \\ &= \|T_{w \rightarrow z}(K)\|_\infty^{-1} \operatorname{Re} \operatorname{Tr} (T_{w \rightarrow z}(K, j\omega)^H Q Y Q^H G_{12}(K, j\omega) \delta K(j\omega) G_{21}(K, j\omega)) \\ &= \|T_{w \rightarrow z}(K)\|_\infty^{-1} \operatorname{Re} \operatorname{Tr} (G_{21}(K, j\omega) T_{w \rightarrow z}(K, j\omega)^H Q Y Q^H G_{12}(K, j\omega) \delta K(j\omega)). \end{aligned} \quad (11)$$

In consequence, for a static K , the Clarke subdifferential of $f(K) := \|T_{w \rightarrow z}(K)\|_\infty$ at K consists of all subgradients Φ_Y of the form

$$\Phi_Y = \|T_{w \rightarrow z}(K)\|_\infty^{-1} \operatorname{Re} (G_{21}(K, j\omega) T_{w \rightarrow z}(K, j\omega)^H Q Y Q^H G_{12}(K, j\omega))^T, \quad (12)$$

where $Y \in \mathbb{B}$. Recall that Φ_Y is now an element of the same matrix space as K and acts on test vectors δK through $\langle \Phi_Y, \delta K \rangle = \operatorname{Tr}(\Phi_Y^T \delta K)$.

This formula is easily adapted if the H_∞ norm is attained at a finite number of frequencies $\omega_1, \dots, \omega_q$. In this more general situation, subgradients of f at K are of the form

$$\Phi_Y = \|T_{w \rightarrow z}(K)\|_\infty^{-1} \sum_{\nu=1}^q \operatorname{Re} (G_{21}(K, j\omega_\nu) T_{w \rightarrow z}(K, j\omega_\nu)^H Q Y_\nu Q^H G_{12}(K, j\omega_\nu))^T, \quad (13)$$

where $Y = (Y_1, \dots, Y_q) \in \mathbb{B}^q$.

At this stage, it is important to stress that expressions (11), (12) and (13) are general and can accommodate any problem discussed in previous sections. Below we resume and expand this list by considering more examples of practical interest.

Example 5. Dynamic controllers. Assume now that the controller is dynamic as in (2). The subgradient set is again obtained via formula (13) by performing the substitutions:

$$\begin{aligned} K &\rightarrow \begin{bmatrix} A_K & B_K \\ C_K & D_K \end{bmatrix}, \quad A \rightarrow \begin{bmatrix} A & 0 \\ 0 & 0_k \end{bmatrix}, \quad B_1 \rightarrow \begin{bmatrix} B_1 \\ 0 \end{bmatrix}, \\ C_1 &\rightarrow [C_1 \quad 0], \quad B_2 \rightarrow \begin{bmatrix} 0 & B_2 \\ I_k & 0 \end{bmatrix}, \quad C_2 \rightarrow \begin{bmatrix} 0 & I_k \\ C_2 & 0 \end{bmatrix}, \\ D_{12} &\rightarrow [0 \quad D_{12}], \quad D_{21} \rightarrow \begin{bmatrix} 0 \\ D_{21} \end{bmatrix}. \end{aligned} \quad (14)$$

The entire Clarke subdifferential is then described by the set of subgradients in $\mathbb{M}_{k+m_2, k+p_2}$

$$\Phi_Y := \begin{bmatrix} \Phi_{Y, A_K} & \Phi_{Y, B_K} \\ \Phi_{Y, C_K} & \Phi_{Y, D_K} \end{bmatrix},$$

where $Y = (Y_1, \dots, Y_q) \in \mathbb{B}^q$ and Φ_Y is derived through formulas (12) or (13). $\square$

Example 6. Structured controllers. In practice, it is sometimes required that some entries in the controller gain be put to zero, while the others may be freely assigned. This is the case in decentralized control, where the controller must enjoy

a block-diagonal structure. Consider a pattern matrix W with entries $W_{ij} \in \{0, 1\}$, where $W_{ij} = 0$ means that the controller gain $K_{ij} = 0$ must be zero, whereas $W_{\ell k} = 1$ means that $K_{\ell k}$ can be freely assigned. The Clarke subdifferential of $f = \|\cdot\|_\infty \circ \mathcal{F}$ at K is then of the form

$$W \odot \Phi_Y,$$

where $\Phi_Y \in \partial\|\cdot\|_\infty(\mathcal{F}(K))$ is as in (13) and where $\odot$ denotes the entry-wise Hadamard or Schur product [34]. $\square$

Example 7. PID controllers. PID control is one of the most classical approaches in control system design. The controller is generally written as

$$K(s) = K_P + \frac{1}{s}K_I + \frac{s}{1 + \tau s}K_D,$$

where K_P , K_I and K_D are matrix static gains to be computed, and τ is a small positive scalar. Using the general formula (11), the subgradients with respect to $\kappa := [K_P; K_I; K_D]$ at $K(s)$ are obtained through

$$\Phi_{Y,\kappa} = \|T_{w \rightarrow z}(K)\|_\infty^{-1} \sum_{\nu=1}^q \operatorname{Re} (G_{21}(K, j\omega_\nu) T_{w \rightarrow z}(K, j\omega_\nu)^H Q Y_\nu Q^H G_{12}(K, j\omega_\nu) [I \frac{1}{j\omega_\nu} I \frac{j\omega_\nu}{1 + \tau j\omega_\nu} I])^T,$$

where as before $Y = (Y_1, \dots, Y_q) \in \mathbb{B}^q$.

The above approach could be generalized by making τ an additional design parameter. Then a constraint $\tau \geq 0$ should be added. Notice also that the above formula readily extends to arbitrary basis functions $\{Q_j(s)\}_{j=1, \dots, r}$

$$K(s) := \sum_{j=1}^r K_j Q_j(s),$$

where the K_j 's are the design variables.

Example 8. Matrix fraction representations. An alternative representation of controllers is via matrix fraction descriptions. For instance, the left matrix fraction representation is given as

$$K(s) = N(s)D(s)^{-1} = (\mathcal{N}J_N(s))(\mathcal{D}J_D(s))^{-1},$$

with

$$J_N(s) := [I \ sI \ \dots \ s^n I]^T, \quad J_D(s) := [I \ sI \ \dots \ s^d I]^T$$

and

$$\mathcal{N} := [N_0 \ \dots \ N_n], \quad \mathcal{D} := [D_0 \ \dots \ D_n]$$

Now, $\mathcal{N}$ and $\mathcal{D}$ are the design variables. As before, it is immediate to show that partial subgradients with respect to $\mathcal{N}$ and $\mathcal{D}$ are given as

$$\begin{aligned} \Phi_{Y,\mathcal{N}} &= \|T_{w \rightarrow z}(K)\|_\infty^{-1} \sum_{\nu=1}^q \operatorname{Re} (J_N(j\omega_\nu) (D(j\omega_\nu)J_D(j\omega_\nu))^{-1} G_{21}(K, j\omega_\nu) T_{w \rightarrow z}(K, j\omega_\nu)^H Q Y_\nu Q^H G_{12}(K, j\omega_\nu))^T \\ \Phi_{Y,\mathcal{D}} &= -\|T_{w \rightarrow z}(K)\|_\infty^{-1} \sum_{\nu=1}^q \operatorname{Re} (J_D(j\omega_\nu) (D(j\omega_\nu)J_D(j\omega_\nu))^{-1} G_{21}(K, j\omega_\nu) T_{w \rightarrow z}(K, j\omega_\nu)^H Q Y_\nu Q^H G_{12}(K, j\omega_\nu) K(j\omega_\nu))^T, \end{aligned}$$

respectively. $\square$

Example 9. Multiple performance channels. Practical specifications often impose that several closed-loop channels $i =$

$1, \dots, r$ be minimized simultaneously. One way to address multi-objective optimization of this type is to solve a program of the form

$$\text{minimize}_K \max \{ \gamma_i \|T_{w \rightarrow z}^i(K)\|_\infty : i = 1, \dots, r \},$$

where $T_{w \rightarrow z}^i$ is the i th performance specification to be optimized. Since the maximum of a finite number of maximum eigenvalue functions is itself a maximum eigenvalue function of a block diagonal operator $T = \operatorname{diag}(T_{w \rightarrow z}^1, \dots, T_{w \rightarrow z}^r)$, the Clarke subgradients could be obtained directly from (13). When the usual max formula is used, the result is the same, i.e., subgradients are of the form

$$\phi_{(Y,\tau)} = \sum_{i \in I(K)} \tau_i \gamma_i \phi_{Y_i},$$

where $I(K)$ are the $i = 1, \dots, r$ which are active at K , $\tau_i \geq 0$, $\sum_{i \in I(K)} \tau_i = 1$ and $\phi_{Y_i} \in (T_{w \rightarrow z}^i)^* \partial\|\cdot\|_\infty(K)$ as specified in (13). $\square$

Before going further, it is worth mentioning that our methodology carries over to a wide range of controller structures of practical interest. This is in particular the case when the structural constraint is of the form $\mathcal{K} = \{K : K = \mathcal{S}(\ell), \ell \in \mathcal{L}\}$, where $(\mathcal{S}, \mathcal{L})$ is a suitable differentiable parametrization of the class $\mathcal{K}$. This includes for instance observed-based controllers, feed-forward compensators, controllers defined through Youla parameterizations and much else.

V. STEEPEST DESCENT METHOD

Nonsmooth techniques have been used before in algorithms for controller synthesis. For instance, E. Polak and co-workers have proposed a variety of techniques suited for eigenvalue or singular-value optimization and for extensions to the semi-infinite case, covering in particular the H_∞ -norm (see [35], [36] and the citations given there). Another reference is [28], where the authors exploit the Youla parameterization via convex nondifferentiable analysis to derive the cutting plane and ellipsoid algorithms.

