21st International Symposium on
Mathematical Theory of Networks and Systems
July 7-11, 2014. Groningen, The Netherlands

Structured Singular Values versus Diagonal Scaling: the Noncommutative Setting

Joseph A. Ball[1], Gilbert Groenewald[2], and Sanne ter Horst[3]

Abstract— **The structured singular value (often referred to simply as μ) was introduced independently by Doyle and Safanov as a tool for analyzing robustness of system stability and performance in the presence of structured uncertainty in the system parameters. While the structured singular value provides a necessary and sufficient criterion for robustness with respect to a structured ball of uncertainty, it is notoriously difficult to actually compute. The method of diagonal (or simply "D") scaling, on the other hand, provides an easily computable upper bound for the structured singular value, but provides an exact evaluation of μ (or even a useful upper bound for μ) only in special cases. However it was discovered in the 1990s that enhancement of the uncertainly structure to allow what can be interpreted as time-varying uncertainty (equivalently, letting the uncertainty parameters be freely noncommuting operators on an infinite-dimensional separable Hilbert space) resulted in the D-scaling procedure leading to an exact evaluation of μ. This report discusses recent refinements of these results.**

I. INTRODUCTION

We recall the definition of the structured singular values introduced independently by Doyle [10] and Safanov [20]; see [23] for a thorough more recent treatment. Let N be a positive integer with a partitioning $N = n_1 + \cdots + n_K + m_1 + \cdots + m_F$ for positive integers n_i ($i = 1, \dots, K$) and m_j ($j = 1, \dots, F$). We let $\mathbf{\Delta}$ denote the set of $N \times N$ matrices of the form

$$\begin{aligned}\mathbf{\Delta} =\{&\operatorname{diag}[\delta_1 I_{n_1}, \dots, \delta_K I_{n_K}, \Delta_1, \dots, \Delta_F]\colon \\ &\delta_i \in \mathbb{C},\, \Delta_j \in \mathbb{C}^{m_j \times m_j}\}.\end{aligned} \tag{1}$$

For an $N \times N$ matrix $M \in \mathbb{C}^{N\times N}$, we define the *structured singular value* of M by

$$\mu_{\mathbf{\Delta}}(M) := \frac{1}{\min\{\|\Delta\|\colon \Delta \in \mathbf{\Delta},\, 1 \in \sigma(M\Delta)\}}, \tag{2}$$

where in general $\sigma(X)$ denotes the spectrum of the square matrix X. Motivation for this notion comes from robust control theory (see [23], [11]). In the case where $K = 0$ and $F = 1$, the structured singular value $\mu_{\mathbf{\Delta}}(M)$ collapses to the largest singular value $\overline{\sigma}_1(M)$ of M or, equivalently, the induced operator norm of M as an operator on $\mathbb{C}^N$, where $\mathbb{C}^N$ is given the standard 2-norm. A key property of the largest singular value from the point of view of systems and control follows from the Small Gain Theorem.

Theorem 1.1 (Small Gain Theorem): Let $M \in \mathbb{C}^{N\times N}$ such that $\overline{\sigma}_1(M) < 1$. Then $I - \Delta M$ is invertible for all Δ with $\|\Delta\| \le 1$.

The systems and control interpretation of this result is that $\overline{\sigma}_1(M) < 1$ implies that perturbation of the 'plant' M with a multiplicative perturbation Δ does not affect stability of the closed-loop feedback as long as $\|\Delta\| \le 1$.

There are many applications in which the uncertainty parameter Δ is known to carry some structure, as in (1). In these cases it is enough that the structured singular value $\mu_{\mathbf{\Delta}}(M)$ be less than 1 to guarantee the maintenance of stability against structured multiplicative perturbations $\Delta \in \mathbf{\Delta}$ with $\|\Delta\| \le 1$.

However, it turns out that the structured singular value $\mu_{\mathbf{\Delta}}(M)$ is notoriously difficult to compute in a computationally efficient and reliable way. Indeed, computing the exact structured singular value $\mu_{\mathbf{\Delta}}(M)$ is an NP-hard problem (see [8]).

There is a convenient upper bound for $\mu_{\mathbf{\Delta}}(M)$, namely $\widehat{\mu}_{\mathbf{\Delta}}(M)$, defined as follows. We let $\mathbf{\Delta}'$ be the *commutant* of $\mathbf{\Delta}$:

$$\mathbf{\Delta}' = \{D \in \mathbb{C}^{N\times N}\colon D\Delta = \Delta D \text{ for all } \Delta \in \mathbf{\Delta}\}. \tag{3}$$

We then set

$$\widehat{\mu}_{\mathbf{\Delta}}(M) = \inf\{\|DMD^{-1}\|\colon D \in \mathbf{\Delta}' \text{ and } D \text{ invertible}\}.$$

It turns out that $\widehat{\mu}_{\mathbf{\Delta}}(M)$ can be computed accurately and efficiently. Indeed, to test whether $\widehat{\mu}_{\mathbf{\Delta}}(M) < 1$ it suffices to find a positive definite matrix $P \in \mathbf{\Delta}'$ which solves the structured Stein inequality

$$M^*PM - P \prec 0.$$

Note that the condition $P \in \mathbf{\Delta}'$ is equivalent to P having the block diagonal form

$$P = \operatorname{diag}[P_1, \dots, P_K, p_1 I_{m_1}, \dots, p_F I_{m_F}],$$

where P_i is a positive definite matrix of size $n_i \times n_i$ (for $i = 1, \dots, K$) and p_j a positive number (for $j = 1, \dots, F$). This puts the computation of $\widehat{\mu}_{\mathbf{\Delta}}$ within the framework of the MATLAB LMI toolbox.

While the general inequality $\mu_{\mathbf{\Delta}}(M) \le \widehat{\mu}_{\mathbf{\Delta}}(M)$ is easily derived, actual equality holds only in very special cases. In particular, equality holds for all M with respect to a given choice of structure specified by nonnegative integers K and F as in (1) if and only if $2K + F \le 3$ (see [18], [23], [11]).

---

[1] Joseph A. Ball is with the Department of Mathematics, Virginia Tech, Blacksburg, VA 24061 USA `joball@math.vt.edu`

2 Gilbert Groenewald is with the School of Computer, Statistical and Mathematical Sciences, North West University, Potchefstroom Campus, Potchefstroom 2520, SOUTH AFRICA `gilbert.groenewald@nwu.ac.za`

[3] Sanne ter Horst is with the Unit for Business Mathematics and Informatics, North West University, Potchefstroom Campus, Potchefstroom 2520, SOUTH AFRICA `Sanne.TerHorst@nwu.ac.za`

Moreover, even with K and F in (1) fixed, there is in general no bound on the gap between $\mu_{\mathbf{\Delta}}(M)$ and its upper bound $\widehat{\mu}_{\mathbf{\Delta}}(M)$; see [22]. Thus the compromise of using $\widehat{\mu}_{\mathbf{\Delta}}(M)$ as a substitute for $\mu_{\mathbf{\Delta}}(M)$ can be arbitrarily conservative.

