

SPARSE NONCOMMUTATIVE POLYNOMIAL OPTIMIZATION

IGOR KLEP, VICTOR MAGRON, AND JANEZ POVH

ABSTRACT. This article focuses on optimization of polynomials in noncommuting variables, while taking into account sparsity in the input data. A converging hierarchy of semidefinite relaxations for eigenvalue and trace optimization is provided. This hierarchy is a noncommutative analogue of results due to Lasserre [SIAM J. Optim. 17(3) (2006), pp. 822–843] and Waki et al. [SIAM J. Optim. 17(1) (2006), pp. 218–242]. The Gelfand-Naimark-Segal (GNS) construction is applied to extract optimizers if flatness and irreducibility conditions are satisfied. Among the main techniques used are amalgamation results from operator algebra. The theoretical results are utilized to compute lower bounds on minimal eigenvalue and trace of noncommutative polynomials from the literature.

1. INTRODUCTION

The goal of this article is to handle a specific class of sparse polynomial optimization problems with noncommuting variables (e.g., polynomials in matrices). Applications of interest include control theory and linear systems in engineering [SIG98], quantum theory and quantum information science [NPA08]. For example, in the latter context, noncommutative polynomial optimization provides upper bounds on the maximum violation level of Bell inequalities [PV09]. Further motivations relate to the generalized Lax conjecture [Lax58], with computational proof attempts relying on noncommutative sums of squares (in Clifford algebras) [NT14]. The problem of verifying noncommutative polynomial inequalities has also occurred in a conjecture formulated by Bessis, Moussa and Villani (BMV) in 1975 [BMV75], and restated by Lieb and Seiringer [LS04]. The initial conjecture boils down to verifying that the (univariate) polynomial $t \mapsto \operatorname{tr}(A + tB)^m$ has only nonnegative coefficients, for all positive semidefinite matrices A and B , and all $m \in \mathbb{N}$. Even though the BMV conjecture has been established by Stahl [Sta13] for all m , one can rely on computational proofs for a fixed value of m . Schweighofer and the first author derived a computational proof [KS08] of the conjecture for $m \leq 13$. Recently, noncommutative polynomial optimization has been used in [GdLL18] to study optimization problems related to bipartite quantum correlations, and in [GdLL19] to

2010 *Mathematics Subject Classification.* 90C22; 47N10; 13J10.

Key words and phrases. noncommutative polynomial; sparsity pattern; semialgebraic set; semidefinite programming; eigenvalue optimization; trace optimization; GNS construction.

IK was supported by the Slovenian Research Agency grants J1-8132, N1-0057 and P1-0222. Partially supported by the Marsden Fund Council of the Royal Society of New Zealand.

VM was supported by the FMJH Program PGM0 (EPICS project) and EDF, Thales, Orange et Criteo, as well as from the Tremplin ERC Stg Grant ANR-18-ERC2-0004-01 (T-COPS project).

JP was supported by the Slovenian Research Agency program P2-0256 and grants J1-8132, J1-8155, N1-0057 and N1-0071.

derive hierarchies of lower bounds for several matrix factorization ranks, including nonnegative rank, positive semidefinite rank as well as their symmetric analogues.

In the commutative setting, *polynomial optimization* focuses on minimizing or maximizing a polynomial over a *semialgebraic* set, that is, a set defined by a finite conjunction/disjunction of polynomial inequalities. In general, computing the exact solution of a polynomial optimization problem is an NP-hard problem [Lau09b]. In practice, one can at least compute an approximation of the solution by considering a relaxation of the problem. In the seminal 2001 paper [Las01], Lasserre introduced a nowadays famous hierarchy of relaxations, called the *moment-sums of squares hierarchy* allowing us to obtain a converging sequence of lower bounds for the minimum of a polynomial over a compact semialgebraic set. Each lower bound is computed by solving a semidefinite program (SDP). In SDP, one optimizes a linear function under a linear matrix inequality constraint. SDP itself appears in a wide range of applications (combinatorial optimization [GLV09], control theory [BEGFB94], matrix completion [Lau09a], etc.) and can be solved efficiently (up to a few thousand optimization variables) by freely available software, e.g. SeDuMi [Stu99], SDPT3 [TTT03], SDPA [YFK03] or Mosek [Mos]. For optimization problems involving n -variate polynomials of degree less than d , the size of the matrices involved at step $k \geq d$ of Lasserre’s hierarchy of SDP relaxations is proportional to $\binom{n+k}{n}$. Therefore, the size of the SDP problems arising from the hierarchy grows rapidly.

For unconstrained problems involving a large number of variables n , a remedy consists of reducing the size of the SDP matrices by discarding the monomials which never appear in the support of the SOS decompositions. This technique, based on a result by Reznick [Rez78], consists of computing the *Newton polytope* of the input polynomial (that is, the convex hull of the support of this polynomial) and selecting only monomials with support lying in half of this polytope. For constrained optimization, another workaround is based on exploiting a potential sparsity/symmetry pattern arising in the input polynomials. In [Las06] (see also [WKKM06] and the related **SparsePOP** solver [WKK⁺09]), the author derives a sparse version of Putinar’s representation [Put93] for polynomials positive on compact semialgebraic sets. This variant can be used for cases where the objective function can be written as a sum of polynomials, each of them involving a small number of variables. Sparse polynomial optimization techniques enable us to successfully handle various concrete applications. The frameworks [MCD17, Mag18], coming with the **Real2Float** software library, rely on these techniques to produce a hierarchy of upper bounds converging to the absolute roundoff error of a numerical program involving polynomial operations. In energy networks, it is now possible to compute the solution of large-scale power flow problems with up to thousand variables [Jos16]. In [TWLH19], the authors derive the sparse analogue of [HLS09] to obtain a hierarchy of upper bounds for the volume of large-scale semialgebraic sets. In the same spirit, the symmetry pattern of the polynomial optimization problem can be exploited [RTAL13]. More recent progress focused on the use of alternative hierarchies, including the so-called *bounded degree SOS hierarchy* (BSOS) [LTY17]. Here, one represents a positive polynomial as the sum of two terms: an SOS polynomial of a priori fixed degree, and a term lying in the set of *Krivine-Stengle* representations [Kri64], that is, a combination of positive linear cross-products of the polynomials involved in the set of constraints. The BSOS hierarchy can handle bigger instances than the standard moment-SOS hierarchy. In addition, sparsity

can be exploited in the same way as for the sparse SOS hierarchy, which allows us to tackle even larger problems [WLT18].

In the noncommutative context, a given noncommutative polynomial in n variables and degree d is positive semidefinite if and only if it decomposes as a *sum of hermitian squares* (SOHS) [Hel02, McC01]. In practice, an SOHS decomposition can be computed by solving an SDP of size $O(n^d)$, which is even larger than the size of the matrices involved in the commutative case. SOHS decompositions are also used for constrained optimization, either to minimize eigenvalues or traces of noncommutative polynomial objective functions, under noncommutative polynomial (in)equality constraints. The optimal value of such constrained problems can be approximated, as closely as desired, while relying on the noncommutative analogue of Lasserre’s hierarchy [PNA10, CKP12, BCKP13]. The `NCSOStools` [CKP11, BKP16] library can compute such approximations for optimization problems involving polynomials in noncommuting variables. By comparison with the commutative case, the size $O(n^k)$ of the SDP matrices at a given step k of the noncommutative hierarchy becomes intractable even faster, typically for $k, n \simeq 6$ on a standard laptop.

A remedy for unconstrained problems is to rely on the adequate noncommutative analogue of the standard Newton polytope method, which is called the *Newton chip method* (see, e.g. [BKP16, §2.3]) and can be further improved with the *augmented Newton chip method* (see, e.g., [BKP16, §2.4]), by removing certain terms which can never appear in an SOHS decomposition of a given input. As in the commutative case, the Newton polytope method cannot be applied for constrained problems. When one cannot go from step k to step $k + 1$ in the hierarchy because of the computational burden, one can always consider matrices indexed by all terms of degree k plus a fixed percentage of terms of degree $k + 1$. This is used for instance to compute tighter upper bounds for maximum violation levels of Bell inequalities [PV09]. Another trick, implemented in the `Ncpol2sdpa` library [Wit15], consists of exploiting simple equality constraints, such as “ $X^2 = Y$ ”, to derive substitution rules for variables involved in the SDP relaxations. Similar substitutions are performed in the commutative case by `Gloptipoly 3` [HLL09].

Apart from such heuristic procedures, there is, to the best of our knowledge, no general method to exploit additional structure, such as sparsity, of (un)constrained noncommutative polynomial optimization problems.