Let us consider the problem of minimizing $f(x) = \|\mathcal{G}(x)\|_\infty$, where x regroups the controller data, referred to as K in the previous section, and where $\mathcal{G}$ maps $\mathbb{R}^n$ smoothly into a space $\mathbf{H}_\infty$ of stable transfer functions. We write $\mathcal{G}(x, s)$ or $\mathcal{G}(x, j\omega)$ when the complex argument of $\mathcal{G}(x) \in \mathbf{H}_\infty$ needs to be specified.

A necessary condition for optimality is $0 \in \partial f(x) = \mathcal{G}'(x)^* \partial\|\cdot\|_\infty(\mathcal{G}(x))$. It is therefore reasonable to consider the program

$$\begin{aligned} d &= -\frac{g}{\|g\|}, \\ g &= \operatorname{argmin}\{\|\phi_Y\| : Y = (Y_1, \dots, Y_q) \in \mathbb{B}^q\} \end{aligned} \quad (15)$$

which either shows $0 \in \partial f(x)$, or produces the direction d of steepest descent at x if $0 \notin \partial f(x)$, and where the ϕ_Y are as in (13). If we vectorize $y = \operatorname{vec}(Y)$, $Y = (Y_1, \dots, Y_q)$, then we may represent ϕ_Y by a matrix vector product, $\phi_Y = \Phi y$,

with a suitable matrix Φ . Program (15) is then equivalent to the following SDP:

$$\begin{aligned} & \text{minimize} && t \\ & \text{subject to} && \begin{bmatrix} t & y^T \Phi^T \\ \Phi y & tI \end{bmatrix} \succeq 0 \\ & && Y_i \succeq 0, i = 1, \dots, q \\ & && e^T y = 1 \end{aligned} \quad (16)$$

where $e^T y = 1$ encodes the constraint $\sum_i \text{Tr}(Y_i) = 1$. The direction d of steepest descent at x is then obtained as $d = -\Phi y / \|\Phi y\|$, where (t, y) is solution of (16) with $y \neq 0$. This suggests the following:

Steepest descent method for the H_∞ -norm

- 1) If $0 \in \partial f(x)$ stop. Otherwise:
- 2) Solve (16) and compute the direction d of steepest descent at x .
- 3) Perform a line search and find a descent step $x^+ = x + t d$.
- 4) Replace x by x^+ and go back to step 1.

The drawback of this approach is that it may fail to converge due to the nonsmoothness of f . We believe that a descent method should at least give the weak convergence certificate that accumulation points of the sequence of iterates are critical. This is not guaranteed by the above scheme. The reason is that the steepest descent direction at x does not depend continuously on x . This is why modifications of the steepest descent scheme are discussed in the next section.

Remark. Spectral abscissa versus H_∞ norm. Notice that the tangent program (16) is convenient because it leads to relatively small size SDPs. Indeed, the matrices Y_ν , $\nu = 1, \dots, q$, are of the size of the multiplicity of $\lambda_1(\mathcal{F}(x, j\omega_\nu))$, and our experiments indicate that $\dim(y)$ in (16) rarely exceeds 30. The situation is very different for the spectral abscissa (7). In [29] we have derived a tangent program for (7). The difficulty is that Lyapunov variables re-enter the scene. Indeed, x^* is a local minimum of the composite function $\alpha \circ \mathcal{F}$ like in (7) with value $\alpha(\mathcal{F}(x^*)) = t^*$ if and only if (x^*, t^*, X^*) is a local minimum of the optimization program:

$$\begin{aligned} & \text{minimize} && t \\ (P) \quad & \text{subject to} && X \succeq \sigma I, X \preceq I \\ & && \mathcal{F}(x)^T X + X \mathcal{F}(x) - 2tX \preceq 0 \end{aligned}$$

for some small fixed $0 < \sigma \ll 1$. An optimality test for $\alpha \circ \mathcal{F}$ is therefore derived from an optimality test for program (P), as shown in [29]. This leads to a SDP with unknown variable of $\dim(x, X)$, which may be prohibitively large. This is one of the reasons why our present approach privileges the use of the H_∞ norm (5) over (7). $\square$

VI. FIRST-ORDER DESCENT METHOD

In this section we devise a first-order algorithm for composite functions of the H_∞ norm. Along with $\mathcal{G} : \mathbb{R}^n \rightarrow \mathbf{H}_\infty$ we consider the symmetrized operators $\mathcal{F}(x, s) = \mathcal{G}(x, s)\mathcal{G}(x, s)^H$ respectively $\mathcal{F}(x, s) = \mathcal{G}(x, s)^H \mathcal{G}(x, s)$. We represent $f(x)$ as

$$f(x) = \sup_{\omega \in \mathbb{R}} f(x, \omega), \quad f(x, \omega) = \lambda_1(\mathcal{F}(x, j\omega)), \quad (17)$$

and solve the optimization program

$$\min_{x \in \mathbb{R}^n} f(x), \quad f(x) = \|\mathcal{G}(x)\|_\infty^2 = \sup_{\omega \in \mathbb{R}} \lambda_1(\mathcal{F}(x, j\omega))$$

Notice that for fixed $\omega \in \mathbb{R}$, $x \mapsto \mathcal{F}(x, j\omega)$ is a smooth operator into the space of Hermitian $m \times m$ matrices, $\mathbb{S}_m$, while $\lambda_1 : \mathbb{S}_m \rightarrow \mathbb{R}$ is the maximum eigenvalue function. Similar techniques could be applied to broader classes with a structure like (17). Deriving the method will require three steps, which we regroup into subsections. We start with the important special case $f = \lambda_1 \circ \mathcal{F}$, where $\mathcal{F}$ maps $\mathbb{R}^n$ smoothly into $\mathbb{S}_m$.

A. Preparation

The function (17) is subject to two sources of nonsmoothness. The nonsmooth character of the maximum eigenvalue function, and the nonsmoothness introduced by the operator $\sup$, which in the case of $\|\cdot\|_\infty$ is even infinite. Each individual function $f(\cdot, \omega)$ will be analytic at x if the multiplicity of $\lambda_1(\mathcal{F}(x, j\omega))$ is one, but nonsmoothness needs to be taken into account as soon as eigenvalues coalesce.

From a practical point of view it is reasonable to make the following additional hypothesis.

- (H) The maximum $f(x)$ is always attained on a finite set of frequencies. This set is denoted by $\Omega(x)$ and may contain $\omega = \pm\infty$.

Assumption (H) is for instance satisfied when the multiplicity of λ_1 is 1, as $\omega \mapsto f(x, \omega)$ is then analytic in typical control applications.

Let us introduce some more notation. For $\Omega \subset \mathbb{R} \cup \{-\infty, +\infty\}$ we define $f_\Omega \leq f$ as

$$f_\Omega(x) = \max_{\omega \in \Omega} f(x, \omega).$$

Notice that $f_\Omega(x) = f(x)$ as soon as $\Omega(x) \subset \Omega$. Next recall the definition of the ϵ -subdifferential of the maximum eigenvalue function [23]

$$\partial_\epsilon \lambda_1(X) = \{Z \in \mathbb{S}_m : \text{Tr}(ZX) \geq \lambda_1(X) - \epsilon\}$$

which is an important analytical tool in nonsmooth analysis. Since $\partial_\epsilon \lambda_1(X)$ is difficult to compute, we follow Cullum et al. [23] and Oustry [25] and introduce a modification $\delta_\epsilon \lambda_1(X)$ of $\partial_\epsilon \lambda_1(X)$, called the ϵ -enlarged subdifferential for the maximum eigenvalue function. For $\epsilon > 0$ and $X \in \mathbb{S}_m$ let $r(\epsilon, X)$ the index such that

$$\lambda_1 \geq \dots \geq \lambda_{r(\epsilon, X)} \geq \lambda_1 - \epsilon > \lambda_{r(\epsilon, X)+1} \geq \dots \geq \lambda_m.$$

The index $r(\epsilon, X)$ is also called the ϵ -multiplicity of $\lambda_1(X)$. Let Q_ϵ be a $r(\epsilon, X) \times m$ -matrix whose columns form an orthonormal basis of the invariant subspace of X associated with the first $r(\epsilon, X)$ eigenvalues. Then we define

$$\delta_\epsilon \lambda_1(X) = \{Q_\epsilon Y Q_\epsilon^H : Y \in \mathbb{B}_{r(\epsilon, X)}\}.$$

By construction $\partial \lambda_1(X) \subset \delta_\epsilon \lambda_1(X) \subset \partial_\epsilon \lambda_1(X)$, so $\delta_\epsilon \lambda_1(X)$ is an enlargement of $\partial \lambda_1(X)$ and an inner approximation of the ϵ -subdifferential (see also [25]). The gap associated with the choice ϵ is $\Delta(\epsilon, X) = \lambda_{r(\epsilon, X)}(X) - \lambda_{r(\epsilon, X)+1}(X) > 0$. If

r is the multiplicity of $\lambda_1(X)$, then choosing ϵ small enough gives $r(\epsilon, X) = r$. In this case we have $\delta_\epsilon \lambda_1(X) = \partial \lambda_1(X)$.

The following is an important step toward the analysis of (17). Consider a differentiable mapping $\mathcal{F} : \mathbb{R}^n \rightarrow \mathbb{S}_m$. We extend $\partial_\epsilon \lambda_1(X)$ and the enlarged subdifferential $\delta_\epsilon \lambda_1(X)$ to the composite function $f = \lambda_1 \circ \mathcal{F}$ by setting

$$\partial_\epsilon f(x) = \mathcal{F}'(x)^* [\partial_\epsilon \lambda_1(X)], \quad X = \mathcal{F}(x),$$

and similarly

$$\delta_\epsilon f(x) = \mathcal{F}'(x)^* [\delta_\epsilon \lambda_1(X)], \quad X = \mathcal{F}(x).$$

Here $\mathcal{F}'(x)^*$ is the adjoint of the linear operator $\mathcal{F}'(x)$. Finally, going one step further, this approach allows us to consider $\partial_\epsilon f(x, \omega)$ and $\delta_\epsilon f(x, \omega)$, which are applied to the variable x .