However, if the structure is relaxed by letting the uncertainty parameters $\delta_i \in \mathbb{C}$ and the matrix entries $[\Delta_j]_{\alpha,\beta} \in \mathbb{C}$ be operators on a separable infinite-dimensional Hilbert space, say ℓ^2, then the modified μ is equal to its easily computable upper bound. To make this precise, we introduce the enhanced structure

$$\widetilde{\mathbf{\Delta}} = \{\mathrm{diag}[\widetilde{\delta}_1 \otimes I_{\mathbb{C}^{n_1}}, \dots, \widetilde{\delta}_K \otimes I_{\mathbb{C}^{n_K}}, \widetilde{\Delta}_1, \dots, \widetilde{\Delta}_F]\} \quad (4)$$

where each $\widetilde{\delta}_i \in \mathcal{L}(\ell^2)$ and each $\widetilde{\Delta}_j \in \mathcal{L}([\ell^2]^{m_j})$. Here $[\ell^2]^{m_j}$ denotes the m_j-fold direct sum of ℓ^2 with itself. We replace $M \in \mathbb{C}^{N\times N}$ with $\widetilde{M} = I_{\ell^2} \otimes M \in \mathcal{L}([\ell^2]^N)$ and define a new variation on $\mu(M)$ by

$$\widetilde{\mu}_{\mathbf{\Delta}}(M) := \mu_{\widetilde{\mathbf{\Delta}}}(I_{\ell^2} \otimes M).$$

It turns out that the two notions of $\widehat{\mu}$ are the same:

$$\widehat{\mu}_{\widetilde{\mathbf{\Delta}}}(I_{\ell^2} \otimes M) = \widehat{\mu}_{\mathbf{\Delta}}(M).$$

and hence the common value $\widehat{\mu}_{\mathbf{\Delta}}(M)$ is easily computable. The remarkable result is that this relaxed structured singular value is always equal to its easily computable upper bound, i.e.,

$$\widetilde{\mu}_{\mathbf{\Delta}}(M) = \widehat{\mu}_{\mathbf{\Delta}}(M). \quad (5)$$

This result can be found in the dissertation of Paganini [19]; the complete proof, as thoroughly elucidated in the book [11] (at least for the case where $K = 0$) draws on earlier ideas and results from Megretski-Treil [17] and Shamma [21]. Also there is an interpretation of the quantity $\widetilde{\mu}$ as robustness with respect to an enlarged block-structured uncertainty (see [11, page 246]).

As elegant as this result is, it is incomplete from a conceptual point of view since the structure given by (1) is limited in two respects:

(L1) There is an asymmetry between the scalar blocks and the full blocks in (1). A scalar block $\delta_i I_{n_i}$ can be considered as a full block with size $m_i = 1$, but with a repetition (or multiplicity) of n_i possibly larger than 1 allowed. On the other hand, the full blocks Δ_j are considered to be independently arbitrary with no repetitions allowed.

(L2) All blocks are considered square. There are interesting multidimensional input/state/output systems where this same structure occurs but with nonsquare blocks (see [4], [5]).

To remedy these limitations, we introduce a more general uncertainty structure based on a specific kind of bipartite graph in the next section, and we define the analogues of μ_Δ and $\widehat{\mu}_\Delta$ for this setting. In Section III we illustrate this structure with two examples. Section IV contains the definitions of the extended structure and the analogue of $\widetilde{\mu}_\Delta$, and it is argued here that the equality between the relaxed μ_Δ and its upper bound $\widehat{\mu}_\Delta$ still holds in this more general framework. Complete details will appear in [7]. The paper concludes with a section that consists of general remarks and observations.

II. THE GRAPH FORMALISM SET-UP

A. The graph

Throughout let $G = (V, E)$ be a finite simple bipartite graph such that each path-connected component of G is a complete bipartite graph. Here V denotes the set of vertices and E the set of edges. Since G is a bipartite graph, the vertex set V admits a decomposition $V = S \cup R$, with $S \cap R = \emptyset$, such that each edge $e \in E$ has one vertex in S (the source side) and one vertex in R (the range side); the vertex in S will be denoted by $s(e)$ and the vertex in R by $r(e)$. We write P for the set consisting of G's path-connected components, and for a vertex $v \in V$ we write $[v]$ to indicate the path-connected component that contains v. For each $p \in P$ we denote its vertex set and edge set by V_p and E_p, respectively. Each path-connected component $p \in P$ of G is also a simple bipartite graph and its vertex set V_p can be decomposed as $V_p = S_p \cup R_p$ with $S_p = S \cap V_p$ and $R_p = R \cap V_p$. By assumption, each $p \in P$ consists of all possible (nonoriented) edges connecting a vertex in S_p with a vertex in R_p, and no edge of G connects a vertex in $S_p \cup R_p$ with a vertex in $S_{p'} \cup R_{p'}$ if $p \neq p'$.

B. The uncertainty structure.

Next, with each path-connected component $p \in P$ we associate a C^*-algebra $\mathbf{\Delta}_p$ represented concretely as a C^*-subalgebra of $\mathcal{L}(\mathcal{H}_p)$ for some separable Hilbert space $\mathcal{H}_p$. We will also encounter the commutant of $\mathbf{\Delta}_p$ in $\mathcal{L}(\mathcal{H}_p)$ which we shall denote by $\mathbf{\Delta}_p'$:

$$\mathbf{\Delta}_p' = \{\Gamma \in \mathcal{L}(\mathcal{H}_p) :\ \Delta\Gamma = \Gamma\Delta, \text{ for each } \Delta \in \mathbf{\Delta}_p\}.$$

We set $\mathcal{H}_v = \mathcal{H}_{[v]}$ for each $v \in V$ and we further introduce the spaces

$$\begin{aligned} \mathcal{H}_S &= \bigoplus_{s\in S} \mathcal{H}_s, \quad \mathcal{H}_{S_p} = \bigoplus_{s\in S_p} \mathcal{H}_s \quad (p \in P),\\ \mathcal{H}_R &= \bigoplus_{r\in R} \mathcal{H}_r, \quad \mathcal{H}_{R_p} = \bigoplus_{r\in R_p} \mathcal{H}_r \quad (p \in P). \end{aligned} \quad (6)$$