Contributions. We state and prove in Section 3 a sparse variant of the noncommutative version of Putinar’s Positivstellensatz, under the same sparsity pattern assumptions as the ones used in the commutative case [Las06, WKKM06]; these conditions are known as the *running intersection property* (RIP) in graph theory [FKMN01, NFF+03]. Our proof relies on amalgamation results for operator algebras. Then, we present in Section 4 the sparse Gelfand-Naimark-Segal (GNS) construction yielding representations for linear functionals positive w.r.t. sparsity. This allows us to extract minimizers, providing that flatness and irreducibility conditions hold. We rely on this sparse representation to design algorithms performing eigenvalue optimization (Section 5) and trace optimization of noncommutative sparse polynomials (Section 6), both in the unconstrained and constrained case. Along the way we exhibit an example showing that the Helton-McCullough [Hel02, McC01] Sum of Squares theorem (every positive nc polynomial is a sum of hermitian squares) fails in the sparse setting, see Example 5.2. Finally, we provide

in Section 7 experimental comparisons between the bounds given by the dense relaxations and the ones produced by our algorithms, currently implemented in the `NCSOStools` software library.

2. NOTATION AND DEFINITIONS

This section gives the basic definitions and introduces notation used in the rest of the article.

2.1. Noncommutative polynomials. Let us denote by $\mathbb{M}_n(\mathbb{R})$ (resp. \mathbb{S}_n) the space of all real (resp. symmetric) matrices of order n , and by \mathbb{S}_n^k the set of k -tuples $\underline{A} = (A_1, \dots, A_k)$ of symmetric matrices A_i of order n . Let \mathbf{I}_n stands for the identity matrix of order n . For a fixed $n \in \mathbb{N}$, we consider a finite alphabet X_1, \dots, X_n and generate all possible words of finite length in these letters. The empty word is denoted by 1. The resulting set of words is $\langle \underline{X} \rangle$, with $\underline{X} = (X_1, \dots, X_n)$. We denote by $\mathbb{R}\langle \underline{X} \rangle$ the set of real polynomials in noncommutative variables, abbreviated as *nc polynomials*. A *monomial* is an element of the form $a_w w$, with $a_w \in \mathbb{R} \setminus \{0\}$ and $w \in \langle \underline{X} \rangle$. The *degree* of an nc polynomial $f \in \mathbb{R}\langle \underline{X} \rangle$ is the length of the longest word involved in f . For $d \in \mathbb{N}$, $\langle \underline{X} \rangle_d$ is the set of all words of degree at most d . Let us denote by \mathbf{W}_d the vector of all words of $\langle \underline{X} \rangle_d$ w.r.t. to the lexicographic order. Note that the dimension of $\mathbb{R}\langle \underline{X} \rangle_d$ is the length of \mathbf{W}_d , which is $\sigma(n, d) := \sum_{i=0}^d n^i = \frac{n^{d+1}-1}{n-1}$. The set $\mathbb{R}\langle \underline{X} \rangle$ is equipped with the involution \star that fixes $\mathbb{R} \cup \{X_1, \dots, X_n\}$ point-wise and reverses words, so that $\mathbb{R}\langle \underline{X} \rangle$ is the \star -algebra freely generated by n symmetric letters X_1, \dots, X_n . The set of all *symmetric elements* is defined as $\text{Sym } \mathbb{R}\langle \underline{X} \rangle := \{f \in \mathbb{R}\langle \underline{X} \rangle : f = f^\star\}$.

Sums of hermitian squares. An nc polynomial of the form $g^\star g$ is called a *hermitian square*. A given $f \in \mathbb{R}\langle \underline{X} \rangle$ is a *sum of hermitian squares* (SOHS) if there exist nc polynomials $h_1, \dots, h_r \in \mathbb{R}\langle \underline{X} \rangle$, with $r \in \mathbb{N}$, such that $f = h_1^\star h_1 + \dots + h_r^\star h_r$. Let $\Sigma\langle \underline{X} \rangle$ stands for the set of SOHS. We denote by $\Sigma\langle \underline{X} \rangle_d \subseteq \Sigma\langle \underline{X} \rangle$ the set of SOHS polynomials of degree at most $2d$. We now recall how to check whether a given $f \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ is an SOHS. The existing procedure, known as the *Gram matrix method*, relies on the following proposition (see, e.g., [Hel02, §2.2]):

Proposition 2.1. *Assume that $f \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle_{2d}$. Then $f \in \Sigma\langle \underline{X} \rangle$ if and only if there exists $G_f \succeq 0$ satisfying*

$$(2.1) \quad f = \mathbf{W}_d^\star G_f \mathbf{W}_d.$$

Conversely, given such $G_f \succeq 0$ with rank r , one can construct $g_1, \dots, g_r \in \mathbb{R}\langle \underline{X} \rangle_d$ such that $f = \sum_{i=1}^r g_i^\star g_i$.

Any symmetric matrix G_f (not necessarily positive semidefinite) satisfying (2.1) is called a *Gram matrix* of f .

Semialgebraic sets and quadratic modules. Given a positive integer m and $S = \{s_1, \dots, s_m\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$, the *semialgebraic set* \mathcal{D}_S associated to S is defined as follows:

$$(2.2) \quad \mathcal{D}_S := \bigcup_{k \in \mathbb{N}} \{ \underline{A} = (A_1, \dots, A_n) \in \mathbb{S}_k^n : s_j(\underline{A}) \succeq 0, \quad j = 1, \dots, m \}.$$

When considering only tuples of $N \times N$ symmetric matrices, we use the notation $\mathcal{D}_S(N) := \mathcal{D}_S \cap \mathbb{S}_N^n$. The *operator semialgebraic set* \mathcal{D}_S^∞ is the set of all bounded self-adjoint operators \underline{A} on a Hilbert space making $g(\underline{A})$ a positive semidefinite operator, for all $g \in S$. The *quadratic module* M_S , generated by S , is defined by

$$(2.3) \quad M_S := \left\{ \sum_{i=1}^K a_i^* s'_i a_i : K \in \mathbb{N}, a_i \in \mathbb{R}\langle \underline{X} \rangle, s'_i \in S \cup \{1\} \right\}.$$

Given $d \in \mathbb{N}$, the truncated *quadratic module* $M_{S,d}$ of order d , generated by S , is

$$(2.4) \quad M_{S,d} := \left\{ \sum_{i=1}^K a_i^* s'_i a_i : K \in \mathbb{N}, a_i \in \mathbb{R}\langle \underline{X} \rangle, s'_i \in S \cup \{1\}, \deg(a_i^* s'_i a_i) \leq 2d \right\}.$$

A quadratic module M is called *archimedean* if for each $a \in \mathbb{R}\langle \underline{X} \rangle$, there exists $N \in \mathbb{N}$ such that $N - a^* a \in M$, e.g., $N - \sum_{i=1}^n X_i^2 \in M$. The noncommutative analog of Putinar's Positivstellensatz [Put93] describing noncommutative polynomials positive on \mathcal{D}_S^∞ with archimedean M_S is due to Helton and McCullough:

Theorem 2.2 ([HM04, Theorem 1.2]). *Let $S \cup \{f\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and assume that M_S is archimedean. If $f(\underline{A}) \succ 0$ for all $\underline{A} \in \mathcal{D}_S^\infty$, then $f \in M_S$.*

2.2. Sparsity patterns. Let $I_0 := \{1, \dots, n\}$. For $p \in \mathbb{N}$ consider $I_1, \dots, I_p \subseteq I_0$ satisfying $\bigcup_{k=1}^p I_k = I_0$. Let n_k be the size of I_k , for each $k = 1, \dots, p$.