B. Descent Step Generator

In this section we discuss a very simple mechanism which generates descent steps in such a way that the following weak form of convergence can be guaranteed: every accumulation point of the sequence of iterates is a stationary point.

Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a locally Lipschitz function, and let $\partial f(x)$ denote its Clarke subdifferential [22]. Suppose we can exhibit a mechanism $s : \mathbb{R}^n \rightarrow \mathbb{R}^n$, the *descent step generator*, such that the following rules are satisfied:

- (i) Whenever $0 \notin \partial f(x)$, then $f(s(x)) < f(x)$.
- (ii) When $0 \notin \partial f(x)$, then there exists a neighborhood $B(x, \epsilon)$ of x and some $\delta > 0$ such that $f(s(x')) \leq f(x') - \delta$ for every $x' \in B(x, \epsilon)$.

While (i) simply means that $s(x)$ is a descent step away from x , we can interpret (ii) as some weak form of continuity of $s(\cdot)$. Indeed, when the mapping $s(\cdot)$ describing descent is continuous, then (i) implies (ii) without further work. Axiom (ii) is weaker than asking $s(\cdot)$ to be continuous. Clearly, (ii) always implies (i).

Example. In order to understand the idea behind axiom (ii), consider a C^1 -function f and let $s(x)$ denote the steepest descent step at x , obtained by an Armijo line search. If we define formally $s(x) = x - t(x)f'(x)$, where $t(x)$ is the smallest step satisfying an Armijo rule, then $s(x)$ turns out continuous, but it is clear that in practice we would accept *any* step $t(x)$ satisfying the Armijo condition, without insisting on a continuous dependence of $t(x)$ on x . In that case $s(x)$ will not be continuous, but property (ii) will still hold true. $\square$

Clearly the situation we have in mind is when f is nonsmooth, so that a steepest descent step in tandem with Armijo search would typically fail even when $t(x)$ was continuous. This is explained by the fact that under nonsmoothness, defining $s(\cdot)$ along the lines above would miss axiom (ii), because the steepest descent direction $-f'(x)$ behaves very discontinuously. Indeed, examples where this happens are easily produced.

Proposition 1: Suppose the descent step generator $s(\cdot)$ for f satisfies axioms (i) and (ii). Let x_k be a sequence of iterates defined as $x_{k+1} = s(x_k)$. Then every accumulation point $\bar{x}$ of x_k is a critical point, that is, satisfies $0 \in \partial f(\bar{x})$.

Proof: Let $\mathcal{N} \subset \mathbb{N}$ be an infinite sequence such that $x_k \rightarrow \bar{x}$, $k \in \mathcal{N}$. Then by monotonicity, $f(x_{k+1}) \rightarrow f(\bar{x})$. That means

$$f(x_{k+1}) - f(x_k) \rightarrow 0 \quad (k \in \mathcal{N}).$$

Now use axiom (ii) at the limit point $\bar{x}$. There exist $\epsilon, \delta > 0$ such that $f(s(x)) - f(x) \leq -\delta < 0$ for all $x \in B(\bar{x}, \epsilon)$. Since $x_k \in B(\bar{x}, \epsilon)$ for $k \in \mathcal{N}$ large enough, and since $x_{k+1} = s(x_k)$, we should have $f(x_{k+1}) - f(x_k) \leq -\delta$ for $k \in \mathcal{N}$ large enough, a contradiction. $\blacksquare$

C. Eigenvalue Optimization

How can we define a descent step generator $s(\cdot)$ with properties (i) and (ii) for a maximum eigenvalue function $f(x) = \lambda_1(\mathcal{F}(x))$? Suppose we are at a point x where $0 \notin \partial f(x)$. Let the eigenvalues of $X = \mathcal{F}(x) \in \mathbb{S}_m$ be arranged into groups:

$$\lambda_1 = \dots = \lambda_{k_2-1} > \lambda_{k_2} = \dots = \lambda_{k_3-1} > \lambda_{k_3} = \dots$$

where $k_1 = 1$ and where k_i are the group leaders. Consequently, eigenvalue gaps occur between $k_i - 1$ and k_i . Let Q_1 be an orthonormal bases of the eigenspace associated with the first block $\lambda_1(X), \dots, \lambda_{k_2-1}(X)$, Q_2 an orthonormal basis containing Q_1 associated with the first two blocks $\lambda_1(X), \dots, \lambda_{k_3-1}(X)$, and so on. At $X = \mathcal{F}(x) \in \mathbb{S}_m$ we compute the quantities

$$\Delta_i(X) = \lambda_{k_{i+1}-1}(X) - \lambda_{k_{i+1}}(X) > 0,$$

$$d_i(x) = \min\{\|\mathcal{F}'(x)^* Q_i Y_i Q_i^H\| : Y_i \in \mathbb{B}_{k_{i+1}-1}\}$$

and keep those $i = 1, \dots, r$ where $d_i(x) > 0$. Notice that $d_i(x) = \text{dist}(0, \delta_{\epsilon_i(x)} f(x))$, where $\epsilon_i(x) > 0$ cuts into the i th gap, that is $\lambda_1(X) - \epsilon_i(x) \in [\lambda_{k_{i+1}-1}(X), \lambda_{k_{i+1}}(X)]$. Put differently, the $\epsilon_i(x)$ -multiplicity of $\lambda_1(X)$ is $k_{i+1} - 1$. Moreover, $d_1(x) \geq d_2(x) \geq \dots \geq d_r(x)$ and $d_{r+1}(x) = d_{r+2}(x) = \dots = 0$ eventually. We compute the quantity

$$M(x) = \max_{i=1 \dots r} \Delta_i(X) d_i(x)^2 \quad (18)$$

Now we use the following

Lemma 1: Let $0 \notin \partial f(x)$ and $R > 0$. Let k_i be the leader of a group of eigenvalues of $X = \mathcal{F}(x) \in \mathbb{S}_m$ such that $d_i(x) = \text{dist}(0, \delta_{\epsilon_i(x)} f(x)) > 0$. Let $h_i(x)$ be the direction of steepest $\epsilon_i(x)$ enlarged descent, that is $h_i(x) = -\frac{g_i}{\|g_i\|}$ with

$$g_i = \text{argmin}\{\|g\| : g = \mathcal{F}'(x)^* Q_i Y Q_i^H, Y \in \mathbb{B}_{k_{i+1}-1}\},$$

where the columns of Q_i are an orthonormal basis of the invariant subspace of X associated with the eigenvalues up to $k_{i+1} - 1$. Then there exists a descent step $s_i(x)$ away from x in direction $h_i(x)$, which decreases the value of f by at least

$$f(s_i(x)) - f(x) \leq -\kappa(x) \Delta_i(X) d_i(x)^2 < 0.$$

Here $\kappa(x) > 0$ depends on $\sup\{\|\mathcal{F}'(x')\| : \|x - x'\| \leq R\}$ and in particular continuously on x . The line search required to compute this step is finite.

Proof: In the case of an affine operator $\mathcal{A} : \mathbb{R}^n \rightarrow \mathbb{S}_m$, Oustry [25, Theorem 5] shows that $f = \lambda_1 \circ \mathcal{A}$ may be decreased by the quantity

$$f(x + t(x)h_i(x)) - f(x) \leq -\frac{1}{4\|\mathcal{A}_*\|} \Delta_{\epsilon_i(x)}(X) d_i(x)^2$$

where $\mathcal{A}_*$ is the linear part of $\mathcal{A}$. Here $\kappa = 1/4\|\mathcal{A}_*\|$ is even independent of x . Moreover, $t(x) > 0$ is computed by a line search which terminates after a finite number of steps, see [25, sect. 3.2].

In [37], [38], this result is generalized to nonconvex maximum eigenvalue functions $f = \lambda_1 \circ \mathcal{F}$, with a constant $\kappa(x)$ now depending on the Lipschitz constant of $\mathcal{F}$ on a bounded region around x , like for instance $B(x, R)$ with some fixed $R > 0$. In the nonconvex case, the line search procedure locating $t(x)$ is more complicated than for $f = \lambda_1 \circ \mathcal{A}$, but finite termination is still guaranteed (cf. [37, sect. 3.7]). The constant $\kappa(x)$ may be computed via the formulae (15), (19) and (20) of that reference. ■

Following the lead of [37], it is now clear how to obtain our step $s(x)$. Choose i so that the maximum (18) is attained and take $s(x) = s_i(x)$. Then by Lemma 1, $s(x)$ gives a guaranteed decrease of

$$f(s(x)) - f(x) \leq -\kappa(x) \Delta_{\epsilon_i(x)}(X) d_i(x)^2 = -\kappa(x) M(x),$$

where the constant $\kappa(x) > 0$ depends continuously on x as argued above. What remains to be checked is

Lemma 2: This choice of $s(x)$ and $M(x)$ guarantees property (ii).

Proof: Consider a sequence $x_k \rightarrow x$. We need to compare $M(x_k)$ to $M(x)$. Now observe that due to the continuity of the eigenvalue functions $\lambda_j(X)$, every eigenvalue gap at X is also an eigenvalue gap at X_k as soon as X_k is sufficiently close to X . On the other hand, X_k may (and will) have many more gaps than X . But notice that the maximum (18) is over all eigenvalue gaps. So each gap in $M(x)$ will occur in the computation of $M(x_k)$. More precisely, it will be approximated by some of the gaps considered in $M(x_k)$. Put differently, for the i th eigenvalue gap of X we have $\Delta_{i_k}(X_k) d_{i_k}(x_k)^2 \rightarrow \Delta_i(X) d_i(x)^2$, where i_k is the index of the eigenvalue gap of X_k corresponding to the i th gap of X . Here $\Delta_{i_k}(X_k) \rightarrow \Delta_i(X)$ and $d_{i_k}(x_k) \rightarrow d_i(x)$ rely on the fact that gaps at X_k remain gaps at X . Since $d_{i_k}(x_k)$ calls for all the $Q_1, Q_2, \dots$ up to the i_k th gap at X_k , and since these converge to the corresponding basis Q regrouping the gap i , the result follows.