For $s \in S$ we write ι_s for the canonical embedding of $\mathcal{H}_{[s]}$ into $\mathcal{H}_S$ that maps $\mathcal{H}_{[s]}$ onto the component in the direct sum $\mathcal{H}_s$ in (6) indexed by s via $\iota_s h = \oplus_{s'}(\delta_{s',s} h)$ for $h \in \mathcal{H}_{[s]}$ (with $\delta_{s',s}$ equal to the Kronecker delta). Similarly, for $r \in R$ we write ι_r for the canonical embedding of $\mathcal{H}_{[r]}$ into $\mathcal{H}_R$. Note that ι_s (resp. ι_r) acts on $\mathcal{H}_{[s]}$ (resp. $\mathcal{H}_{[r]}$) and not on $\mathcal{H}_s$ (resp. $\mathcal{H}_r$), so that for an $e \in E$ the product $\iota_{s(e)}\iota_{r(e)}^*$ is properly defined.

We now define the set $\mathbf{\Delta}_G$ associated with the graph G and C^*-algebras $\mathbf{\Delta}_p$, $p \in P$, by

$$\mathbf{\Delta}_G = \{Z_G(\Delta) :\ \Delta \in \mathbf{\Delta}_E\} \subset \mathcal{L}(\mathcal{H}_R, \mathcal{H}_S), \quad (7)$$

where

$$\mathbf{\Delta}_E = \{(\Delta_e)_{e\in E} :\ \Delta_e \in \mathbf{\Delta}_{[s(e)]}\}$$

and $Z_G : \mathbf{\Delta}_E \to \mathcal{L}(\mathcal{H}_R, \mathcal{H}_S)$ is the map given by

$$Z_G(\Delta) = \sum_{e\in E} \iota_{s(e)}\Delta_e\iota^*_{r(e)} \quad (\Delta = (\Delta_e)_{e\in E} \in \mathbf{\Delta}_E). \tag{8}$$

Note that this notation is somewhat condensed: the set of operators $\mathbf{\Delta}_G$ in (7) depends not only on the admissible bipartite graph G but also on the choice of the C^*-algebra $\mathbf{\Delta}_p$ and assignment of Hilbert space $\mathcal{H}_p$ such that $\mathbf{\Delta}_p$ is represented as a C^*-subalgebra in $\mathcal{L}(\mathcal{H}_p)$ for each connected component p of the graph G.

The set $\mathbf{\Delta}_G$ forms the replacement of $\mathbf{\Delta}$ in (1). Its elements are in general non-square, hence it takes care of the limitation mentioned in (L2). That the limitation of (L1) is also taken care of will become clear from the examples discussed in Section III.

Since the elements of $\mathbf{\Delta}_G$ can be non-square, we cannot work with its commutant, like we did with $\mathbf{\Delta}$ in (3). Instead we will make use of the intertwining space

$$\mathbf{\Delta}'_G = \{(X,Y) \in \mathcal{L}(\mathcal{H}_R)\times\mathcal{L}(\mathcal{H}_S) \colon ZX = YZ,\ Z \in \mathbf{\Delta}_G\}. \tag{9}$$

In this case, one can prove that the operators X and Y in a pair $(X,Y) \in \mathbf{\Delta}'_G$ take the form

$$X = \sum_{r\in R} \iota_r\Gamma_{[r]}\iota_r^* \quad \text{and} \quad Y = \sum_{s\in S} \iota_s\Gamma_{[s]}\iota_s^*,$$

where for each $p \in P$ the operator Γ_p in in $\mathbf{\Delta}'_p$. Moreover, any pair (X,Y) with X and Y of this form is in $\mathbf{\Delta}'_G$.

C. The structured singular value and its upper bound

With $\mathbf{\Delta}_G \subset \mathcal{L}(\mathcal{H}_R, \mathcal{H}_S)$ the structure defined in the preceding subsection, we define the *μ_G-structured singular value* of an operator of $M \in \mathcal{L}(\mathcal{H}_S, \mathcal{H}_R)$ as in (2) but with $\mathbf{\Delta}_G$ as in (7) in place of $\mathbf{\Delta}$:

$$\mu_G(M) = \frac{1}{\min\{\|\Delta\| \colon \Delta \in \mathbf{\Delta}_G,\ 1 \in \sigma(M\Delta)\}}. \tag{10}$$

The analogue of the D-scaled version of μ is defined as

$$\widehat{\mu}_{\mathbf{\Delta}}(M) = \inf\{\|XMY^{-1}\| \colon (X,Y) \in \mathbf{\Delta}'_G\}.$$

As in the classical case, $\widehat{\mu}_{\mathbf{\Delta}}(M)$ has the following properties:

- Computation of $\widehat{\mu}_G(M)$ can be reduced to a C^*-algebra LMI computation: $\widehat{\mu}_G(M) < 1$ *if and only if there exists a positive definite structured solution* $(X,Y) \in \mathbf{\Delta}'_G$ *of the structured Stein equation*

$$M^*XM - Y \prec 0. \tag{11}$$

- $\widehat{\mu}_G(M)$ is always an upper bound for $\mu_G(M)$:

$$\mu_G(M) \le \widehat{\mu}_G(M).$$

III. EXAMPLES: NON SQUARE AND HIGHER-MULTIPLICITY BLOCKS

In this section we specify the graph-formalism structure introduced in the previous paragraph to two special cases. The first special case coincides with the structure $\mathbf{\Delta}$ given in (1).

Example 3.1: In this example we index the connected components of our graph G by the first $K+F$ natural numbers $\{i \colon 1 \le i \le S+F\}$. We take

$$R_k = \{r_k\} \quad \text{and} \quad S_k = \{s_k\} \quad \text{for } k = 1,\dots,K$$

while for $1 \le f \le F$ we set

$$R_{K+f} = \{r_{K+f,\alpha} \colon 1 \le \alpha \le m_f\},$$
$$S_{K+f} = \{s_{K+f,\alpha} \colon 1 \le \alpha \le m_f\},$$