We denote by $\langle \underline{X}(I_k) \rangle$ (resp. $\mathbb{R}\langle \underline{X}(I_k) \rangle$) the set of words (resp. nc polynomials) in the n_k variables $\underline{X}(I_k) = \{X_i : i \in I_k\}$. The dimension of $\mathbb{R}\langle \underline{X}(I_k) \rangle_d$ is $\sigma(n_k, d) = \frac{n_k^{d+1} - 1}{n_k - 1}$. Note that $\mathbb{R}\langle \underline{X}(I_0) \rangle = \mathbb{R}\langle \underline{X} \rangle$. We also define $\text{Sym } \mathbb{R}\langle \underline{X}(I_k) \rangle := \text{Sym } \mathbb{R}\langle \underline{X} \rangle \cap \mathbb{R}\langle \underline{X}(I_k) \rangle$, let $\Sigma\langle \underline{X}(I_k) \rangle$ stands for the set of SOHS in $\mathbb{R}\langle \underline{X}(I_k) \rangle$ and we denote by $\Sigma\langle \underline{X}(I_k) \rangle_d$ the restriction of $\Sigma\langle \underline{X}(I_k) \rangle$ to nc polynomials of degree at most $2d$. Let \mathcal{I}_k be the set of all subsets of I_k , for all $k = 1, \dots, p$. For $u \in \langle \underline{X} \rangle$, $\text{supp } u$ denotes the support of u , that is, the set of all $i \in \{1, \dots, n\}$ for which the letter X_i appears in the word u . As an example, if one considers $n = 6$ and $u = X_1^2 X_5 X_6$, then $\text{supp } u = \{1, 5, 6\}$. In the sequel, we will rely on two specific assumptions. The first one is as follows:

Assumption 2.3 (Boundedness). *Let \mathcal{D}_S be as in (2.2). There is $M \in \mathbb{R}^{>0}$ such that $\sum_{i=1}^n X_i^2 \preceq M$, for all $\underline{X} = (X_1, \dots, X_n) \in \mathcal{D}_S$.*

Then, Assumption 2.3 implies that $\sum_{j \in I_k} X_j^2 \preceq M$, for all $k = 1, \dots, p$. Thus we define

$$(2.5) \quad s_{m+k} := M - \sum_{j \in I_k} X_j^2, \quad k = 1, \dots, p,$$

and set $m' = m + p$ in order to describe the same set \mathcal{D}_S again as:

$$(2.6) \quad \mathcal{D}_S := \{ \underline{A} \in \mathbb{S}^n : s_j(\underline{A}) \succeq 0, \quad j = 1, \dots, m' \}.$$

The second assumption is as follows:

Assumption 2.4 (RIP). *Let \mathcal{D}_S be as in (2.6) and let $f \in \mathbb{R}\langle \underline{X} \rangle$. The index set $J := \{1, \dots, m'\}$ is partitioned into p disjoint sets J_1, \dots, J_p and the two collections $\{I_1, \dots, I_p\}$ and $\{J_1, \dots, J_p\}$ satisfy:*

- (i) *For all $j \in J_k$, $g_j \in \text{Sym } \mathbb{R}\langle \underline{X}(I_k) \rangle$.*
- (ii) *The objective function can be decomposed as $f = f_1 + \dots + f_p$, with $f_k \in \mathbb{R}\langle \underline{X}(I_k) \rangle$, for all $k = 1, \dots, p$.*

(iii) *The running intersection property (RIP) holds, i.e., for all $k = 1, \dots, p-1$, one has*

$$(2.7) \quad I_{k+1} \cap \bigcup_{j \leq k} I_j \subseteq I_\ell \quad \text{for some } \ell \leq k.$$

Even though we assume that I_1, \dots, I_p are explicitly given, one can compute such subsets using the procedure in [WKKM06]. Roughly speaking, this procedure consists of two steps. The first step provides the correlation sparsity pattern (csp) graph of the variables involved in the input polynomial data. The second step computes the maximal cliques of a chordal extension of this csp graph. Even if the computation of all maximal cliques of a graph is an NP hard problem in general, it turns out that this procedure is efficient in practice, due to the properties of chordal graphs (see, e.g., [BP93] for more details on the properties of chordal graphs).

2.3. Hankel and localizing matrices. To $g \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and a linear functional $L : \mathbb{R}\langle \underline{X} \rangle_{2d} \rightarrow \mathbb{R}$, one associates the following two matrices:

- the *noncommutative Hankel matrix* H_L is the matrix indexed by words $u, v \in \langle \underline{X} \rangle_d$, with $(H_L)_{u,v} = L(u^*v)$;
- the *localizing matrix* $H_{L,g}^\uparrow$ is the matrix indexed by words $u, v \in \langle \underline{X} \rangle_{d - \lceil \deg g / 2 \rceil}$, with $(H_{L,g}^\uparrow)_{u,v} = L(u^*gv)$.

The functional L is called *unital* if $L(1) = 1$ and is called *symmetric* if $L(f^*) = L(f)$, for all f belonging to the domain of L . We also recall the following useful facts together with their proofs for the sake of completeness.

Lemma 2.5 ([BKP16, Lemma 1.44]). *Let $g \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and let $L : \mathbb{R}\langle \underline{X} \rangle_{2d} \rightarrow \mathbb{R}$ be a symmetric linear functional. Then, one has:*

- $L(h^*h) \geq 0$ for all $h \in \mathbb{R}\langle \underline{X} \rangle_d$, if and only if, $H_L \succeq 0$;
- $L(h^*gh) \geq 0$ for all $h \in \mathbb{R}\langle \underline{X} \rangle_{d - \lceil \deg g / 2 \rceil}$, if and only if, $H_{L,g}^\uparrow \succeq 0$.

Proof. For $h = \sum_w h_w w \in \mathbb{R}\langle \underline{X} \rangle_d$, let us denote by $\mathbf{h} \in \mathbb{R}^{\sigma(n,d)}$ the vector consisting of all coefficients h_w of h . The first statement now follows from

$$L(h^*h) = \sum_{u,v} h_u h_v L(u^*v) = \sum_{u,v} h_u h_v (H_L)_{u,v} = \mathbf{h}^T H_L \mathbf{h}.$$

Similarly, the second statement follows after checking that $L(h^*gh) = \mathbf{h}^T H_{L,g}^\uparrow \mathbf{h}$. \square

Definition 2.6. *Suppose $L : \mathbb{R}\langle \underline{X} \rangle_{2d+2\delta} \rightarrow \mathbb{R}$ is a linear functional with restriction $\tilde{L} : \mathbb{R}\langle \underline{X} \rangle_{2d} \rightarrow \mathbb{R}$. We associate to L and \tilde{L} the Hankel matrices H_L and $H_{\tilde{L}}$ respectively, and get the block form*

$$H_L = \begin{bmatrix} H_{\tilde{L}} & B \\ B^T & C \end{bmatrix}.$$

We say that L is δ -flat or that L is a flat extension of \tilde{L} , if H_L is flat over $H_{\tilde{L}}$, i.e., if $\text{rank } H_L = \text{rank } H_{\tilde{L}}$.

For a subset $I \subseteq \{1, \dots, p\}$, let us define $H_L(I)$ to be the Hankel submatrix obtained from H_L after retaining only those rows and columns indexed by $w \in \langle \underline{X}(I) \rangle$. When $I \subseteq I_k$ and $g \in \mathbb{R}\langle \underline{X}(I_k) \rangle$, for $k \in \{1, \dots, p\}$, we define the localizing submatrix $H_{L,g}^\uparrow(I)$ in a similar fashion. In particular, $H_L(I_k)$ and $H_{L,g}^\uparrow(I_k)$ can be

seen as Hankel and localizing matrices with rows and columns indexed by a basis of $\mathbb{R}\langle \underline{X}(I_k) \rangle_d$. Both matrices depend only on the variables y_w with $\text{supp } w \in \mathcal{I}_k$.

3. SPARSE REPRESENTATIONS OF NONCOMMUTATIVE POSITIVE POLYNOMIALS

In this section, we prove our main theoretical result, which is a sparse version of the Helton-McCullough archimedean Positivstellensatz (Theorem 2.2). For this, we rely on amalgamation theory for C^* -algebras, see e.g. [Bla78, Voi85].

Given a Hilbert space \mathcal{H} , we denote by $\mathcal{B}(\mathcal{H})$ the set of bounded operators on \mathcal{H} . A C^* -algebra is a complex Banach algebra \mathcal{A} with an involution satisfying $\|xx^*\| = \|x\|^2$ for all $x \in \mathcal{A}$. Equivalently, it is a norm closed subalgebra with involution of $\mathcal{B}(\mathcal{H})$ for some Hilbert space \mathcal{H} .

Theorem 3.1 ([Bla78] or [Voi85, Section 5]). *Let (\mathcal{A}, φ_0) and $\{(\mathcal{B}_k, \varphi_k) : k \in I\}$ be C^* -algebras with states (positive linear functionals), and let ι_k be a state-preserving embedding of \mathcal{A} into \mathcal{B}_k , for each $k \in I$. Then there exists a C^* -algebra D amalgamating the $(\mathcal{B}_k, \varphi_k)$ over (\mathcal{A}, φ_0) . That is, there is a state φ on D , and state-preserving homomorphisms $j_k : \mathcal{B}_k \rightarrow D$, such that $\iota_k \circ j_k = \iota_i \circ j_i$, for all $k, i \in I$, and such that $\bigcup_{k \in I} j_k(\mathcal{B}_k)$ generates D .*

We also recall the construction by Gelfand-Naimark-Segal (GNS) establishing a correspondence between \star -representations of a C^* -algebra and positive linear functionals on it. In our context, the next result [BKP16, Theorem 1.27] restricts to linear functionals on $\mathbb{R}\langle \underline{X} \rangle$ which are positive on an archimedean quadratic module.