That means $\limsup M(x_k) \geq M(x)$, so we can assume $M(x_k) \geq \frac{1}{2}M(x)$ from some index k on. This in turn implies

$$f(s(x_k)) - f(x_k) \leq -\kappa(x_k) M(x_k) \leq -\frac{1}{4}\kappa(x) M(x)$$

as soon as $\kappa(x_k)$ is close enough to $\kappa(x)$, proving property (ii). ■

The method outlined in this section gives a convergence certificate because *all* eigenvalue gaps are included in the computation of (18). This may seem inconvenient for very large size matrices $\mathcal{F}(x)$. If we decide to truncate and consider only some of the eigenvalue gaps among the largest eigenvalues, the

theoretical convergence properties of the method are weaker, even though convergence may still be guaranteed e.g. when f is convex (see [37]). Notice however that even for large m the quantity $M(x)$ in (18) may be computed fairly reliably by considering $\Delta_{\epsilon_i(x)}(X) d_i(x)^2$ for some of the first i only. As $d_i(x) \rightarrow 0$ rather quickly, the higher $\Delta_{\epsilon_i(x)}(X) d_i(x)^2$ as a rule do not contribute to the computation of $M(x)$. Also notice that since the sequence $d_i(x)$ is monotone, computing $d_i(x)$ may often be avoided, for instance when $\Delta_{\epsilon_{i-1}(x)}(X) \geq \Delta_{\epsilon_i(x)}(X)$ or $\Delta_{\epsilon_i(x)}(X) \leq \Delta_{\epsilon_{i-\nu}(x)}(X)$ when the current best value is located at index $i - \nu$.

D. Descent by a local model

In this section we present an alternative way to obtain a descent step generator $s(\cdot)$ for the maximum eigenvalue function $f(x) = \lambda_1(\mathcal{F}(x))$. We start by constructing an intermediate function $\theta(x)$, which serves as an optimality test, and secondly, will allow us to quantify decrease. In this section our method follows the line of [36, Thm. 2.1.6].

Let $X = \mathcal{F}(x) \in \mathbb{S}_m$. Let $\mu_1(X) > \mu_2(X) > \dots > \mu_r(X)$ be the eigenvalues of X without repetitions. That means $\mu_i = \lambda_{k_i}$ in our old terminology, where we agree that there are $r \leq m$ distinct eigenvalues. For some fixed $\sigma > 0$ we define the criticality measure

$$\theta(x) = \inf_{h \in \mathbb{R}^n} \sup_{i=1, \dots, r} \sup_{Y_i \in \mathbb{B}} \left\{ -f(x) + \mu_i(X) + \text{Tr}(Y_i Q_i^H [\mathcal{F}'(x)h] Q_i) + \frac{1}{2} \sigma \|h\|^2 \right\} \quad (19)$$

Here Q_i has $k_{i+1} - 1$ columns which form an orthonormal basis of the invariant subspace of X associated with the first $k_{i+1} - 1$ eigenvalues. It is immediately clear that $\theta(x) \leq 0$, because putting $h = 0$ gives the upper bound $\theta(x) \leq \sup_{i=1, \dots, r} -f(x) + \mu_i(X) = -f(x) + \mu_1(X) = 0$.

Lemma 3: We have $\theta(x) = 0$ if and only if $0 \in \partial f(x)$.

Proof: The easiest way to see this is to swap max and min in (19). This requires that we first replace the inner double supremum in (19) by a double supremum over the convex hull of $Y_i \succeq 0$, $\text{Tr}(Y_i) = 1$, $i = 1, \dots, r$, a manoeuvre which does not change the value $\theta(x)$. Then we use Fenchel duality to interchange the inner (double) supremum and the outer infimum, which goes again without changing the value. The now inner infimum is unconstrained and may be computed explicitly. For fixed Y_i and convex coefficients τ it is attained at $h(x) = -\sigma^{-1} \sum_{i=1}^r \tau_i \mathcal{F}'(x)^* Q_i Y_i Q_i^H$. Substituting this back into (19) leaves the dual expression

$$\theta(x) = \sup_{\tau_i \geq 0, \sum_i \tau_i = 1} \sup_{Y_i \in \mathbb{B}} \left\{ -f(x) + \sum_{i=1}^r \tau_i \mu_i(X) - \frac{1}{2\sigma} \left\| \sum_{i=1}^r \tau_i \mathcal{F}'(x)^* Q_i Y_i Q_i^H \right\|^2 \right\} \quad (20)$$

which the reader recognizes as a semidefinite program. Since $\mu_i(X) < f(x)$ for $i \geq 2$, equality $\theta(x) = 0$ is only possible when $\tau_2 = \dots = \tau_r = 0$ and hence

$$0 = \theta(x) = \sup_{Y_1 \in \mathbb{B}} -\frac{1}{2\sigma} \left\| \mathcal{F}'(x)^* Q_1 Y_1 Q_1^H \right\|^2$$

But the quantity on the right hand side is only zero when $\mathcal{F}'(x)^* Q_1 Y_1 Q_1^H = 0$, i.e., when $0 \in \partial f(x)$. The latter follows readily from the representation

$$\partial f(x) = \{\mathcal{F}'(x)^* Q_1 Y_1 Q_1^H : Y_1 \in \mathbb{B}\}$$

of the subdifferential $\partial f(x)$. ■

As a byproduct of the proof via duality we have the following

Corollary 1: The infimum in (19) is attained at

$$h(x) = -\frac{1}{\sigma} \sum_{i=1}^r \tau_i \mathcal{F}'(x)^* Q_i Y_i Q_i^H, \quad (21)$$

where (τ, Y) is the solution of the dual program (20). As soon as $\theta(x) < 0$, $h(x)$ is a direction of descent of f at x .

In order to construct our descent step generator, we need to establish two additional properties of θ . Firstly, we need to show that decrease at x may be quantified with the help of θ . Secondly, we have to establish continuity of θ .

Lemma 4: The function θ is continuous.

Proof: Notice that by the dual formula (20), $\theta(x)$ is the supremum of an infinite family of functions of the form $-f(x) + \sum_{i=1}^r \tau_i \mu_i(X) + c(x) = -f(x) + \sum_{j=1}^m \sigma_j \lambda_j(X) + c(x)$ for $\sigma_j = \tau_i / (k_{i+1} - k_i)$, where $k_i \leq j < k_{i+1}$, and where $c(x) = -1/2\sigma \|\sum_i \tau_i \mathcal{F}'(x)^* Q_i Y_i Q_i^H\|^2$ depends continuously on x . Notice that $\mu_i(X) = (k_{i+1} - k_i)^{-1} \sum_{j=k_i}^{k_{i+1}-1} \lambda_j(X)$. This shows that $\theta(x)$ is the supremum of a family of continuous functions, indexed by (τ, Y) . θ is therefore lower semi-continuous. It remains to prove that θ is also upper semi-continuous.

Let $x_j \rightarrow x$ such that $\theta(x_j) \rightarrow \theta$. We have to show $\theta \leq \theta(x)$. We use the representation (20). Let $\epsilon_j > 0$, $\epsilon_j \rightarrow 0$ and choose τ_i^j and Y_i^j such that

$$\theta(x_j) \leq -f(x_j) + \sum_{i=1}^r \tau_i^j \mu_i(X_j) - \frac{1}{2\sigma} \left\| \sum_{i=1}^r \tau_i^j g_i^j \right\|^2 + \epsilon_j,$$

where $g_i^j = \mathcal{F}'(x_j)^* Q_i^j Y_i^j Q_i^j{}^H$. Passing to a subsequence if necessary, we may assume that each X_j has exactly r eigenvalue gaps, which remain in the same places k_i , $i = 1, \dots, r$. That is $\mu_i(X_j) = \lambda_{k_i}(X_j)$. Passing to yet another subsequence, we may assume $\tau_i^j \rightarrow \tau_i$, $Y_i^j \rightarrow Y_i$ and $Q_i^j \rightarrow Q_i$, where the limiting elements are all of the same types and dimensions as the elements at stage j .

However, $\lambda_{k_i}(X)$ are no longer the distinct eigenvalues of X , because some of the distinct $\mu_i(X_j) = \lambda_{k_i}(X_j)$ may coalesce in the limit $j \rightarrow \infty$. Suppose for instance that $\lim_{j \rightarrow \infty} \mu_i(X_j) = \dots = \lim_{j \rightarrow \infty} \mu_{i+t}(X_j)$, so that the blocks $i, i+1, \dots, i+t$ coalesce in the limit, forming a new larger eigenvalue block of X :

$$\lambda_{k_{i-1}}(X) > \lambda_{k_i}(X) = \dots = \lambda_{k_{i+t-1}}(X) > \lambda_{k_{i+t}}(X),$$

which is represented by a certain $\mu_\nu(X)$. Suppose there are N block leaders at X . For each of these $\mu_\nu(X)$ we define

$$\sigma_\nu = \lim_{j \rightarrow \infty} \sum_{s=i}^{i+t} \tau_s^j,$$

then

$$\lim_{j \rightarrow \infty} \sum_{i=1}^r \tau_i^j \mu_i(X_j) = \sum_{\nu=1}^N \sigma_\nu \mu_\nu(X),$$

where $\sum_{\nu=1}^N \sigma_\nu = 1$. This represents the limiting linear term in (20) in a new form suited for the dual representation of $\theta(x)$.