where $m_1,\dots m_F$ are preset positive integers. We then declare $[s_i] = [r_i] = i$ for $1 \le i \le S$ while $[s_{K+j,\alpha}] = [r_{K+j,\alpha}] = K+j$ for $1 \le j \le F$. Thus the connected components G_l of our bipartite graph G are as follows: for $1 \le k \le K$, G_k consists of a single edge connecting the source vertex s_k to the range vertex r_k. For $1 \le f \le F$, the graph G_{K+f} consists of all possible edges $e_{\alpha,\beta}$ connecting a source vertex of the form $s_{K+f,\alpha}$ with a range vertex of the form $r_{K+f,\beta}$, where α and β run from 1 to m_f. The Hilbert space $\mathcal{H}_{G_k}$ we take to be $\mathbb{C}^{n_k}$ for $1 \le k \le S$, again here $n_1,\dots,n_K$ are preset positive integers, and $\mathcal{H}_{G_{K+f}} = \mathbb{C}$ for $1 \le f \le F$. Then it is easily checked that the structure (7) reduces to the standard structure (1) used for the block-uncertainty analysis in robust control theory (see [23]). Moreover, in this square case (where $\mathcal{H}_{R,i} = \mathcal{H}_{S,i}$ for each $i = 1,\dots,K+F$), one can see that the intertwining space $\mathbf{\Delta}'_G$ given by (9) reduces to the image of the commutant $\mathbf{\Delta}'$ (3) under the diagonal embedding: $\Delta \mapsto (\Delta,\Delta)$.

The reader should note that, with other choices of spaces $\mathcal{H}_p$ and C^*-subalgebras Δ_p, one can remedy the deficiencies mentioned at the end of Section I, as is illustrated in the following more general example.

Example 3.2: In this example we take the graph G at the level of generality considered in Section II, but we specify the C^*-algebras to be of the form

$$\mathbf{\Delta}_p = I_{\mathcal{H}_p^\circ} \otimes \mathcal{L}(\widetilde{\mathcal{H}}_p) \quad (p \in P).$$

Hence we have $\mathbf{\Delta}_p \subset \mathcal{L}(\mathcal{H}_p)$ with $\mathcal{H}_p$ factored in the form $\mathcal{H}_p = \mathcal{H}_p^\circ \otimes \widetilde{\mathcal{H}}_p$ with $\mathcal{H}_p^\circ$ and $\widetilde{\mathcal{H}}_p$ separable Hilbert spaces. Set $\widetilde{\mathcal{H}}_v = \widetilde{\mathcal{H}}_{[v]}$ for each $v \in V$ and

$$\widetilde{\mathcal{H}}_{S_p} = \oplus_{s\in S_p}\widetilde{\mathcal{H}}_s, \quad \widetilde{\mathcal{H}}_{R_p} = \oplus_{r\in R_p}\widetilde{\mathcal{H}}_r \quad (p \in P).$$

Since $(\mathcal{X} \otimes \mathcal{Y}_1) \oplus (\mathcal{X} \otimes \mathcal{Y}_2) = \mathcal{X} \otimes (\mathcal{Y}_1 \oplus \mathcal{Y}_2)$ holds for arbitrary Hilbert spaces $\mathcal{X}$, $\mathcal{Y}_1$ and $\mathcal{Y}_2$, we obtain that $\mathcal{H}_{S_p}$ and $\mathcal{H}_{R_p}$ can be written as

$$\mathcal{H}_{S_p} = \mathcal{H}_p^\circ \otimes \widetilde{\mathcal{H}}_{S_p} \quad \text{and} \quad \mathcal{H}_{R_p} = \mathcal{H}_p^\circ \otimes \widetilde{\mathcal{H}}_{R_p}.$$

Hence in this case we have

$$\mathcal{H}_S = \bigoplus_{p\in P}(\mathcal{H}_p^\circ \otimes \widetilde{\mathcal{H}}_{S_p}) \quad \text{and} \quad \mathcal{H}_R = \bigoplus_{p\in P}(\mathcal{H}_p^\circ \otimes \widetilde{\mathcal{H}}_{R_p}).$$

Then the structure space $\mathbf{\Delta}_G$ has the explicit form

$$\mathbf{\Delta}_G = \oplus_{p\in P}(I_{\mathcal{H}_p^\circ} \otimes \mathcal{L}(\widetilde{\mathcal{H}}_{R_p}, \widetilde{\mathcal{H}}_{S_p})),$$

and a pair (X,Y) with $X \in \mathcal{L}(\mathcal{H}_R)$ and $Y \in \mathcal{L}(\mathcal{H}_S)$ is the intertwining space $\mathbf{\Delta}'_G$ if and only if

$$X = \oplus_{p\in P}(\Gamma_p^\circ \otimes I_{\widetilde{\mathcal{H}}_{R_p}}) \quad \text{and} \quad Y = \oplus_{p\in P}(\Gamma_p^\circ \otimes I_{\widetilde{\mathcal{H}}_{S_p}}),$$

where for each $p \in P$ we have $\Gamma_p^\circ \in \mathcal{L}(\mathcal{H}_p^\circ)$.

Remark 3.3: The structured graph and spaces considered in [4], [5], [6] amount to the special case of Example 2 above where one takes $\widetilde{\mathcal{H}}_p = \mathbb{C}$ for all p so that $\mathbf{\Delta}_p = \{\lambda I_{\mathcal{H}_p} \colon \lambda \in \mathbb{C}\}$. If one considers the case of Example 3.2 with all spaces $\mathcal{H}_p$ finite dimensional, as we shall do in the next section, then the set-up of Example 2 can be reduced to the special case of [4], [5], [6] by adjusting the choice of the graph G. This enhanced structure is enough to resolve the limitations (L1) and (L2) of the standard (K, F) structure discussed at end of the introduction.

IV. THE MAIN RESULT

In this section we discuss the special case of $\mathbf{\Delta}_G$ discussed in Example 3.2. We assume in addition that the Hilbert spaces $\mathcal{H}_p$ are all finite dimensional. Let $M \in \mathcal{L}(\mathcal{H}_S, \mathcal{H}_R)$. Just as in the classical case, we can define a relaxed version $\widetilde{\mu}_G(M)$ of $\mu_G(M)$ by tensoring with operators on ℓ^2 as follows. We define $\widetilde{\mathbf{\Delta}}_G$ by

$$\begin{aligned}\widetilde{\mathbf{\Delta}}_G &= \bigoplus_{p\in P} \left(\mathcal{L}(\ell^2)\otimes I_{\mathcal{H}_p^\circ}\otimes \mathcal{L}(\widetilde{\mathcal{H}}_{R_p},\widetilde{\mathcal{H}}_{S_p})\right)\\ &= \mathcal{L}(\ell^2)\otimes\left(\bigoplus_{p\in P}\left(I_{\mathcal{H}_p^\circ}\otimes\mathcal{L}(\widetilde{\mathcal{H}}_{R_p},\widetilde{\mathcal{H}}_{S_p})\right)\right)\\ &= \mathcal{L}(\ell^2)\otimes\mathbf{\Delta}_G \subset \mathcal{L}(\ell^2\otimes\mathcal{H}_R,\ell^2\otimes\mathcal{H}_S).\end{aligned}$$