Theorem 3.2. *Let $S \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ be given such that its quadratic module M_S is archimedean. Let $L : \mathbb{R}\langle \underline{X} \rangle \rightarrow \mathbb{R}$ be a nontrivial linear functional with $L(M_S) \subseteq \mathbb{R}^{\geq 0}$. Then there exists a tuple $\underline{A} = (A_1, \dots, A_n) \in \mathcal{D}_S^\infty$ and a vector \mathbf{v} such that $L(f) = \langle f(\underline{A})\mathbf{v}, \mathbf{v} \rangle$, for all $f \in \mathbb{R}\langle \underline{X} \rangle$.*

For $k = 1, \dots, p$, let us define

$$M_S^k := \left\{ \sum_{i=1}^K a_i^* s_i a_i : K \in \mathbb{N}, a_i \in \mathbb{R}\langle \underline{X}(I_k) \rangle, s_i \in S \cap \text{Sym } \mathbb{R}\langle \underline{X}(I_k) \rangle \cup \{1\} \right\},$$

and

$$(3.1) \quad M_S^{\text{sparse}} := M_S^1 + \dots + M_S^p.$$

Theorem 3.3. *Let $S \cup \{f\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and let \mathcal{D}_S be as in (2.6) with the additional quadratic constraints (2.5). Suppose Assumption 2.4 holds. If $f(\underline{A}) \succ 0$ for all $\underline{A} \in \mathcal{D}_S^\infty$, then $f \in M_S^{\text{sparse}}$.*

Proof. The proof is by contradiction: suppose that $f(\underline{A}) \succ 0$ for all $\underline{A} \in \mathcal{D}_S^\infty$, and that $f \notin M_S^{\text{sparse}}$. By the Hahn-Banach separation theorem (also known as the Eidelheit-Kakutani Theorem [Bar02, Corollary III.1.7] in this context), there exists a linear functional $L : \mathbb{R}\langle \underline{X} \rangle \rightarrow \mathbb{R}$ with $L(f) \leq 0$ and $L(M_S^{\text{sparse}}) \subseteq \mathbb{R}^{\geq 0}$. Without loss of generality $L(1) = 1$.

Here, we cannot directly apply Theorem 3.2 since M_S^{sparse} is not the quadratic module of $\mathbb{R}\langle \underline{X} \rangle$ generated by the polynomials involved in S . Nevertheless, we will prove that there exists a tuple $\underline{A} = (A_1, \dots, A_n) \in \mathcal{D}_S^\infty$ and a nonzero vector \mathbf{w} such that $L(f) = \langle f(\underline{A})\mathbf{w}, \mathbf{w} \rangle$. Since $f \succ 0$ implies that $\langle f(\underline{A})\mathbf{w}, \mathbf{w} \rangle > 0$, this will contradict the fact that $L(f) \leq 0$.

For $k = 1, \dots, p$, let us denote by $L^k : \mathbb{R}\langle \underline{X}(I_k) \rangle \rightarrow \mathbb{R}$ the restriction of L to $\mathbb{R}\langle \underline{X}(I_k) \rangle$. Observe that $L^k(M_S^k) \subseteq \mathbb{R}^{\geq 0}$. Each linear functional L^k induces a sesquilinear form

$$(g, h) \mapsto \langle g, h \rangle_k := L^k(g^*h)$$

on $\mathbb{R}\langle \underline{X}(I_k) \rangle$, which is positive semidefinite since L^k is positive on sums of hermitian squares, allowing us to apply the Cauchy-Schwartz inequality. Let $\mathcal{N}^k := \{h \in \mathbb{R}\langle \underline{X}(I_k) \rangle : \langle h, h \rangle_k = 0\}$ be the nullvectors corresponding to L^k . By using again the Cauchy-Schwartz inequality, one can show that \mathcal{N}^k is a vector subspace of $\mathbb{R}\langle \underline{X}(I_k) \rangle$, and the sesquilinear form L^k induces an inner product on the quotient space $\mathbb{R}\langle \underline{X}(I_k) \rangle / \mathcal{N}^k$. Let us denote by $\mathcal{H}(I_k)$ the Hilbert space completion of $\mathbb{R}\langle \underline{X}(I_k) \rangle / \mathcal{N}^k$. Since $1 \notin \mathcal{N}^k$, $\mathcal{H}(I_k)$ is nontrivial and separable. By using the fact that L^k is nonnegative on the archimedean quadratic module M_S^k , there exists $N \in \mathbb{N}$ such that $L^k(g^*(N - X_i^2)g) \geq 0$, for all $g \in \mathbb{R}\langle \underline{X}(I_k) \rangle$. Therefore, one has

$$(3.2) \quad 0 \leq \langle X_i g, X_i g \rangle_k = L^k(g^* X_i^2 g) \leq N L^k(g^* g),$$

implying that \mathcal{N}^k is a left ideal. Therefore, the left multiplication operator $\hat{X}_i^k : g \mapsto X_i^k g$ is well-defined on $\mathbb{R}\langle \underline{X}(I_k) \rangle / \mathcal{N}^k$, for all $i \in I_k$. By (3.2), this operator is also bounded and can be extended uniquely to $\mathcal{H}(I_k)$. We fix an orthonormal basis of $\mathcal{H}(I_k)$ and denote by \hat{A}_i^k the corresponding representative of the left multiplication by X_i in $\mathcal{B}(\mathcal{H}(I_k))$ with respect to this basis. Let us denote $\hat{A}^k := (\hat{A}_i^k)_{i \in I_k}$. Then, one has for all $g \in \mathbb{R}\langle \underline{X}(I_k) \rangle$

$$(3.3) \quad L^k(g) = \langle g(\hat{A}^k) \mathbf{v}^k, \mathbf{v}^k \rangle_k,$$

where $\mathbf{v}^k \in \mathcal{H}^k$ is the image of the identity polynomial. We denote by φ_k the induced state on $\mathcal{B}(\mathcal{H}(I_k))$, cf. Figure 1.

Now, the proof proceeds by induction on p . With $p = 1$, this corresponds to the dense representation result stated in Theorem 2.2.


Case $p = 2$. First, note that the running intersection property (2.7) holds in this case. If $I_1 \cap I_2 = \emptyset$, then one has $f_k(\underline{A}) \succ 0$, for all $\underline{A} \in \mathcal{D}_S^\infty$. As a consequence of Theorem 2.2, one has $f_k \in M_S^k$, which implies that $f = f_1 + f_2 \in M_S^{\text{sparse}}$.

Next, let us suppose that $I_1 \cap I_2 \neq \emptyset$. Let us define the sesquilinear form $(g, h) \mapsto \langle g, h \rangle_{12} := L^1(g^*h)$ on $\mathbb{R}\langle \underline{X}(I_1 \cap I_2) \rangle$. As above, we obtain $\mathcal{N}^{12} := \{h \in \mathbb{R}\langle \underline{X}(I_1 \cap I_2) \rangle : \langle h, h \rangle_{12} = 0\}$ and the Hilbert space completion $\mathcal{H}(I_1 \cap I_2)$ of $\mathbb{R}\langle \underline{X}(I_1 \cap I_2) \rangle / \mathcal{N}^{12}$. We denote by L^0 the restriction of L^1 (or, equivalently, L^2) to $\mathbb{R}\langle \underline{X}(I_1 \cap I_2) \rangle$, and by φ_0 the induced state on $\mathcal{B}(\mathcal{H}(I_1 \cap I_2))$.

For $k \in \{1, 2\}$, let us denote by $i_k : \mathbb{R}\langle \underline{X}(I_1 \cap I_2) \rangle \rightarrow \mathbb{R}\langle \underline{X}(I_k) \rangle$ the canonical embedding. Next we apply Theorem 3.1 with $I = \{1, 2\}$, $\mathcal{A} = \mathcal{B}(\mathcal{H}(I_1 \cap I_2))$ endowed with φ_0 , $\mathcal{B}_k = \mathcal{B}(\mathcal{H}(I_k))$ endowed with φ_k , and $\iota_k : \mathcal{B}(\mathcal{H}(I_1 \cap I_2)) \rightarrow \mathcal{B}(\mathcal{H}(I_k))$ being the canonical embedding. As displayed in Figure 1, we obtain an amalgamation $D \subseteq \mathcal{B}(\mathcal{H})$, with state φ and homomorphisms $j_k : \mathcal{B}(\mathcal{H}(I_k)) \rightarrow D$ such that $\iota_1 \circ j_1 = \iota_2 \circ j_2$.