We next have to treat the norm square term arising in (20) in much the same way. Notice first that the Q_i^j , $i = 1, \dots, r$ together form a nested sequence of basis vectors adding up to an orthonormal basis of eigenvectors of X_j . Passing to the limit $j \rightarrow \infty$ gives an orthonormal basis of eigenvectors of X . We regroup it according to the eigenvalue gaps of X and rename the corresponding parts $P_1 \subset P_2 \subset \dots \subset P_N$. All that remains to do is to re-write the limit $\lim_{j \rightarrow \infty} \sum_{i=1}^r \tau_i^j Q_i^j Y_i^j Q_i^j{}^H$ in the form $\sum_{\nu=1}^N \sigma_\nu P_\nu Z_\nu P_\nu^H$ for certain $Z_\nu \in \mathbb{B}$. This is done by writing

$$Q_i Y_i Q_i^H = Q_j \begin{bmatrix} Y_i & 0 \\ 0 & 0 \end{bmatrix} Q_j^H = P_\nu \begin{bmatrix} Y_i & 0 \\ 0 & 0 \end{bmatrix} P_\nu^H$$

whenever $i < j$ have to be regrouped in the same P_ν , and where j is the last among the old indices subsumed into the new index ν , so that $Q_j = P_\nu$. Then

$$Z_\nu = \sum_{s=i}^{i+t} \frac{\tau_i}{\sigma_\nu} \begin{bmatrix} Y_i & 0 \\ 0 & 0 \end{bmatrix}$$

is as required, because $\sum_{s=i}^{i+t} \tau_s = \sigma_\nu$, hence $\text{Tr}(Z_\nu) = 1$, while $Z_\nu \succeq 0$ is clear.

The argument shows that $\theta = -f(x) + \sum_{\nu=1}^N \sigma_\nu \mu_\nu(X) - \frac{1}{2\sigma} \left\| \sum_{\nu=1}^N \sigma_\nu \mathcal{F}'(x)^* P_\nu Z_\nu P_\nu^H \right\|^2$, which is to say that θ is now of the form required to be admitted to the supremum (20) defining $\theta(x)$. In other words, $\theta \leq \theta(x)$, and this is what we had to prove. ■

Notice that [36, Thm. 2.1.6 (e)] is obtained as a special case of Lemma 4 if the operator $\mathcal{F}(x)$ is specialized to a diagonal matrix. The extension to multiple eigenvalues is possible because *all* the eigenvalues are taken into account simultaneously. For large matrices, this may again seem inconvenient since it will lead to large SDPs in (20).

Lemma 5: The mapping $h(x)$ defined by (21) is continuous.

The proof follows along the lines of the previous Lemma and is therefore omitted.

Let us now see how $\theta(x)$ may be used to quantify descent of $f = \lambda_1 \circ \mathcal{F}$ at x . Assume $0 \notin \partial f(x)$, so that $\theta(x) < 0$. Using the directional derivative of f at x in direction $h(x)$, we obtain

$$\begin{aligned} f'(x; h(x)) &= \sup_{Y \in \mathbb{B}} h(x)^T (\mathcal{F}'(x)^* Q_1 Y Q_1^H) \\ &= \sup_{Y \in \mathbb{B}} \text{Tr}(Y Q_1^H [\mathcal{F}'(x) h(x)] Q_1) \\ &\leq \theta(x) - \frac{1}{2\sigma} \|h(x)\|^2 < \theta(x) < 0 \end{aligned}$$

which follows readily from the primal formula (19) for θ if we use the fact that $-f(x) + \mu_1(X) = 0$. In consequence we have the following

Lemma 6: . Let $0 < \tau < 1$ be fixed. There exists $\epsilon > 0$ such that $f(x' + th(x')) - f(x') < t\tau\theta(x')$ for every $x' \in B(x, \epsilon)$ and all $t > 0$ sufficiently small.

Proof: By the definition of the directional derivative and the fact that $0 < \tau < 1$ we have

$$\begin{aligned} f(x + th(x)) - f(x) &\leq t\tau f'(x; h(x)) \\ &\leq t\tau \left(\theta(x) - \frac{1}{2\sigma} \|h(x)\|^2 \right) \quad (22) \\ &< t\tau\theta(x) \end{aligned}$$

for some $t_0 > 0$ and all $0 < t < t_0$. Since $h(\cdot)$ and $\theta(\cdot)$ are continuous, we can find a neighborhood $B(x, \epsilon)$ of x such that $f(x' + th(x')) - f(x') < t\tau\theta(x')$ for all $x' \in B(x, \epsilon)$ and every $0 < t < t_0$. This proves the claim. ■

In order to construct $s(\cdot)$, we follow [36, 3c, p. 223] and define $s(x) = x + t(x)h(x)$, where

$$t(x) := \sup\{2^{-k} : k \in \mathbb{N}, f(x + 2^{-k}h(x)) - f(x) < 2^{-k}\tau\theta(x)\}.$$

The supremum is over a nonempty set because of (22), hence $0 < t(x) < +\infty$. Let $k(x)$ be the integer where the supremum $t(x)$ is attained. Let us check property (ii) with $\epsilon > 0$ as in Lemma 6 and $\delta = -2^{-k(x)-1}\tau\theta(x) > 0$. Let $x' \in B(x, \epsilon)$. Since $t(x)$ satisfies (22), we have $f(x' + t(x)h(x')) - f(x') < t(x)\tau\theta(x')$ by Lemma 6. Therefore $t(x) = 2^{-k(x)}$ is admitted for the supremum $t(x')$, which implies $t(x') \geq t(x)$. Therefore $f(s(x')) - f(x') = f(x' + t(x')h(x')) - f(x') < t(x')\tau\theta(x') \leq t(x)\tau\theta(x') \leq t(x)\tau\theta(x)/2 = -2^{-k(x)-1}\tau\theta(x) = -\delta$, when we assume that ϵ is chosen sufficiently small to assure $\theta(x') \leq \theta(x)/2$ for every $x' \in B(x, \epsilon)$. This proves (ii) for the above choice of $t(x)$. Other choices of $t(x)$ like for instance [39] are possible.

E. The semi-infinite case

Our last step is now to address the semi-infinite case. We are in the situation (17). Suppose that for finite $\Omega \subset \mathbb{R} \cup \{-\infty, +\infty\}$ we already dispose of a descent step generator s_Ω for f_Ω , satisfying axioms (i) and (ii). This is naturally the case when the $f(x, \omega)$ are maximum eigenvalue functions, because a finite maximum of maximum eigenvalue functions is itself a maximum eigenvalue function. So here we obtain s_Ω as in sections VI-C or VI-D. Suppose now that we can specify a sequence of finite sets $\Omega_1 \subset \Omega_2 \subset \dots$ such that the following conditions are satisfied:

- (iii) If $0 \notin \partial f(x)$, then $\limsup_{k \rightarrow \infty} f(s_{\Omega_k}(x)) - f(x) < 0$.
- (iv) For every $x \in \mathbb{R}^n$ and $\epsilon > 0$,

$$\lim_{k \rightarrow \infty} \max_{x' \in B(x, \epsilon)} \inf_{\omega' \in \Omega_k} |f(x', \omega) - f(x', \omega')| = 0.$$

Notice that both axioms guarantee that the approximation Ω_k improves with growing k . Axiom (iii) tells that as soon as $0 \notin \partial f(x)$, descent steps may be eventually generated by using approximations f_{Ω_k} of f . Axiom (iv) is simply saying that approximations get better as k increases, and that this happens uniformly on small neighborhoods of each x .

Using these axioms, let us construct a descent step generator $s(\cdot)$ for f . We proceed as follows:

Semi-infinite descent step generator

- 1) If $0 \in \partial f(x)$ then $s(x) = x$ and return. Otherwise put counter $k = 1$ and continue.
- 2) At counter k , if $0 \in \partial f_{\Omega_k}(x)$, then increase k until $0 \notin \partial f_{\Omega_k}(x)$.
- 3) At counter k with $0 \notin \partial f_{\Omega_k}(x)$, compute the descent step $s_{\Omega_k}(x)$ for f_{Ω_k} at x . Let $\epsilon_k, \delta_k > 0$ be such that $f_{\Omega_k}(s_k(x')) - f_{\Omega_k}(x') \leq -\delta_k < 0$ for every $x' \in B(x, \epsilon_k)$, as guaranteed by axiom (ii) for s_{Ω_k} .
- 4) Compute $\eta_k := \inf\{|f(x', \omega) - f(x', \omega')| : \omega' \in \Omega_k, x' \in B(x, \epsilon_k) \cup s(B(x, \epsilon_k))\}$.
- 5) If $3\eta_k < \delta_k$, then let $s(x) = s_{\Omega_k}(x)$ and stop. Otherwise increase k by one and go back to step 3.

We have to make sure that this scheme is well-defined and introduces a step generator $s(\cdot)$ for the infinite maximum function f .

Lemma 7: . The descent step generator $s(\cdot)$ for $f(x) = \sup_{\omega \in \mathbb{R}} f(x, \omega)$ is well-defined and satisfies axioms (i) and (ii). Moreover, each of the above loops ends after a finite number of iterations.

Proof: Notice first that if $0 \notin \partial f(x)$, then descent around x is possible. Since $f_{\Omega_k}(x) \rightarrow f(x)$ as $k \rightarrow \infty$, we can also decrease f_{Ω_k} around x for k sufficiently large. So step 2 ends with a descent step of f_{Ω_k} at x after a finite number of trials. Moreover, this remains so for the following counters k , because $\Omega_k \subset \Omega_{k+1}$.

Next observe that $\eta_k \rightarrow 0$ by axiom (iv), while $\limsup -\delta_k < 0$ by axiom (iii). That means $3\eta_k < \delta_k$ for k sufficiently large, i.e., the procedure ends in step 4 after a finite number of updates $k \rightarrow k + 1$. It remains to check that s is as required.

Suppose $f_{\Omega_k}(s_{\Omega_k}(x')) - f_{\Omega_k}(x') \leq -\delta_k$ for every $x' \in B(x, \epsilon_k)$. Then $|f(s_{\Omega_k}(x')) - f_{\Omega_k}(s_{\Omega_k}(x'))| \leq \eta_k$, $|f(x') - f_{\Omega_k}(x')| \leq \eta_k$, hence $f(s_{\Omega_k}(x')) - f(x') \leq -\delta_k + 2\eta_k \leq -\frac{1}{3}\delta_k$. This proves axiom (ii). ■

The procedure is sufficiently flexible to accommodate a problem oriented step generation. We may adapt the choice of Ω_k to the structure of f , and what is more important, to the local behavior of f around the current x .