We then set $\widetilde{M} = I_{\ell^2}\otimes M \in \mathcal{L}(\ell^2\otimes\mathcal{H}_R, \ell^2\otimes\mathcal{H}_L)$ and define

$$\widetilde{\mu}_{\mathbf{\Delta}_G}(M) := \mu_{\widetilde{\mathbf{\Delta}}_G}(\widetilde{M}).$$

One can check that a pair $(\widetilde{X},\widetilde{Y}) \in \mathcal{L}(\ell^2\otimes\mathcal{H}_R)\times\mathcal{L}(\ell^2\otimes\mathcal{H}_S)$ is in $\widetilde{\delta}'_G$ if and only if

$$\begin{aligned}\widetilde{X} &= I_{\ell^2}\otimes\left(\oplus_{p\in P}(\Gamma_p^\circ\otimes I_{\widetilde{\mathcal{H}}_{R_p}})\right),\\ \widetilde{Y} &= I_{\ell^2}\otimes\left(\oplus_{p\in P}(\Gamma_p^\circ\otimes I_{\widetilde{\mathcal{H}}_{S_p}})\right),\end{aligned}$$

with $\Gamma_p^\circ\in\mathcal{L}(\mathcal{H}_p^\circ)$ for each $p\in P$. Hence $(\widetilde{X},\widetilde{Y})\in\widetilde{\Delta}'_G$ if and only if there exists a pair $(X,Y)\in\Delta'_G$ such that

$$\widetilde{X} = I_{\ell^2}\otimes X \quad\text{and}\quad \widetilde{Y} = I_{\ell^2}\otimes Y$$

From this fact it easily follows that

$$\widehat{\mu}_{\widetilde{\Delta}_G}(\widetilde{M}) = \widehat{\mu}_{\Delta_G}(M).$$

Hence $\widehat{\mu}_{\widetilde{\Delta}_G}(\widetilde{M}) < 1$ is characterized by the same (finite dimensional) LMI condition (11) as $\widehat{\mu}_{\Delta_G}(M) < 1$. Our main result is the following extension of the equality (5) to this more general setting.

Theorem 4.1: With assumptions and notation as in the preceding discussion in force, the relaxed μ-singular value $\widetilde{\mu}_{\mathbf{\Delta}_G}(M)$ is equal to its upper bound $\widehat{\mu}_{\Delta_G}(M)$:

$$\widetilde{\mu}_{\mathbf{\Delta}_G}(M) = \widehat{\mu}_{\mathbf{\Delta}_G}(M).$$

Complete details on this result will be forthcoming in [7]. Let us include here only some comments on the proof. Actually there are at least two distinct approaches.

A. First proof of Theorem 4.1

The first approach is based on the noncommutative multivariable Bounded Real Lemma from [6]. To state the result we first need some preliminaries. We suppose that G, $\mathcal{H}_p = I_{\mathcal{H}_p^\circ}\otimes\widetilde{\mathcal{H}}_p$, $\mathbf{\Delta}_p = \mathcal{H}_p^\circ\otimes\mathcal{L}(\widetilde{\mathcal{H}}_p)$ for $p\in P$ are as in Example 2 above. We suppose that we are given a colligation matrix $\mathbf{U}$ of the form

$$\mathbf{U} = \begin{bmatrix} A & B\\ C & D\end{bmatrix} \colon \begin{bmatrix}\mathcal{X}\\ \mathcal{U}\end{bmatrix}\to\begin{bmatrix}\mathcal{X}\\ \mathcal{Y}\end{bmatrix}.$$

We shall be interested in the induced operator $\widetilde{\mathbf{U}} := I_{\ell^2}\otimes\mathbf{U} = \left[\begin{smallmatrix}\widetilde{A} & \widetilde{B}\\ \widetilde{C} & \widetilde{D}\end{smallmatrix}\right]$ given by

$$\widetilde{\mathbf{U}} = \begin{bmatrix} I_{\ell^2}\otimes A & I_{\ell^2}\otimes B\\ I_{\ell^2}\otimes C & I_{\ell^2}\otimes D\end{bmatrix} \colon \begin{bmatrix}\ell^2\otimes\mathcal{X}\\ \ell^2\otimes\mathcal{U}\end{bmatrix}\to\begin{bmatrix}\ell^2\otimes\mathcal{X}\\ \ell^2\otimes\mathcal{Y}\end{bmatrix}.$$

and the associated linear-fractional map $\mathcal{F}_{\widetilde{\mathbf{U}}}$:

$$\mathcal{F}_{\widetilde{\mathbf{U}}}[\widetilde{\Delta}] = \widetilde{D} + \widetilde{C}(I - \widetilde{\Delta}\widetilde{A})^{-1}\widetilde{\Delta}\widetilde{B}$$

defined for $\widetilde{\Delta}\in\widetilde{\mathbf{\Delta}}_G$ whenever the inverse exists.

The key result from [6] can be stated as follows.

Proposition 4.2 (Bounded Real Lemma, see [6]): Let G, $\mathcal{H}_p = \mathcal{H}_p^\circ\otimes\widetilde{\mathcal{H}}_p$, $\mathbf{\Delta}_p = I_{\mathcal{H}_p^\circ}\otimes\mathcal{L}(\widetilde{\mathcal{H}}_p)$, for $p\in P$, be as in Example 2 above. Suppose that the colligation matrix

$$\mathbf{U} = \begin{bmatrix} A & B\\ C & D\end{bmatrix} \colon \begin{bmatrix}\mathcal{H}_S\\ \mathcal{U}\end{bmatrix}\to\begin{bmatrix}\mathcal{H}_R\\ \mathcal{Y}\end{bmatrix}$$

has the property that $\|\mathcal{F}_{\widetilde{\mathbf{U}}}[\widetilde{\Delta}]\| < 1$ for all $\widetilde{\Delta}\in\widetilde{\mathbf{\Delta}}_G$ having $\|\widetilde{\Delta}\|\le 1$. Then there exist $(X,Y)\in\mathbf{\Delta}'_G$ so that

$$\begin{bmatrix} A & B\\ C & D\end{bmatrix}^*\begin{bmatrix} X & 0\\ 0 & I_{\mathcal{Y}}\end{bmatrix}\begin{bmatrix} A & B\\ C & D\end{bmatrix} - \begin{bmatrix} Y & 0\\ 0 & I_{\mathcal{Y}}\end{bmatrix}\prec 0.$$