Next perform the GNS construction with (D, φ) . There is a Hilbert space \mathcal{K} , representation $\pi : D \rightarrow \mathcal{B}(\mathcal{K})$ and vector $\xi \in \mathcal{K}$ so that $\varphi(a) = \langle \pi(a)\xi, \xi \rangle$. Then, let us define $\underline{A} := (A_1, \dots, A_n)$, with

$$A_i := \begin{cases} \pi(j_1(\hat{A}_i^1)) & \text{if } i \in I_1, \\ \pi(j_2(\hat{A}_i^2)) & \text{if } i \in I_2, \end{cases}$$

FIGURE 1. Amalgamation for the case $p = 2$.

For all $g \in \mathbb{R}\langle \underline{X} \rangle$, we now set $\tilde{L}(g) := \langle g(\underline{A})\xi, \xi \rangle$. We claim that \tilde{L} extends L^k . Indeed, for $g \in \mathbb{R}\langle \underline{X}(I_k) \rangle$ we have

$$\begin{aligned} \tilde{L}(g) &= \langle g(\underline{A})\xi, \xi \rangle = \langle g(\pi(j_k(\hat{A}^k)))\xi, \xi \rangle = \langle \pi(g(j_k(\hat{A}^k)))\xi, \xi \rangle \\ &= \varphi(g(j_k(\hat{A}^k))) = \varphi(j_k(g(\hat{A}^k))) = \varphi_k(g(\hat{A}^k)) = L^k(g). \end{aligned}$$

Therefore,

$$\langle f(\underline{A})\xi, \xi \rangle = \tilde{L}(f) = \tilde{L}(f_1) + \tilde{L}(f_2) = L^1(f_1) + L^2(f_2) = L(f) \leq 0.$$

It only remains to prove that $\underline{A} \in \mathcal{D}_S^\infty$, i.e., that $s(\underline{A}) \succeq 0$, for all $s \in S$. By assumption, $s \in \text{Sym } \mathbb{R}\langle \underline{X}(I_k) \rangle$ for some $k \in \{1, 2\}$, so

$$(3.4) \quad s(\underline{A}) = s((\pi \circ j_k)(\hat{A}^k)) = (\pi \circ j_k)(s(\hat{A}^k)).$$

Since $\mathbb{R}\langle \underline{X}(I_k) \rangle / \mathcal{N}^k$ is dense in \mathcal{H}^k , one can approximate as closely as desired $\mathbf{u} \in \mathcal{H}^k$ by elements of $\mathbb{R}\langle \underline{X}(I_k) \rangle / \mathcal{N}^k$. We prove that $\langle s(\hat{A}^k)\mathbf{u}, \mathbf{u} \rangle_k \geq 0$, where \mathbf{u} stands for a vector representative of $u \in \mathbb{R}\langle \underline{X}(I_k) \rangle / \mathcal{N}^k$. Given such a vector \mathbf{u} , there exists a polynomial $g \in \mathbb{R}\langle \underline{X}(I_k) \rangle$ (by construction), such that $\mathbf{u} = g(\hat{A}^k)\mathbf{v}^k$. Next, the following holds:

$$\langle s(\hat{A}^k)\mathbf{u}, \mathbf{u} \rangle_k = \langle s(\hat{A}^k)g(\hat{A}^k)\mathbf{v}^k, g(\hat{A}^k)\mathbf{v}^k \rangle_k = \langle (sg)(\hat{A}^k)\mathbf{v}^k, g(\hat{A}^k)\mathbf{v}^k \rangle_k = L^k(g^*sg),$$

where the last equality comes from (3.3). Since $g^*sg \in M_S^k$ and $L^k(M_S^k) \subseteq \mathbb{R}^{\geq 0}$, one has $\langle s(\hat{A}^k)\mathbf{u}, \mathbf{u} \rangle_k \geq 0$, which implies that $c := s(\hat{A}^k) \succeq 0$. Since c is a nonnegative element of the C^* -algebra $\mathcal{B}(\mathcal{H}(I_k))$, there exists $b \in \mathcal{B}(\mathcal{H}(I_k))$ such that $c = b^*b$. Eventually by (3.4), one has $s(\underline{A}) = (\pi \circ j_k)(s(\hat{A}^k)) = \pi(j_k(c)) = \pi(j_k(b^*b)) = \pi(j_k(b))^* \pi(j_k(b)) \succeq 0$, yielding $\underline{A} \in \mathcal{D}_S^\infty$, the desired result.

General case. Now assume $p > 2$. For each $m \leq p$ we will construct a tuple $\hat{A}^m \in \mathcal{D}_S^\infty$ and a unit vector ξ^m with

$$(3.5) \quad L(f) = \langle f(\hat{A}^m)\xi^m, \xi^m \rangle$$

for all $f \in \sum_{j \leq m} \mathbb{R}\langle \underline{X}(I_j) \rangle$. Observe that the right-hand side of (3.5) defines a linear functional on the entire $\mathbb{R}\langle \underline{X} \rangle$, which we denote by \tilde{L}^m .

The basis for the induction, $p \leq 2$, has been established above. Let $p > m \geq 2$. By the running intersection property (2.7), there is $k \leq m$ with $(\cup_{j \leq m} I_j) \cap I_{m+1} \subseteq I_k$. We denote by L^m the restriction of \tilde{L}^m to $\mathbb{R}\langle \underline{X}(\cup_{j \leq m} I_j) \rangle$, by L^{m+1} the restriction of L to $\mathbb{R}\langle \underline{X}(I_{m+1}) \rangle$, and by L^0 the restriction of L (or, equivalently, of \tilde{L}^m) to $\mathbb{R}\langle \underline{X}(I_k) \rangle$. As before, Theorem 3.2 produces Hilbert spaces $\mathcal{H}_{m+1}, \mathcal{H}_0$ and states φ_{m+1}, φ_0 on $\mathcal{B}(\mathcal{H}_{m+1}), \mathcal{B}(\mathcal{H}_0)$, respectively. By the induction hypothesis there is a Hilbert space \mathcal{H}_m and a state φ_m on $\mathcal{B}(\mathcal{H}_m)$.

Now amalgamate $(\mathcal{B}(\mathcal{H}_m), \varphi_m)$ and $(\mathcal{B}(\mathcal{H}_{m+1}), \varphi_{m+1})$; if $(\cup_{j \leq m} I_j) \cap I_{m+1} = \emptyset$, then we amalgamate them over \mathbb{R} , and otherwise over $(\mathcal{B}(\mathcal{H}_0), \varphi_0)$. Applying the GNS construction to the amalgamated C^* -algebra (D, φ) then yields the desired $\hat{A}^{m+1} \in \mathcal{D}_S^\infty$ and a unit vector ξ with $L(f) = \langle f(\hat{A}^{m+1})\xi, \xi \rangle$ for all $f \in \sum_{j \leq m+1} \mathbb{R}\langle \underline{X}(I_j) \rangle$. \square

The reader will notice that the RIP property is used subtly in the two last paragraphs of the proof of Theorem 3.3. Next, we provide an example demonstrating that sparsity without a RIP-type condition is not sufficient to deduce sparsity in SOHS decompositions.

Example 3.4. Consider the case of three variables $\underline{X} = (X_1, X_2, X_3)$ and the polynomial

$$\begin{aligned} f &= (X_1 + X_2 + X_3)^2 \\ &= X_1^2 + X_2^2 + X_3^2 + X_1X_2 + X_2X_1 + X_1X_3 + X_3X_1 + X_2X_3 + X_3X_2 \in \Sigma\langle \underline{X} \rangle. \end{aligned}$$

Then $f = f_1 + f_2 + f_3$, with

$$\begin{aligned} f_1 &= \frac{1}{2}X_1^2 + \frac{1}{2}X_2^2 + X_1X_2 + X_2X_1 \in \mathbb{R}\langle X_1, X_2 \rangle, \\ f_2 &= \frac{1}{2}X_1^2 + \frac{2}{1}X_3^2 + X_1X_3 + X_3X_1 \in \mathbb{R}\langle X_1, X_3 \rangle, \\ f_3 &= \frac{1}{2}X_2^2 + \frac{1}{2}X_3^2 + X_2X_3 + X_3X_2 \in \mathbb{R}\langle X_2, X_3 \rangle. \end{aligned}$$

However, the sets $I_1 = \{1, 2\}$, $I_2 = \{1, 3\}$ and $I_3 = \{2, 3\}$ do not satisfy the RIP condition (2.7) and $f \notin \Sigma\langle \underline{X} \rangle^{\text{sparse}} := \Sigma\langle X_1, X_2 \rangle + \Sigma\langle X_1, X_3 \rangle + \Sigma\langle X_2, X_3 \rangle$ since it has a unique Gram matrix by homogeneity.