F. First-order algorithm for the H_∞ norm

We are now ready to present our algorithm, which follows the lines of the previous sections. We start with a version based on sections VI-C and VI-E.

First-order algorithm for the H_∞ norm: variant I

- 0) Fix $0 < \tau < 1$ and choose initial point x_0 .
- 1) Given x_k choose a finite set Ω_k containing $\Omega(x_k)$.
- 2) Compute $M(x_k)$ for f_{Ω_k} according to (18). If $M(x_k) = 0$ then stop, because $0 \in \partial f(x)$. Otherwise
- 3) Use a line search to find a step t_k such that the predicted decrease satisfies $\pi_k = f_{\Omega_k}(x_k + t_k h_k) - f(x_k) \leq -\kappa(x_k)M(x_k) < 0$. Compare with the actual decrease $\alpha_k = f(x_k + t_k h_k) - f(x_k)$. If $\alpha_k \leq \tau\pi_k$, accept t_k , put $x_{k+1} = x_k + t_k h_k$ and goto step 5.

- 4) If $\alpha_k > \tau\pi_k$ then reject t_k and add nodes to Ω_k to obtain the finer mesh Ω_{k+1} . Increase counter k by one and go back to step 2.
- 5) Increase counter k by one and go back to step 1.

In steps 3 and 4 of the algorithm we recognize the mechanism of the previous section, which obtains a descent step generator for the semi-infinite function by using those of the finite models f_{Ω_k} . We accept the step for f if it exceeds a small fraction of the descent predicted by the finite model f_{Ω_k} . This is essentially the same procedure as in VI-E.

Next comes the version based on section VI-D in tandem with section VI-E. The semi-infinite case is handled in exactly the same fashion, but the descent step generators s_{Ω_k} are different.

First-order algorithm for the H_∞ norm: variant II

- 0) Fix $0 < \tau_1, \tau_2 < 1$ and $\delta > 0$ and choose initial point x_0 .
- 1) Given x_k choose a finite set Ω_k containing $\Omega(x_k)$.
- 2) Compute the value θ_k and the solution (τ^k, Y^k) of the SDP (19). If $\theta_k = 0$, then stop because $0 \in \partial f(x_k)$. Otherwise compute the descent direction h_k for f_{Ω_k} at x_k according to (21).
- 3) Using a line search, find a step t_k such that the predicted decrease satisfies $\pi_k = f_{\Omega_k}(x_k + t_k h_k) - f(x_k) < \tau_1 t_k \theta_k < 0$. Compare with the actual decrease $\alpha_k = f(x_k + t_k h_k) - f(x_k)$. If $\alpha_k \leq \tau_2 \pi_k$, accept t_k , put $x_{k+1} = x_k + t_k h_k$ and goto step 5.
- 4) If $\alpha_k > \tau_2 \pi_k$ then reject t_k and add nodes to Ω_k to obtain the finer mesh Ω_{k+1} . Increase counter k by one and go back to step 2.
- 5) Increase counter k by one and go back to step 1.

G. Practical aspects

In this section we comment on the salient features of the nonsmooth first-order descent algorithm and address some of the practical aspects.

In our testing we have observed that often the leading eigenvalues $\lambda_1(\mathcal{F}(x_k, j\omega_\nu))$ have multiplicity 1 at all frequencies $\omega_\nu \in \Omega(x_k)$. In this situation step 2 of the descent algorithm variant II could be simplified. Suppose at the current iterate x_k we have selected a finite set of frequencies $\Omega_k = \{\omega_1, \dots, \omega_q\}$ containing $\Omega(x_k)$. Suppose all $\lambda_1(\mathcal{F}(x_k, j\omega_\nu))$ have multiplicity 1. Let $g_\nu = \mathcal{F}'(x_k, j\omega_\nu)^* e_\nu e_\nu^T = f'(x_k; \omega_\nu)$, where e_ν is the normalized eigenvector associated with $\lambda_1(\mathcal{F}(x_k, j\omega_\nu))$. Then the semidefinite program (16) respectively (20) simplifies to a convex quadratic program $\theta(x_k) =$

$$\sup_{\tau \geq 0, \sum_{\nu=1}^q \tau_\nu = 1} -f_\Omega(x_k) + \sum_{\nu=1}^q \tau_\nu f(x_k, \omega_\nu) - \frac{1}{2\sigma} \left\| \sum_{\nu=1}^q \tau_\nu g_\nu \right\|^2$$

and the associated direction of descent $h_k = h(x_k)$ is

$$h(x_k) = -\frac{1}{\sigma} \sum_{\nu=1}^q \tau_\nu g_\nu,$$

where τ is the optimal solution of the quadratic program. Observe that for $\Omega_k = \Omega(x_k)$, $h(x_k)$ coincides with the steepest descent direction for a finite max function.

Let us now specify in which way we select the frequency set Ω_k at each step. The finite set of frequencies $\Omega(x_k)$ where the H_∞ norm is attained is computed via the Hamiltonian technique [10]. We then form an enriched set Ω_k of frequencies by adding to the peak frequencies a collection of logarithmically spaced frequencies ω_ν such that

$$\|T_{w \rightarrow z}(K)\|_\infty - \overline{\sigma}(T_{w \rightarrow z}(K, j\omega_\nu)) \leq \epsilon_\omega \|T_{w \rightarrow z}(K)\|_\infty,$$

where ϵ_ω is a user-specified tolerance. We usually limit the set to the first 50 frequencies with largest singular values, as this appears to work well on a broad range of numerical tests. Typical values for ϵ_ω range from 0.05 to 0.5. The algorithm requires that this set be iteratively refined when descent steps cannot be computed, but in practice our choice is usually satisfactory, and numerical problems due to exceedingly fine Ω_k can be avoided.

VII. NUMERICAL EXPERIMENTS

In this section we test our nonsmooth algorithms on a variety of synthesis problems from the *COMPlib* collection by F. Leibfritz [40]. Computations were performed on a (low-level) SUN-Blade Sparc with 256 RAM and a 650 MHz sparcv9 processor. LMI-related computations for search directions used the LMI Control Toolbox [41] or our home made SDP code [6] while QP computations are based on Schittkowski's code [42].

Our algorithm is a first-order method. Not surprisingly, it may be slow in the neighborhood of a local solution. We have implemented various stopping criteria to ensure that an adequate approximation of a solution has been found and to avoid unwarranted computational efforts as is often the case with a first-order algorithm. The first of these termination criteria is an absolute stopping test, which provides a criticality assessment

$$\inf\{\|g\| : g \in \partial f(x)\} < \epsilon_1, \quad (23)$$

This is reasonable, as $0 \in \partial f(x)$ indicates a critical point. It is also mandatory to use relative stopping criteria to reduce the dependence on the problem scaling. The test

$$\|T_{w \rightarrow z}(K)\|_\infty - \|T_{w \rightarrow z}(K^+)\|_\infty < \epsilon_2(1 + \|T_{w \rightarrow z}(K)\|_\infty), \quad (24)$$

compares the progress achieved relatively to the current H_∞ performance, while

$$\|K^+ - K\| < \epsilon_3(1 + \|K\|) \quad (25)$$

compares the step-length to the controller gains. The tolerances

$$\epsilon_1 = 1e-5, \epsilon_2 = 1e-3, \epsilon_3 = 1e-3$$

have been used in our numerical testing. For stopping we required that either the first two tests or the third one are satisfied. For the enriched set, the number of frequencies has been limited to 50. They are selected according to our discussion in section VI-G. It is sometimes possible to employ fewer frequencies, but generally better steps are performed when richer sets are used. Our choice appears reasonable and has been validated on numerous experiments. It does not restrict efficiency since QP codes are very efficient up to 500 variables.

A. Stabilization

Pure stabilization may be regarded as H_∞ synthesis under the special form (6). The optimization program is (5), but we stop the algorithm as soon as a stabilizing controller is obtained. Iterating until a local optimum of (5) is reached does not seem to improve any of the usual performance specifications of the stabilizing feedback controller. Such questions have to be addressed in a second procedure, where the stabilizing controller will serve as a starting point. The spectral abscissa (7) is used to check stability of K , but is not used as a cost function. Previous experiments where (7) was used to find stabilizing controllers are reported in [29].

Table I displays results obtained for static stabilization problems borrowed from the literature. The triple (n, m, p) gives the number of states, inputs and outputs. Column 'iter' corresponds to the number of iterations required to meet the combined stopping tests. Column α displays the final closed-loop spectral abscissa (7). Negative $\alpha < 0$ indicates a stable closed-loop system. Column 'cpu' gives the cputime in seconds. Note that the reported cputimes are only indicative as the α -version of our code includes a number of extra tests and even graphics that have been exploited for fine tuning of the various algorithm parameters.

The initial shift for the H_∞ norm minimization was chosen according to the following rule

$$a_0 := \max(\alpha(A) * (1 + 10\%), 1e-2), \quad (26)$$

with a zero initial static controller $K_0 = 0$.

From the table, we observe that in all cases but one, the original choice a_0 of the shift was successful. As indicated, there is no need to run the algorithm until convergence. A stabilizing solution is therefore obtained very early. A single SDP or convex QP with line search suffices in most cases.

In the second example, we had to reduce the shift three times before stability was reached. This was done according to rule (26), applied to the closed-loop dynamics. In this example the H_∞ criterion is flat about a (global) optimum, which allows longer steps and renders the step-length-based termination criterion more stringent. Note that in this case the solution is globally optimal because a zero H_∞ norm is reached.

TABLE I
STATIC OUTPUT-FEEDBACK STABILIZATION

problem	(n, m, p)	iter	α	cpu (sec.)
Transport airplane [43]	(9, 1, 5)	1	$-2.22e-2$	2.26
Horisberger's example [44]	(9, 1, 4)	4	-0.01	12.0
VTOL helicopter [45]	(4, 2, 1)	1	$-6.00e-2$	1.90
Chemical reactor [46]	(4, 2, 2)	1	-1.73	2.26
Piezoelectric actuator [47]	(5, 1, 3)	1	$-9.95e-1$	2.88
Boeing 767 [48]	(55, 2, 2)	1	$-2.33e-2$	6.47

steepest descent method - $\epsilon_\omega = 0.05$

For test reasons, the pure stabilization problems have been solved with the steepest descent method in section V. As expected, this technique is less stable than algorithmic variants I and II. Specifically, the choice of ϵ_ω is rather critical and finding a general selection rule appears difficult. Together

with our analysis in section VI, this encouraged us to use algorithmic variant II in the syntheses problems below.