Proposition 4.2 can be applied to arrive at a proof of Theorem 4.1 as follows. Since it is always the case that $\widetilde{\mu}_{\mathbf{\Delta}}(M)\le\widehat{\mu}_{\mathbf{\Delta}}(M)$, the proof of Theorem 4.1 is complete once we show:

$$\widetilde{\mu}_{\mathbf{\Delta}}(M) < 1 \Rightarrow \widehat{\mu}_{\mathbf{\Delta}}(M) < 1. \tag{12}$$

The hypothesis in (12) translates to the statement

$$(I - \widetilde{\Delta}M)^{-1}\text{ exists for all }\widetilde{\Delta}\in\widetilde{\mathbf{\Delta}}_G\text{ with }\|\widetilde{\Delta}\|\le 1. \tag{13}$$

A tricky shifting technique due to Megretski-Treil [17] (see also the proof of Proposition B.1 in [11]) enable one to conclude that condition (13) actually implies a uniform bound on the resolvent:

$$\textit{There exists } K<\infty \textit{ so that } \|(I-\widetilde{\Delta}M)^{-1}\|\le K \;\textit{for all } \widetilde{\Delta}\in\widetilde{\mathbf{\Delta}}_G \textit{ with } \|\widetilde{\Delta}\|\le 1. \tag{14}$$

It follows that we can choose $\epsilon > 0$ $(0<\epsilon<1/K)$ so that $\|\epsilon(I-\widetilde{\Delta}M)^{-1}\| < 1$ for all $\widetilde{\Delta}\in\widetilde{\mathbf{\Delta}}_G$ with $\|\widetilde{\Delta}\|\le 1$. Note that we can write $\epsilon(I-\Delta M)^{-1}$ in the form $\mathcal{F}_{\widetilde{\mathbf{U}}}[\widetilde{\Delta}]$ if we choose $\mathbf{U} = \left[\begin{smallmatrix} M & M\\ \epsilon I_{\mathcal{H}_S} & \epsilon I_{\mathcal{H}_S}\end{smallmatrix}\right]$. We now apply Proposition 4.2 to this $\mathbf{U}$ to conclude that there is a $(X,Y)\in\mathbf{\Delta}'_G$ so that

$$\begin{bmatrix} M & M\\ \epsilon I_{\mathcal{H}_S} & \epsilon I_{\mathcal{H}_S}\end{bmatrix}^*\begin{bmatrix} X & 0\\ 0 & I\end{bmatrix}\begin{bmatrix} M & M\\ \epsilon I_{\mathcal{H}_S} & \epsilon I_{\mathcal{H}_S}\end{bmatrix} - \begin{bmatrix} Y & 0\\ 0 & I\end{bmatrix}\prec 0.$$

In particular, from the (1,1)-block entry we see that $M^*XM + \epsilon^2 I - Y \prec 0$ and (12) follows.

These connections between the Bounded Real Lemma and Theorem 4.1 were already observed in [6]; however the results there were incomplete due to unawareness of the Megretski-Treil argument for proceeding from (13) to (14).

B. Second proof of Theorem 4.1

The second approach is to adapt the proof for the standard structure laid out in the book [11] for the more general higher-multiplicity setting considered in Example 3.2 above. The elegant cone-separation argument from [11] applies virtually without change to the higher-multiplicity situation. However, there is one extra ingredient required which we describe first.

We recall the well-known Douglas lemmas (see [9]).

Lemma 4.3: Given Hilbert space operators $A \in \mathcal{L}(\mathcal{Y}, \mathcal{Z})$ and $B \in \mathcal{L}(\mathcal{X}, \mathcal{Z})$, there exists an operator $X \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ with $AX = B$ and $\|X\| \leq 1$ if and only if $BB^* - AA^* \preceq 0$.

It turns out that the required additional ingredient for this second proof of Theorem 4.1 following the approach of [11] is a structured version of the Douglas lemma given in Lemma 4.4 below.

We first review some preliminaries. We follow the notation as in Example 3.2 above. Fix a connected component $p \in P$ of the graph G. For notational convenience we set

$$\begin{aligned} &\mathcal{K}_{R_p} = \ell^2 \otimes \mathcal{H}_{R_p}, \quad \mathcal{K}_{S_p} = \ell^2 \otimes \mathcal{H}_{S_p}, \\ &\widetilde{\mathcal{K}}_{R_p} = \ell^2 \otimes \widetilde{\mathcal{H}}_{R_p} \quad \text{and} \quad \widetilde{\mathcal{K}}_{S_p} = \ell^2 \otimes \widetilde{\mathcal{H}}_{S_p}. \end{aligned}$$

Note that we can then identify $\mathcal{K}_{R_p}$ and $\mathcal{K}_{S_p}$ with $\mathcal{H}_p^\circ \otimes \widetilde{\mathcal{K}}_{R_p}$ and $\mathcal{H}_p^\circ \otimes \widetilde{\mathcal{K}}_{R_p}$, respectively.

There is a natural identification between the tensor product space $\mathcal{K}_{R_p} = \mathcal{H}_p^\circ \otimes \widetilde{\mathcal{K}}_{R_p}$ and the space $\mathcal{C}^2(\widetilde{\mathcal{K}}_{R_p}, \mathcal{H}_p^\circ)$ of Hilbert-Schmidt operators from $\widetilde{\mathcal{K}}_{R_p}$ to $\mathcal{H}_p^\circ$. On a pure tensor $k_p^\circ \otimes \widetilde{k}_{R_p}$, the identification is given by

$$V \colon k_p^\circ \otimes \widetilde{k}_{R_p} \mapsto k_p^\circ (\widetilde{k}_{R_p})^\top$$

where $(\widetilde{k}_{R_p})^\top$ is the *transpose* of $\widetilde{k}_{R_p}$ and is viewed as a linear functional in $\mathcal{L}(\widetilde{\mathcal{K}}_{R_p}, \mathbb{C})$; here k_p° is viewed as a linear operator $k_p^\circ \in \mathcal{L}(\mathbb{C}, \mathcal{H}_p^\circ)$ so that the composite $k_0^\circ(\widetilde{k}_{R_p})^\top$ is a rank 1 operator from $\widetilde{\mathcal{K}}_{R_p}$ into $\mathcal{H}_p^\circ$. Note that $(\widetilde{k}_{R_p})^\top$ can be defined as