Now consider $S = \{1 - X_1^2, 1 - X_2^2, 1 - X_3^2\}$. Then \mathcal{D}_S is as in (2.6), M_S^{sparse} is as in (3.1) and $f|_{\mathcal{D}_S^\infty} \succeq 0$. However, we claim that $f + \lambda \in M_S^{\text{sparse}}$ iff $\lambda \geq 3$. Clearly,

$$f + 3 = (X_1 + X_2)^2 + (X_1 + X_3)^2 + (X_2 + X_3)^2 + (1 - X_1^2) + (1 - X_2^2) + (1 - X_3^2) \in M_S^{\text{sparse}}.$$

Hence it suffices to show there exist linear functionals $L_k : \mathbb{R}\langle \underline{X}(I_k) \rangle \rightarrow \mathbb{R}$ satisfying

- (1) $L_k(1) = 1$;
- (2) The corresponding (infinite) Hankel matrix H_{L_k} and its localizing variants $H_{L_k, g}^\uparrow$ are positive semidefinite for all $g \in S$;

- (3) $L_j|_{\mathbb{R}\langle \underline{X}(I_j \cap I_k) \rangle} = L_k|_{\mathbb{R}\langle \underline{X}(I_j \cap I_k) \rangle}$ for all j, k ;
 (4) $\sum_k L_k(f_k) = -3$.

Define

$$A = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad B = -A$$

and let

$$L_j(g) = \text{tr } g(A, B) \quad \text{for } g \in \mathbb{R}\langle \underline{X}(I_j) \rangle.$$

Since $L_j(f_j) = -1$, properties (1), (2) and (4) are easily verified. For property (3), given, say $h \in \mathbb{R}\langle \underline{X}(I_1) \rangle \cap \mathbb{R}\langle \underline{X}(I_3) \rangle = \mathbb{R}\langle X_2 \rangle$, we have

$$\begin{aligned} L_2(h) &= \text{tr } h(B), \\ L_3(h) &= \text{tr } h(A). \end{aligned}$$

But matrices A and B are orthogonally equivalent as $UAU^T = B$ for

$$U = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix},$$

whence $h(B) = h(UAU^T) = Uh(A)U^T$ and $h(A)$ have the same trace.

4. SPARSE GNS CONSTRUCTION AND OPTIMIZER EXTRACTION

The aim of this section is to provide a general algorithm to extract solutions of sparse noncommutative optimization problems. We will apply this algorithm below to eigenvalue optimization (Section 5) and trace optimization (Section 6). For this purpose, we first present sparse noncommutative versions of theorems by Curto and Fialkow. In the commutative case, Curto and Fialkow provided sufficient conditions for linear functionals on the set of degree $2d$ polynomials to be represented by integration with respect to a nonnegative measure. The main sufficient condition to guarantee such a representation is flatness (see Definition 2.6) of the corresponding Hankel matrix.

In the dense case, [BKP16, Theorem 1.69] provides a first noncommutative variant. As this will be needed in the sequel, we recall this theorem and a sketch of its proof, which relies on a finite-dimensional GNS construction.

Theorem 4.1. *Let $S \cup \{f\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle_{2d}$ and set $\delta := \max\{\lceil \deg(g)/2 \rceil : g \in S \cup \{f\}\}$. Let $L : \mathbb{R}\langle \underline{X} \rangle_{2d+2\delta} \rightarrow \mathbb{R}$ be a unital linear functional satisfying $L(M_{S, 2d+2\delta}) \subseteq \mathbb{R}^{\geq 0}$. If L is δ -flat, then there exist $\hat{A} \in \mathcal{D}_S(r)$ for some $r \leq \sigma(n, d)$ and a unit vector \mathbf{v} such that*

$$(4.1) \quad L(f) = \langle f(\hat{A})\mathbf{v}, \mathbf{v} \rangle.$$

Proof. Let $r := \text{rank } H_L$. Since H_L is a positive semidefinite matrix, we obtain the Gram matrix decomposition $H_L = [\langle \mathbf{u}, \mathbf{w} \rangle]_{u,w}$ with vectors $\mathbf{u}, \mathbf{w} \in \mathbb{R}^r$, where the labels are words of degree at most $d + \delta$. Then, we define the following finite-dimensional Hilbert space

$$\mathcal{H} := \text{span} \{ \mathbf{w} \mid \deg w \leq d + \delta \} = \text{span} \{ \mathbf{w} \mid \deg w \leq d \},$$

where the equality comes from the flatness assumption. Afterwards, one can directly consider the operators \hat{A}_i representing the left multiplication by X_i on \mathcal{H} , i.e., $\hat{A}_i \mathbf{w} = \mathbf{X}_i \mathbf{w}$. Thanks to the flatness assumption, the operators \hat{A}_i are well-defined and one can show that they are symmetric, whence $\hat{A} = (\hat{A}_1, \dots, \hat{A}_n) \in \mathcal{D}_S(r)$ and the representation given in (4.1) holds with $\mathbf{v} = \mathbf{1}$. \square

We now give the sparse version of Theorem 4.1.

Theorem 4.2. *Let $S \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle_{2d}$, and assume \mathcal{D}_S is as in (2.6) with the additional quadratic constraints (2.5). Suppose Assumption 2.4(i) holds. Set $\delta := \max\{\lceil \deg(g)/2 \rceil : g \in S \cup 1\}$. Let $L : \mathbb{R}\langle \underline{X} \rangle_{2d+2\delta} \rightarrow \mathbb{R}$ be a unital linear functional satisfying $L(M_{S,d}^{\text{sparse}}) \subseteq \mathbb{R}^{\geq 0}$. Assume that the following holds:*

(H1) $H_L(I_k)$ and $H_L(I_k \cap I_j)$ are δ -flat, for all $j, k \in \{1, \dots, p\}$.

Then, there exist finite-dimensional Hilbert spaces $\mathcal{H}(I_k)$ with dimension r_k , for all $k \in \{1, \dots, p\}$, Hilbert spaces $\mathcal{H}(I_j \cap I_k) \subseteq \mathcal{H}(I_j), \mathcal{H}(I_k)$ for all pairs (j, k) with $I_j \cap I_k \neq 0$, and operators \hat{A}^k, \hat{A}^{jk} , acting on them, respectively. Further, there are unit vectors $\mathbf{v}^j \in \mathcal{H}(I_j)$ and $\mathbf{v}^{jk} \in \mathcal{H}(I_j \cap I_k)$ such that

$$(4.2) \quad \begin{aligned} L(f) &= \langle f(\hat{A}^j) \mathbf{v}^j, \mathbf{v}^j \rangle \quad \text{for all } f \in \mathbb{R}\langle \underline{X}(I_j) \rangle_{2d}, \\ L(g) &= \langle g(\hat{A}^{jk}) \mathbf{v}^{jk}, \mathbf{v}^{jk} \rangle \quad \text{for all } g \in \mathbb{R}\langle \underline{X}(I_j \cap I_k) \rangle_{2d}. \end{aligned}$$

Assuming that for all pairs (j, k) with $I_j \cap I_k \neq 0$, one has

(H2) the matrices $(\hat{A}_i^{jk})_{i \in I_j \cap I_k}$ have no common complex invariant subspaces, then there exist $\underline{A} \in \mathcal{D}_S(r)$, with $r := r_1 \cdots r_p$, and a unit vector \mathbf{v} such that

$$(4.3) \quad L(f) = \langle f(\underline{A}) \mathbf{v}, \mathbf{v} \rangle$$

for all $f \in \sum_j \mathbb{R}\langle \underline{X}(I_j) \rangle_{2d}$.

In the proof of Theorem 4.2 we will make use of the following simple linear algebra observation.

Lemma 4.3. *Let $Z \in M_n(\mathbb{R})$. If $\text{tr}(ZA) = 0$ for all $A \in M_n(\mathbb{R})$, then $Z = 0$.*

Proof. We have $\text{tr}(ZZ^T) = 0$ whence $ZZ^T = 0$ and thus $Z = 0$. \square


Proof of Theorem 4.2. Start by applying Theorem 4.1 to $L|_{\mathbb{R}\langle \underline{X}(I_j) \rangle}$ and $L|_{\mathbb{R}\langle \underline{X}(I_j \cap I_k) \rangle}$ to obtain the desired Hilbert spaces $\mathcal{H}(I_j), \mathcal{H}(I_j \cap I_k)$, unit vectors $\mathbf{v}^j, \mathbf{v}^{jk}$ and operators \hat{A}^j, \hat{A}^{jk} satisfying (4.2). Note that we may assume $\mathcal{H}(I_j \cap I_k) \subseteq \mathcal{H}(I_j), \mathcal{H}(I_k)$ as the map $f(\hat{A}^{jk}) \mathbf{v}^{jk} \mapsto f(\hat{A}^j) \mathbf{v}^j$ is an isometry by construction. Then

$$(4.4) \quad \hat{A}^{jk} = \hat{A}^j|_{\mathcal{H}(I_j \cap I_k)} = \hat{A}^k|_{\mathcal{H}(I_j \cap I_k)}.$$

Let us denote by $\mathcal{A}(I_j)$ and $\mathcal{A}(I_j \cap I_k)$ the algebras generated by \hat{A}^j , and \hat{A}^{jk} , respectively. By (4.4), the map $\hat{A}_\ell^{jk} \mapsto \hat{A}_\ell^j$ is a \star -homomorphism $\mathcal{A}(I_j) \rightarrow \mathcal{A}(I_j \cap I_k)$. With $r_k = \dim \mathcal{H}(I_k)$ and $r_{jk} = \dim \mathcal{H}(I_j \cap I_k)$, one has $\mathcal{A}(I_k) \subseteq \mathbb{M}_{r_k}(\mathbb{R})$ and $\mathcal{A}(I_j \cap I_k) \subseteq \mathbb{M}_{r_{jk}}(\mathbb{R})$.