B. H_∞ synthesis problems

This section is devoted to H_∞ synthesis problems. The synthesis procedure is based on the scheme (3) and must be initialized with a stabilizing controller. This initial phase I is described in the previous section, but alternative techniques to find an initial stabilizing controller K may prove useful [29]. All examples are extracted from the *COMPl_eib* collection [40]. Column 'problem' now indicates the *COMPl_eib* acronym attached to each example.

Notice that for a static K , our program could formally be solved as

$$\begin{aligned} &\text{minimize} \quad \|T_{w \rightarrow z}(K)\|_\infty, \quad K \in \mathbb{M}_{k+m_2, k+p_2} \\ &\text{subject to} \quad \|C_2(sI - \mathcal{A}(K))^{-1} B_2\|_\infty \leq M \end{aligned} \quad (27)$$

where $M > 0$ is a constant e.g. such that the initial stabilizing controller K_0 satisfies

$$\|C_2(sI - \mathcal{A}(K_0))^{-1} B_2\|_\infty \ll M < +\infty$$

and where $\mathcal{A}(K)$ refers to the closed-loop system matrix. Note that the extra constraint in (27) maintains asymptotic stability of the closed-loop system during the optimization of K .

It often happens that stability of the performance channel $T_{w \rightarrow z}(K)$ alone already implies stability of the closed-loop system, so that the (internal) stability constraint in (27) is redundant. This observation is significant, because this constraint can be dropped, and the problem becomes unconstrained. In those cases where this does not hold true, we may if we want to avoid the big- M constraint in (27), introduce the composite problem

$$\text{minimize}_K \left\| \begin{bmatrix} T_{w \rightarrow z}(K) & 0 \\ 0 & \epsilon_s C_2(sI - \mathcal{A}(K))^{-1} B_2 \end{bmatrix} \right\|_\infty$$

where the lower-right term enforces internal stability, with $\epsilon_s > 0$ a small enough parameter. The new problem can now be handled by the techniques discussed so far without change. Similar modifications apply to design problem with additional structural constraints on the controller.

We compare the results of our nonsmooth algorithm variant II in columns 'nonsmooth H_∞ ' to older results obtained with the specialized augmented Lagrangian (AL) algorithm described in [17], displayed in columns ' H_∞ AL', and to results obtained with the Frank & Wolfe (FW) algorithm described in [14], column 'FW'. In column ' H_∞ full', we display the gain obtained with a full-order feedback controller, synthesized via the usual Riccati-based DGKF technique [2]. Note that it gives the best achievable H_∞ performance and thus provides a lower bound for other techniques.

The results which we obtain with our nonsmooth technique are usually close to those previously obtained with the augmented Lagrangian method [17], except for problems with large state dimension as 'AC10' (55 states), 'BDT2' (82 states), 'HF1' (130 states) and 'CM4' (240 states), where the augmented Lagrangian method fails, while the present nonsmooth method is still functional. For these large systems, 'AC10', 'BDT2', 'HF1' and 'CM4', we have observed that

TABLE II
 H_∞ SYNTHESIS WITH NONSMOOTH ALGORITHMIC VARIANT II

problem	(n, m, p)	order	iter	cpu (sec.)	nonsmooth H_∞	H_∞ AL	FW	H_∞ full
AC8	(9, 1, 5)	0	20	45	2.005	2.02	2.612	1.62
HE1	(4, 2, 1)	0	4	7	0.154	0.157	0.215	0.073
REA2	(4, 2, 2)	0	31	51	1.192	1.155	1.263	1.141
AC10	(55, 2, 2)	0	15	294	13.11	*	*	3.23
AC10	(55, 2, 2)	1	46	408	10.21	*	*	3.23
BDT2	(82, 4, 4)	0	44	1501	0.8364	*	*	0.2340
HF1	(130, 1, 2)	0	11	1112	0.447	*	*	0.447
CM4	(240, 1, 2)	0	2	3052	0.816	*	*	*

$\epsilon_\omega = 0.05$, '*': problem is intractable

TABLE III
 H_∞ CONTROLLERS FOR LARGE PROBLEMS WITH ALGORITHM II

problem	order	$K(s)$					
AC10	0	$K =$		$-9.0747e-1$	$2.1249e-5$		
				4.3042	$2.2467e-5$		
AC10	1	$\begin{bmatrix} -6.110e-2 \\ -2.361e-2 \end{bmatrix}$	$(s + 2.633)^{-1} [$	$-6.476e-2$	$1.245e-2]$	$+ \begin{bmatrix} 2.698e-1 & 1.389e-5 \\ -3.285e-01 & 5.226e-5 \end{bmatrix}$	
BDT2	0	$\begin{bmatrix} -1.0402 & 1.2997 & 1.4684 & 3.0555 \\ -3.7306 & -3.5572 & -1.9894 & 2.2705 \\ -2.8162 & 2.1944 & 1.8146 & 9.6202 \\ -9.8493 & -5.5376 & -2.4885 & 5.1699 \end{bmatrix}$					
HF1	0	$\begin{bmatrix} -0.1907 & -1.3093 \end{bmatrix}$					
CM4	0	$\begin{bmatrix} -4.5684 & -0.6667 \end{bmatrix}$					

even Riccati or LMI solvers encounter serious difficulties or even break down.

Notice that our present nonsmooth technique (NS) and the AL method are rigorous in the sense that they converge to local minima (critical points). It is therefore not surprising that NS and AL often achieve the same H_∞ performance at the same K , for which optimality was established in [18].

Note that in its original form, the FW method cannot solve problems where performance is optimized under constraints on the order of the controller. We have therefore encapsulated FW into a dichotomy search in order to assure the best possible performance. According to our numerical experiments, AL and FW, which solve SDPs at every iteration, are no longer functional even for medium size systems such as the Boeing problem 'AC10'. Higher-order problems like 'BDT2', 'HF1', and 'CM4' are completely intractable with these techniques, while the nonsmooth method continues to produce valid solutions. We observed that FW, when functional, is outperformed both by AL and NS. We attribute this to the fact that, as opposed to AL and NS, FW is not supported by a sound convergence theory, and therefore often stops at iterates K which are not even locally optimal. Examples of K where this is the case are easily identified. It suffices to start AL or NS at iterates K where FW stops which almost always leads to further improvement. As an example, when we initialized NS with the FW solution in example 'AC8', the H_∞ performance was improved from 2.612 to 2.005. A phenomenon that we have also observed in examples 'HE1' and 'REA2'. This is a strong argument in favor of those optimization techniques, which generate steps based on local convergence theory. It means that NS and AL can be used to certify criticality of controllers obtained with alternative methods.

As an illustration of the nonsmooth technique, Fig. 1 shows the evolution of the maximum singular value of $T_{w \rightarrow z}(K)$ for example 'AC8' during the first 5 iterations. The stars indicate the frequencies ω_i that were regrouped in the set Ω_k to construct a bundle of Clarke subgradients at iterate x_k . Also, Fig. 2 depicts the evolution of the absolute value of the criticality measure of algorithm variant II (19) and (20) versus iterations. As theoretically expected, $\theta(x_k)$ gradually tends to zero until a local minimum is reached. Finally, controllers for large problems 'AC10', 'BDT2', 'HF1' and 'CM4' are given in Table III.

VIII. CONCLUSION

We have proposed several new algorithms to minimize the H_∞ norm subject to structural constraints on the controller dynamics. The proposed method uses nonsmooth techniques suited for H_∞ synthesis and for semi-infinite eigenvalue or singular value optimization programs. Variant I and variant II of our algorithm are supported by global convergence theory, a crucial parameter for the reliability of an algorithm in practice. Variant II has been shown to perform satisfactorily on a number of difficult examples. In particular, four examples with large state dimension ($n = 55$, $n = 82$, $n = 130$ and $n = 240$) have been solved.

Note that the proposed tools and techniques easily extend to multidisk problems and synthesis problems on prescribed frequency intervals [49]. More importantly, they pave the way for investigating an even larger scope of synthesis problems, characterized through frequency domain inequalities of the form $\lambda_1(H(x, \omega)) \leq 0$, $\omega \geq 0$, where $H(x, \omega)$ is Hermitian-valued and x stands for controller parameters and possibly multiplier variables, as is the case when IQC formulations are

Fig. 1. Max singular values (transport airplane 'AC8') versus frequency - first 5 iterations - '*' selected frequencies

Fig. 2. Criticality measure $\theta(x)$ versus iterations for 'BDT2'

used [50]. This is a strong incentive for further development and research. Also, a second-order version of our technique with enhanced asymptotic convergence is currently under investigation [51].

ACKNOWLEDGMENTS

Thanks to Friedemann Leibfritz, Universität Trier, for providing the *COMPl_eib* collection and for fruitful discussions.