$$(\widetilde{k}_{R_p})^\top = \left(\mathcal{C}(\widetilde{k}_{R_p})\right)^*$$

where $\mathcal{C}$ is a fixed choice of conjugation operator on $\widetilde{\mathcal{K}}_{R_p}$. One can check that V extends via linearity and continuity to a unitary isomorphism from $\mathcal{H}_p^\circ \otimes \widetilde{\mathcal{K}}_{R_p}$ onto $\mathcal{C}^2(\widetilde{\mathcal{K}}_{R_p}, \mathcal{H}_p^\circ)$. Exactly the same story goes through of course with $\mathcal{K}_{S_p} = \mathcal{H}_p^\circ \otimes \widetilde{\mathcal{K}}_{S_p}$ in place of $\mathcal{K}_{R_p} = \mathcal{H}_p^\circ \otimes \widetilde{\mathcal{K}}_{R_p}$. We use the same notation V for either case. Also, we define $V_{k_p} \in \mathcal{C}^2(\mathcal{K}_{R_p}, \mathcal{H}_p^\circ)$ to denote the image of the vector $k_p \in \mathcal{K}_{R_p}$ under V, and similarly we define $V_{h_p} \in \mathcal{C}^2(\mathcal{K}_{S_p}, \mathcal{H}_p^\circ)$ for $h_p \in \mathcal{K}_{S_p}$. We can now state the new structured version of the Douglas lemma; this can be taken as the replacement for Lemma 8.4 in [11].

Lemma 4.4: Given vectors $k \in \ell^2 \otimes \mathcal{H}_R = \oplus_{p \in P} \mathcal{K}_{R_p}$ and $h \in \ell^2 \otimes \mathcal{H}_S = \oplus_{p \in P} \mathcal{K}_{S_p}$, there exists an operator $\Delta \in \widetilde{\boldsymbol{\Delta}}_G$ such that

$$\Delta k = h \quad \text{and} \quad \|\Delta\| \leq 1$$

if and only if the inequality

$$V_{P_{\mathcal{K}_{R_p}} k} \left(V_{P_{\mathcal{K}_{R_p}} k}\right)^* - V_{P_{\mathcal{K}_{S_p}} h} (V_{P_{\mathcal{K}_{S_p}} h})^* \succeq 0$$

holds for all $p \in P$. Here $P_{\mathcal{K}_{R_p}}$ denotes the orthogonal projection of $\ell^2 \otimes \mathcal{H}_R$ onto $\mathcal{K}_{R_p}$ and similarly $P_{\mathcal{K}_{S_p}}$ denotes the orthogonal projection of $\ell^2 \otimes \mathcal{H}_S$ onto $\mathcal{K}_{S_p}$.

We remark that the key ingredient in the proof of this structured Douglas lemma are the identities

$$\begin{aligned} V_{(I_{\mathcal{H}_p^\circ} \otimes \widetilde{\Delta}_p) k_p} &= V_{k_p} (\widetilde{\Delta}_p)^\top \quad (k_p \in \mathcal{K}_{R_p},\, \widetilde{\Delta}_p \in \mathcal{L}(\widetilde{\mathcal{K}}_{R_p})), \\ V_{(I_{\mathcal{H}_p^\circ} \otimes \widetilde{\Delta}_p) h_p} &= V_{h_p} (\widetilde{\Delta}_p)^\top \quad (h_p \in \mathcal{K}_{S_p},\, \widetilde{\Delta}_p \in \mathcal{L}(\widetilde{\mathcal{K}}_{S_p})). \end{aligned}$$

In this way, the existence of a contraction with a repeated block structure satisfying a Douglas lemma constraint can be reduced to the existence of an non-structured contraction.

Now, given $M \in \mathcal{L}(\mathcal{H}_S, \mathcal{H}_R)$ we set $\widetilde{M} = I_{\ell^2} \otimes M$ and we define for each $p \in P$ the function $\phi_p : \ell^2 \otimes \mathcal{H}_S \to \mathcal{L}(\mathcal{H}_p)$ by

$$\phi_p(h) = V_{P_{\mathcal{K}_{R_p}} \widetilde{M} h} V^*_{P_{\mathcal{K}_{R_p}} \widetilde{M} h} - V_{P_{\mathcal{K}_{S_p}} h} V^*_{P_{\mathcal{K}_{S_p}} h} \;\; (h \in \ell^2 \otimes \mathcal{H}_S).$$

In addition we define the sets

$$\nabla = \{(\phi_p(h))_{p \in P} :\ h \in \ell^2 \otimes \mathcal{H}_S\ \|h\| = 1\}, \tag{15}$$

$$\Pi = \{(L_p)_{p \in P} :\ L_p \in \mathcal{L}(\mathcal{H}_p^\circ),\ L_p \succ 0,\ p \in P\}. \tag{16}$$

The structured Douglas lemma now makes it possible to prove the following result.

Proposition 4.5: Assume G, $\mathcal{H}_p = \mathcal{H}_p^\circ \otimes \widetilde{\mathcal{H}}_p$, $\boldsymbol{\Delta}_p = I_{\mathcal{H}_p^\circ} \otimes \mathcal{L}(\widetilde{\mathcal{H}}_p)$, for $p \in P$, are as in Example 2 above with $\mathcal{H}_S$ and $\mathcal{H}_R$ finite dimensional. Let $M \in \mathcal{L}(\mathcal{H}_S, \mathcal{H}_R)$. Then the following are equivalent:

(i) $\widehat{\mu}_G(M) < 1$;

(ii) *There exists a pair $(X, Y) \in \boldsymbol{\Delta}'_G$ with $X \succ 0$ and $Y \succ 0$ such that*

$$M^* X M - Y \prec 0;$$

(iii) *There exist positive definite $\Gamma_p \in \mathcal{L}(\mathcal{H}_p)$, $p \in P$, and $\alpha, \beta \in \mathbb{R}$ such that*

$$\sum_{p \in P} \text{trace}\,(\Gamma_p K_p) \leq \alpha < \beta \leq \sum_{p \in P} \text{trace}\,(\Gamma_p L_p)$$

for all $(K_p)_{p \in P} \in \nabla$ and $(L_p)_{p \in P} \in \Pi$.

Hence (iii) shows that $\widehat{\mu}_G(M) < 1$ is equivalent to strict separation of the sets ∇ and Π by a hyperplane in the direct sum of $\#(P)$ copies of $\mathcal{L}(\mathcal{H}_p^\circ)$, where $\#(P)$ denotes the number of elements in P. It is not hard to show that $\widetilde{\mu}_G(M) < 1$ implies that $\nabla \cap \Pi = \emptyset$. The remainder of the proof involves showing that if $\widetilde{\mu}_G(M) < 1$ holds, then ∇ and Π can be strictly separated by a hyperplane; see the comments in the next section.

V. CONCLUSIONS

We conclude with some general remarks and observations.