We next want to find a finite-dimensional C^* -algebra \mathcal{A} , i.e., a subalgebra of $\mathbb{M}_m(\mathbb{R})$ for some $m \in \mathbb{N}$ making the diagram in Figure 2 commute. (From now on we focus on the $p = 2$ case, as the general case follows by a simple inductive argument.) By the amalgamation property of C^* -algebras, Theorem 3.1, we can always find such an *infinite-dimensional* \mathcal{A} . However, as shown in Example 4.4, there may not be a suitable finite-dimensional \mathcal{A} . To ensure this, we assume that (H2) holds, namely that the matrices $(\hat{A}_i^{12})_{i \in I_1 \cap I_2}$ have no common complex invariant subspaces, i.e., $\mathcal{A}(I_1 \cap I_2) = \mathbb{M}_{r_{12}}(\mathbb{R})$. Then, for all $A \in \mathcal{A}(I_1 \cap I_2) = \mathbb{M}_{r_{12}}(\mathbb{R})$, $\iota_k(A)$ is just a direct sum of copies of A , up to orthogonal equivalence (by the Skolem-Noether theorem [BO13, Section III.3]), i.e., there are orthogonal matrices U_k such that

$$\iota_k(A) = U_k^T (\mathbf{I}_{r_k/r_{12}} \otimes A) U_k,$$

FIGURE 2. Amalgamation of finite-dimensional C^* -algebras

for all $k \in \{1, 2\}$. By replacing \hat{A}^k with their conjugates $U_k^T \hat{A}^k U_k$, and \mathbf{v}^k by $U_k^T \mathbf{v}^k$, we may without loss of generality assume $\iota_k(A) = \mathbf{I}_{r_k/r_{12}} \otimes A$.

The linear functional L induces linear functionals $\check{L}^k, \check{L}^{12}$ on $\mathcal{A}(I_k), \mathcal{A}(I_1 \cap I_2)$ given by $B \mapsto \text{tr}(B\mathbf{v}^k(\mathbf{v}^k)^T)$ and $C \mapsto \text{tr}(C\mathbf{v}^{12}(\mathbf{v}^{12})^T)$, respectively. Write $\mathbf{v}^k = \sum_{j=1}^{r_k/r_{12}} e_j^k \otimes u_j^k$ for the standard basis vectors $e_j^k \in \mathbb{R}^{r_k/r_{12}}$ and some vectors $u_j^k \in \mathbb{R}^{r_{12}}$. Then for $C \in \mathcal{A}(I_1 \cap I_2) = M_{r_{12}}(\mathbb{R})$ we have

$$\begin{aligned}
\check{L}^{12}(C) &= \text{tr}(C\mathbf{v}^{12}(\mathbf{v}^{12})^T) \\
&= \check{L}^k(\mathbf{I}_{r_k/r_{12}} \otimes C) = \text{tr}((\mathbf{I} \otimes C)\mathbf{v}^k(\mathbf{v}^k)^T) \\
&= \text{tr}\left((\mathbf{I} \otimes C)\left(\sum_{j=1}^{r_k/r_{12}} e_j^k \otimes u_j^k\right)\left(\sum_{j=1}^{r_k/r_{12}} e_j^k \otimes u_j^k\right)^T\right) \\
&= \text{tr}\left((\mathbf{I} \otimes C) \sum_{i,j=1}^{r_k/r_{12}} (e_j^k (e_i^k)^T) \otimes (u_j^k (u_i^k)^T)\right) \\
&= \sum_{i,j=1}^{r_k/r_{12}} \text{tr}((e_j^k (e_i^k)^T) \otimes C u_j^k (u_i^k)^T) \\
&= \sum_{i,j=1}^{r_k/r_{12}} \text{tr}(e_j^k (e_i^k)^T) \text{tr}(C u_j^k (u_i^k)^T) \\
&= \text{tr}\left(C \sum_{j=1}^{r_k/r_{12}} u_j^k (u_j^k)^T\right).
\end{aligned}$$

From the equality $\text{tr}(C\mathbf{v}^{12}(\mathbf{v}^{12})^T) = \text{tr}(C \sum_{j=1}^{r_k/r_{12}} u_j^k (u_j^k)^T)$ for all $C \in M_{r_{12}}(\mathbb{R})$ we deduce using Lemma 4.3 that $\mathbf{v}^{12}(\mathbf{v}^{12})^T = \sum_{j=1}^{r_k/r_{12}} u_j^k (u_j^k)^T$. Since the left-hand side outer product is rank one, each of the u_j^k must be a scalar multiple of \mathbf{v}^{12} , say $u_j^k = \lambda_{jk} \mathbf{v}^{12}$. Thus $\mathbf{v}^k = \sum_{j=1}^{r_k/r_{12}} \lambda_{jk} e_j^k \otimes \mathbf{v}^{12}$ and $\sum_j \lambda_{jk}^2 = 1$ since $\|\mathbf{v}^k\| = 1$.

Now set $\mathcal{A} := \mathbb{M}_{r_1 r_2}(\mathbb{R})$ and define $j_1(A) := \mathbf{I}_{r_2} \otimes A$, for all $A \in \mathcal{A}(I_1)$, and

$$(4.5) \quad j_2(B) := (U^T \otimes \mathbf{I}_{r_{12}})(\mathbf{I}_{r_1} \otimes B)(U \otimes \mathbf{I}_{r_{12}}),$$

for all $B \in \mathcal{A}(I_2)$. Here U is an $r_1 r_2 / r_{12}$ orthogonal matrix to be determined later. This amalgamates the diagram in Figure 2 (independently of the choice of U).

Each extension of the linear functional \check{L}^k to a linear functional on \mathcal{A} is of the form

$$(4.6) \quad C \mapsto \operatorname{tr} \left(C \sum_{\ell=1}^{r_{3-k}} \mu_{\ell k} e_{\ell}^{3-k} \otimes \mathbf{v}^k \right) = \operatorname{tr} \left(C \underbrace{\sum_{\ell} \sum_{j=1}^{r_k/r_{12}} \mu_{\ell k} e_{\ell}^{3-k} \otimes \lambda_{jk} e_j^k \otimes \mathbf{v}^{12}}_{\mathbf{w}^k} \right),$$

where $\sum_{\ell} \mu_{\ell k}^2 = 1$. Since the vectors \mathbf{w}^k are norm one, there is a unitary U with $U \mathbf{w}_1 = \mathbf{w}_2$. Using this U in the definition (4.5), the extension (4.6) of \check{L}^1 to a linear functional $\check{L} : \mathcal{A} \rightarrow \mathbb{R}$ also extends \check{L}^2 (via j_2).

Now define the operators $\underline{A} := (A_1, \dots, A_n)$, with

$$A_i := \begin{cases} j_1(\hat{A}_i^1) & \text{if } i \in I_1, \\ j_2(\hat{A}_i^2) & \text{if } i \in I_2. \end{cases}$$

Then $L(f) = \langle f(\underline{A})(\mathbf{w}^1 \otimes \mathbf{v}^{12}), \mathbf{w}^1 \otimes \mathbf{v}^{12} \rangle$ for all $f \in \mathbb{R}\langle \underline{X}(I_1) \rangle_{2d} + \mathbb{R}\langle \underline{X}(I_2) \rangle_{2d}$.