REFERENCES

- [1] P. Gahinet and P. Apkarian, "A linear matrix inequality approach to H_∞ control," *Int. J. Robust and Nonlinear Control*, vol. 4, pp. 421–448, 1994.
- [2] J. Doyle, K. Glover, P. Khargonekar, and B. A. Francis, "State-Space Solutions to Standard H_2 and H_∞ Control Problems," *IEEE Trans. Aut. Control*, vol. AC-34, no. 8, pp. 831–847, august 1989.
- [3] A. Nemirovskii, "Several NP-Hard problems arising in robust stability analysis," *Mathematics of Control, Signals, and Systems*, vol. 6, no. 1, pp. 99–105, 1994.
- [4] V. Blondel and M. Gevers, "Simultaneous stabilizability question of three linear systems is rationally undecidable," *Mathematics of Control, Signals, and Systems*, vol. 6, no. 1, pp. 135–145, 1994.
- [5] D. Henrion, M. Kocvara, and M. Stingl, "Solving simultaneous stabilization BMI problems with PENNON," in *IFIP Conference on System Modeling and Optimization*, vol. 7, Sophia Antipolis, France, July 2003.
- [6] P. Apkarian, D. Noll, J. B. Thevenet, and H. D. Tuan, "A spectral quadratic-SDP method with applications to fixed-order H_2 and H_∞ synthesis," in *Asian Control Conference*, Melbourne, AU, 2004.
- [7] J. Thevenet, D. Noll, and P. Apkarian, "Nonlinear spectral SDP method for BMI-constrained problems: Applications to control design," *European J. of Control*, vol. 10, no. 6, pp. 527–538, 2004.
- [8] F. Leibfritz and E. M. E. Mostafa, "An interior point constrained trust region method for a special class of nonlinear semi-definite programming problems," *SIAM J. on Control and Optimization*, vol. 12, pp. 1048–1074, 2002.
- [9] F. Leibfritz and E. M. E. Mostafa, "Trust region methods for solving the optimal output feedback design problem," *International Journal of Control*, vol. 76, pp. 501–519, 2000.
- [10] S. Boyd, V. Balakrishnan, and P. Kabamba, "A bisection method for computing the H_∞ norm of a transfer matrix and related problems," *Mathematics of Control, Signals, and Systems*, vol. 2, no. 3, pp. 207–219, 1989.
- [11] S. Boyd and V. Balakrishnan, "A regularity result for the singular values of a transfer matrix and a quadratically convergent algorithm for computing its L_∞ -norm," *Syst. Control Letters*, vol. 15, pp. 1–7, 1990.
- [12] P. Gahinet and P. Apkarian, "Numerical computation of the L_∞ norm revisited," in *Proc. IEEE Conf. on Decision and Control*, 1992.
- [13] M. C. de Oliveira and J. C. Geromel, "Numerical comparison of output feedback design methods," in *Proc. American Control Conf.*, Albuquerque, NM, June 1997, pp. 72–76.
- [14] L. ElGhaoui, F. Oustry, and M. AitRami, "An algorithm for static output-feedback and related problems," *IEEE Trans. Aut. Control*, vol. 42, no. 8, pp. 1171–1176, 1997.
- [15] K. M. Grigoriadis and R. E. Skelton, "Low-order control design for LMI problems using alternating projection methods," *Automatica*, vol. 32, no. 8, pp. 1117–1125, Sept. 1996.
- [16] M. Frank and P. Wolfe, "An algorithm for quadratic programming," *Naval Res. Log. Quart.*, vol. 3, pp. 95–110, 1956.
- [17] P. Apkarian, D. Noll, and H. D. Tuan, "Fixed-order H_∞ control design via an augmented Lagrangian method," *Int. J. Robust and Nonlinear Control*, vol. 13, no. 12, pp. 1137–1148, 2003.
- [18] D. Noll, M. Torki, and P. Apkarian, "Partially augmented Lagrangian method for matrix inequality constraints," *SIAM J. on Optimization*, vol. 15, no. 1, pp. 161–184, 2004.
- [19] X. Xin, "Reduced-order controllers for the H_∞ control problem with unstable invariant zeros," *Automatica*, vol. 40, no. 2, pp. 319–326, Feb. 2004.
- [20] X. Xin, L. Guo, and C. Feng, "Reduced-order controllers for continuous and discrete-time singular H_∞ control problems based on LMI," *Automatica*, vol. 32, no. 11, pp. 1581–1585, Nov. 1996.
- [21] A. A. Stoorvogel, A. Saberi, and B. M. Chen, "A reduced-order observer based controller design for H_∞ -optimization," *IEEE Trans. Aut. Control*, vol. 39, no. 2, pp. 355–360, Feb. 1994.
- [22] F. H. Clarke, *Optimization and Nonsmooth Analysis*, ser. Canadian Math. Soc. Series. New York: John Wiley & Sons, 1983.
- [23] J. Cullum, W. Donath, and P. Wolfe, "The minimization of certain nondifferentiable sums of eigenvalues of symmetric matrices," *Math. Programming Stud.*, vol. 3, pp. 35 – 55, 1975.

- [24] C. Lemaréchal and F. Oustry, "Nonsmooth algorithms to solve semidefinite programs," *SIAM Advances in Linear Matrix Inequality Methods in Control series*, ed. L. El Ghaoui & S.-I. Niculescu, 2000.
- [25] F. Oustry, "A second-order bundle method to minimize the maximum eigenvalue function," *Math. Programming Series A*, vol. 89, no. 1, pp. 1–33, 2000.
- [26] B. Anderson and S. Vongpanitlerd, *Network analysis and synthesis: a modern systems theory approach*. Prentice-Hall, 1973.
- [27] C. A. Desoer and M. Vidyasagar, *Feedback Systems: Input-Output Properties*. Academic Press, New York, 1975.
- [28] S. Boyd and C. Barratt, *Linear Controller Design: Limits of Performance*. Prentice-Hall, 1991.
- [29] P. Apkarian and D. Noll, "Controller design via nonsmooth multidirectional search," *SIAM J. on Control and Optimization*, 2005.
- [30] J. Burke, A. Lewis, and M. Overton, "Two numerical methods for optimizing matrix stability," *Linear Algebra and its Applications* 351–352, pp. 147–184., 2002.
- [31] —, "Robust stability and a criss-cross algorithm for pseudospectra," *IMA Journal of Numerical Analysis*, vol. 23, pp. 1–17, 2003.
- [32] —, "A robust gradient sampling algorithm for nonsmooth, nonconvex optimization," *SIAM J. Optimization*, vol. 15, pp. 751–779, 2005.
- [33] K. Glover, "All optimal Hankel-norm approximations of linear multivariable systems and their L_∞ -error bounds," *Int. J. Control*, vol. 39, pp. 1115–1193, 1984.
- [34] R. A. Horn and C. A. Johnson, *Matrix Analysis*. Cambridge University Press, 1985.
- [35] E. Polak, "On the mathematical foundations of nondifferentiable optimization in engineering design," *SIAM Rev.*, vol. 29, pp. 21–89, Mar. 1987.
- [36] —, *Optimization : Algorithms and Consistent Approximations*. Applied Mathematical Sciences, 1997.
- [37] D. Noll and P. Apkarian, "Spectral bundle methods for nonconvex maximum eigenvalue functions: first-order methods," *Mathematical Programming Series B*, 2005.
- [38] —, "Spectral bundle methods for nonconvex maximum eigenvalue functions: second-order methods," *Mathematical Programming Series B*, 2005.
- [39] O. Pironneau and E. Polak, "On the rate of convergence of a class of methods of centers," *Mathematical Programming*, vol. 2, no. 2, pp. 230–258, 1972.
- [40] F. Leibfritz, "COMPL_eIB, CONstraint Matrix-optimization Problem Library - a collection of test examples for nonlinear semidefinite programs, control system design and related problems," Universität Trier, Tech. Rep., 2003.
- [41] P. Gahinet, A. Nemirovski, A. J. Laub, and M. Chilali, *LMI Control Toolbox*. The MathWorks Inc., 1995.
- [42] K. Schittkowski, "QLD: A Fortran code for quadratic programming," Mathematisches Institut, Universität Bayreuth, Germany, Tech. Rep., 1986.
- [43] D. Gangsaas, K. Bruce, J. Blight, and U.-L. Ly, "Application of modern synthesis to aircraft control: Three case studies," *IEEE Trans. Aut. Control*, vol. AC-31, no. 11, pp. 995–1014, Nov. 1986.
- [44] H. P. Horisberger and P. R. Belanger, "Solution of the optimal constant output feedback problem by conjugate gradients," *IEEE Trans. Aut. Control*, vol. 19, pp. 434–435, 1974.
- [45] L. H. Keel, S. P. Bhattacharyya, and J. W. Howze, "Robust control with structured perturbations," *IEEE Trans. Aut. Control*, vol. 36, pp. 68–77, 1988.
- [46] Y. S. Hung and A. G. J. MacFarlane, *Multivariable feedback: A classical approach*, ser. Lectures Notes in Control and Information Sciences. Springer Verlag, New York, Heidelberg, Berlin, 1982.
- [47] B. M. Chen, *H_∞ Control and Its Applications*, ser. Lectures Notes in Control and Information Sciences. Springer Verlag, New York, Heidelberg, Berlin, 1998, vol. 235.
- [48] E. J. Davison, (ed.), "Benchmark problems for control system design," Oxford, Pergamon Press, IFAC Technical Committee Reports, Tech. Rep., 1990.
- [49] P. Apkarian and D. Noll, "Nonsmooth optimization for multidisk H_∞ synthesis," *submitted*, 2005.
- [50] A. Megretski and A. Rantzer, "System analysis via integral quadratic constraints," *IEEE Trans. Aut. Control*, vol. 42, no. 6, pp. 819–830, June 1997.
- [51] V. Bompert, D. Noll, and P. Apkarian, "Second-order nonsmooth optimization for H_∞ and H_2 syntheses," *in preparation*, 2005.

Pierre Apkarian received the Ph.D degree in control engineering from the "Ecole Nationale Supérieure de l'Aéronautique et de l'Espace" (ENSAE) in 1988. He was qualified as a professor from the Université Paul Sabatier in both control engineering and applied mathematics in 1999 and 2001, respectively.

Since 1988, he is research scientist at ONERA-CERT and associate Professor in the Maths department of Université Paul Sabatier in Toulouse. His research interests include robust and gain-scheduling control theory, LMI techniques and mathematical programming and applications in aeronautics. Pierre Apkarian has served as an associate editor for the IEEE Transactions on Automatic Control.

Dominikus Noll received his PhD in 1983 and his *venia legendi* in 1989 from the Universität Stuttgart (Germany). Since 1995 he is professor of applied mathematics at Université Paul Sabatier in Toulouse (France). D. Noll's research interests are in continuous optimization, medical imaging, and automatic control.