- The proof of the Bounded Real Lemma from [6] (see Proposition 4.2 above) used in the First Proof of Theorem 4.1 relies on the realization result for the noncommutative Schur-Agler class from [5] together with the State-Space Similarity Theorem for structured noncommutative multidimensional linear systems from [4]. The proof of the realization result for structured noncommutative multidimensional systems in turn relies on an infinite-dimensional cone-separation argument modeled on the argument for the commutative case due to Agler [1]. On the other hand, the cone-separation part of the argument borrowed from [11] for the Second Proof is a completely elementary finite-dimensional cone-separation structure. The nontrivial part of the argument involves two main ingredients:

  (i) to show that the set ∇ defined in (15) has a convex closure, and

  (ii) to show that the closure of ∇ is strictly separated from the second convex set Π defined in (16) so as to be able to apply the Hahn-Banach cone-separation theorem.

  Verifications of (i) and (ii) both involve tricky arguments using shift-invariance of the relevant operators adapting arguments used in slightly different contexts from [17] and [21] respectively.

- The argument in the First Proof of Theorem 4.1 can be reversed: assuming that one knows the realization theorem for affine structured noncommutative multidimensional linear systems as in [4] (results of this type go back to the work of Fliess [12] in the 1970s), one can arrive at the realization theorem for the noncommutative Schur-Agler class as an application of Theorem 4.1 ($\widetilde{\mu} = \widehat{\mu}$) and a simple application of the Main Loop Theorem [23, Theorem 11.7]. In this way one can trace simultaneous independent origins of the Schur-Agler theory in the engineering and the mathematical communities.

- One of the results from [3] is that the noncommutative Schur-Agler class remains the same if one replaces the d-tuple of operators $\widetilde{\delta}_j$ $(1 \leq j \leq d)$ on ℓ^2 with a d-tuple of $n \times n$ matrices, and let n run over all all possible finite sizes $n = 1, 2, 3, \dots$. This is the starting point for noncommutative function theory which is currently an active area of research (see in particular [13], [14], [15], [16]). In particular the linear pencils studied in [13], [14] are more flexible versions of the maps $\widetilde{\delta} \mapsto Z_G(\widetilde{\delta})$ as in (8) appearing in [4], [5], [6]. The general theory of noncommutative functions (see [16]) starts with more general axioms to be satisfied by functions of tuples of $n \times n$ matrices involving compatibility conditions as one moves between different values of n. One then arrives at local power series representations as a result. Interesting recent work of Agler and McCarthy [2] arrives at the same type of results as in [5], but starting with the axiomatic noncommutative function theory framework rather than assuming power series representations at the start as in [5].

REFERENCES

[1] J. Agler, On the representation of certain holomorphic functions defined on a polydisk, in: *Topics in Operator Theory: Ernst D. Hellinger Memorial Volume*, pp. 47–66, Oper. Theory Adv. Appl. **48**, Birkhäuser-Verlag, Basel, 1990.

[2] J. Agler and J.E. McCarthy, Global holomorphic functions in several non-commuting variables, arXiv:1305.1636v1.

[3] D. Alpay, D. Kalyuzhnyi-Verbovetzkiıı, Matrix-J-unitary non-commutative rational formal power series, In: *The State Space Method: Generalizations and Applications*, pp. 49–113, Oper. Theory Adv. Appl. **161**, Birkhäuser, Basel-Berlin-Boston, 2006.

[4] J.A. Ball, G. Groenewald and T. Malakorn, Structured noncommutative multidimensional linear systems, *SIAM J. Control Optim.* **44** (2005) No. 4, 1474–1528.

[5] J.A. Ball, G. Groenewald and T. Malakorn, Conservative structured noncommutative multidimensional linear systems, in: *The State Space Method Generalizations and Applications*, pp. 179–223, Oper. Theory Adv. Appl. **161**, Birkhäuser, Basel-Berlin-Boston, 2006.

[6] J.A. Ball, G. Groenewald and T. Malakorn, Bounded real lemma for structured noncommutative multidimensional linear systems and robust control, *Multidimens. Sys. Signal Process.* **17** (2006), 119–150.

[7] J.A. Ball, G. Groenewald, and S. ter Horst, Bounded Real Lemma and structured singular value versus diagonal scaling: the freely noncommutative setting, in preparation.

[8] R. Braatz, P. Young, J.C. Doyle, and M. Morari, Computational complexity of μ calculations, *IEEE Transactions on Automatic Control* **39** (1994), 1000–1002.

[9] R.G. Douglas, On majorization, factorization, and range inclusion of operators on Hilbert space, *Proc. Amer. Math. Soc.* **17** (1966), 413–415.

[10] J.C. Doyle, Analysis of feedback systems with structured uncertainties, *IEEE Proceedings* **129** (1982), 242–250.

[11] G.E. Dullerud and F. Paganini, *A Course in Robust Control Theory: A Convex Approach*, Texts in Applied Mathematics Vol. **36**, Springer-Verlag, New York, 2000.

[12] M. Fliess, Matrices de hankel, *J. Math. Pures. Appl.* **53** (1974), 197–222.

[13] J.W. Helton, I. Klep and S. McCullough, Noncommutative ball maps, *J. Functional Analysis* **257** (2009), 47–87.

[14] J.W. Helton, I. Klep and S. McCullough, Analytic mappings between noncommutative pencil balls, *J. Math. Anal. Appl.* **376** (2011), 407–428.

[15] J.W. Helton, I. Klep and S. McCullough, Proper analytic free maps, *J. Functional Analysis* **260** (2011), 1476–1490.

[16] D.S. Kaliuzhnyi-Verbovetskyi and V. Vinnikov, *Foundations of Noncommutative Function Theory*, arXiv:1212:6345.

[17] A. Megretski and S. Treil, Power distribution inequalities in optimization and robustness of uncertain systems, *J. Mathematical Systems, Estimation, and Control* **3** (1993), 301–319.

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

[19] F. Paganini, *Sets and Constraints in the Analysis of Uncertain Systems*, Thesis submitted to California Institute of Technology, Pasadena, 1996.

[20] M.G. Safanov, *Stability Robustness of Multivariable Feedback Systems*, MIT Press, Cambridge, MA, 1980.

[21] J.S. Shamma, Robust stability with time-varying structured uncertainty, *IEEE Trans. Automat. Control* **39** (1994) No. 4, 714–724.

[22] S. Treil, The gap between the complex structured singular value μ and its upper bound is infinite, preprint.

[23] K. Zhou, J.C. Doyle and K. Glover, *Robust and Optimal Control*, Prentice-Hall, Upper Saddle River, NJ, 1996.