To conclude the proof note that each A_i is symmetric and $\underline{A} \in \mathcal{D}_S(r_1 r_2)$. For the latter we use the fact that each constraint g is either in $\mathbb{R}\langle \underline{X}(I_1) \rangle$ or $\mathbb{R}\langle \underline{X}(I_2) \rangle$, and that \star -subalgebras of matrix algebras admit square roots of positive semidefinite operators. \square

Example 4.4 (Non-amalgamation in the category of finite-dimensional algebras). For given I_1, I_2 , suppose $\mathcal{A}(I_1 \cap I_2)$ is generated by the 2×2 diagonal matrix

$$A^{12} = \begin{pmatrix} 1 & \\ & 2 \end{pmatrix},$$

and assume $\mathcal{A}(I_1) = \mathcal{A}(I_2) = \mathbb{M}_3(\mathbb{R})$. (Observe that $\mathcal{A}(I_1 \cap I_2)$ is the algebra of all diagonal matrices.) For each $k \in \{1, 2\}$, let us define $\iota_k(A) := A \oplus k$, for all $A \in \mathcal{A}(I_1 \cap I_2)$. We claim that there is no finite-dimensional C^* -algebra \mathcal{A} amalgamating the above Figure 2. Indeed, by the Skolem-Noether theorem, every homomorphism $\mathbb{M}_n(\mathbb{R}) \rightarrow \mathbb{M}_m(\mathbb{R})$ is of the form $x \mapsto P^{-1}(x \otimes \mathbf{I}_{m/n})P$ for some invertible P ; in particular, n divides m . If a desired \mathcal{A} existed, then the matrices $(A^{12} \oplus 1) \otimes \mathbf{I}_k$ and $(A^{12} \oplus 2) \otimes \mathbf{I}_k$ would be similar. But they are not as is easily seen from eigenvalue multiplicities.

Remark 4.5. Theorem 4.2 can be seen as a noncommutative variant of the result by Lasserre stated in [Las06, Theorem 3.7], related to the minimizers extraction in the context of sparse polynomial optimization. In the sparse commutative case, Lasserre assumes flatness of each moment matrix indexed in the canonical basis of $\mathbb{R}\langle \underline{X}(I_k) \rangle_d$, for each $k \in \{1, \dots, p\}$, which is similar to our flatness condition (H1). In addition, Lasserre assumes that each moment matrix indexed in the canonical basis of $\mathbb{R}\langle \underline{X}(I_j \cap I_k) \rangle_d$ is rank one, for all pairs (j, k) with $I_j \cap I_k \neq \emptyset$, which is the commutative analog of our irreducibility condition (H2).

4.1. Implementing the Sparse GNS Construction. As in the dense case, we can summarize the sparse GNS construction procedure described in the proof of Theorem 4.2 into an algorithm, called **SparseGNS**, stated below in Algorithm 4.6, for the case $p = 2$ (the general case is similar).

This algorithm describes how to compute the tuple $\underline{A} = (A_1, \dots, A_n)$ of amalgamated matrices acting on $\mathcal{H} = \mathbb{R}^{r_1 r_2} \cong \mathcal{H}(I_1) \otimes \mathbb{R}^{r_2} \cong \mathbb{R}^{r_1} \otimes \mathcal{H}(I_2)$, and a vector \mathbf{v} satisfying (4.3).

Algorithm 4.6. SparseGNS

Input: H_L , Hankel matrix of L .

- 1: Apply the GNS construction to obtain $\mathcal{H}(I_1)$, $\mathcal{H}(I_2)$ and $\mathcal{H}(I_1 \cap I_2)$ of respective dimensions r_1 , r_2 and r_{12} , associated to $H_L(I_1)$, $H_L(I_2)$ and $H_L(I_1 \cap I_2)$, as well as \hat{A}^1 , \hat{A}^2 and \hat{A}^{12} acting on $\mathcal{H}(I_1)$, $\mathcal{H}(I_2)$ and $\mathcal{H}(I_1 \cap I_2)$, respectively. \triangleright the dense GNS algorithm is implemented in e.g. `NCS0Stools` [CKP11]
- 2: Find the corresponding unit vectors $\mathbf{v}^1 \in \mathcal{H}(I_1)$, $\mathbf{v}^2 \in \mathcal{H}(I_2)$ and $\mathbf{v}^{12} \in \mathcal{H}(I_1 \cap I_2)$ so that (4.2) holds.
- 3: **if** The flatness (H1) and irreducibility (H2) conditions from Theorem 4.2 do not hold **then**
- 4: Stop
- 5: **end if**
- 6: **for** $k \in \{1, 2\}$, $i \in I_1 \cap I_2$ **do**
- 7: Compute the block diagonalization $\hat{A}_i^k = \text{diag}(\chi_{i,\ell}^k)_\ell$ \triangleright e.g., by [MKKK10, Algorithm 4.1]
- 8: Compute invertible matrices $(P_\ell)_{\ell>1}$ such that $P_\ell^{-1} \chi_{i,\ell}^k P_\ell = \chi_{i,1}^k$
- 9: Normalize each P_ℓ to make it orthogonal. Use them to change the basis in the blocks $(\chi_{i,\ell}^k)_{\ell>1}$ \triangleright Thus, one has $\hat{A}_i^k = \mathbf{I} \otimes \chi_i^k$
- 10: Compute an orthogonal P such that $P^{-1} \chi_i^k P = \hat{A}_i^{12}$ \triangleright Hence, without loss of generality, $\hat{A}_i^{12} = \chi_i^k$
- 11: Decompose $\mathbf{v}^k = \sum_j \lambda_{jk} e_j^k \otimes \mathbf{v}^{12}$
- 12: **end for**
- 13: Find an orthogonal matrix U sending $e_1^1 \otimes \sum_j \lambda_{j1} e_j^1 \mapsto e_1^2 \otimes \sum_j \lambda_{j2} e_j^2$
- 14: **for** $i \in \{1, \dots, n\}$ **do**
- 15:
$$A_i := \begin{cases} \mathbf{I}_{r_2} \otimes \hat{A}_i^1 & \text{if } i \in I_1, \\ (U^T \otimes \mathbf{I}_{r_{12}})(\mathbf{I}_{r_1} \otimes \hat{A}_i^2)(U \otimes \mathbf{I}_{r_{12}}) & \text{if } i \in I_2 \end{cases}$$
- 16: **end for**
- 17: Compute $\mathbf{v} = e_1^1 \otimes \mathbf{v}^1$

Output: $\underline{A} = (A_1, \dots, A_n)$ and \mathbf{v} .

Corollary 4.7. *The procedure SparseGNS described in Algorithm 4.6 is sound and returns the tuple \underline{A} and the vector \mathbf{v} from Theorem 4.2.*

Proof. Correctness of the algorithm has been essentially established in the proof of Theorem 4.2. Both computation in Line 8 and Line 10 can be performed since the only homomorphisms out of full matrix algebras are ampliations composed with a conjugation (by the Skolem-Noether theorem). One can perform an orthogonal change of basis in Line 9, and $\hat{A}_i^k = \mathbf{I} \otimes \chi_i^k$, for all $k \in \{1, 2\}$ and $i \in I_1 \cap I_2$. Indeed, let us assume that a matrix P is invertible, and the map $\phi : A \mapsto P^{-1}AP$ from $\mathbb{M}_n(\mathbb{R})$ to $\mathbb{M}_n(\mathbb{R})$ preserves transposes. Then, the following equalities

$$\phi(A^T) = P^{-1}A^T P = (P^{-1}AP)^T = P^T A^T P^{-T}$$

imply that PP^T commutes with all $n \times n$ matrices. Therefore, PP^T is a scalar matrix, and P is a scalar multiple of an orthogonal matrix, the desired result. Eventually, each component of the tuple A , given in Line 15, is well defined by

construction and gives rise to the desired amalgamation. Line 17 constructs the vector \mathbf{v} needed for (4.2) to hold. \square

5. EIGENVALUE OPTIMIZATION OF NONCOMMUTATIVE SPARSE POLYNOMIALS

The aim of this section is to provide SDP relaxations allowing one to under-approximate the smallest eigenvalue that a given nc polynomial can attain on a tuple of symmetric matrices from a given semialgebraic set. The unconstrained case is handled in Section 5.1, where we show how to compute a lower bound on the smallest eigenvalue via solving an SDP. The constrained case is handled in Section 5.2, where we derive a hierarchy of lower bounds converging to the minimal eigenvalue, assuming that the quadratic module is archimedean and that RIP holds (Assumption 2.4).

We first recall the celebrated Helton-McCullough Sums of Squares theorem [Hel02, McC01] stating the equivalence between sums of hermitian squares (SOHS) and positive semidefinite nc polynomials.

Theorem 5.1. *Given $f \in \mathbb{R}\langle X \rangle$, we have $f(\underline{A}) \succeq 0$, for all $\underline{A} \in \mathbb{S}^n$, if and only if $f \in \Sigma\langle X \rangle$.*

In contrast with the constrained case where we obtain the analog of Putinar's Positivstellensatz in Theorem 3.3, there is no sparse analog of the Helton-McCullough Sums of Squares theorem, as shown in the following example.

Example 5.2. Let $v = [X_1 \ X_1X_2 \ X_2 \ X_3 \ X_3X_2]$,

$$(5.1) \quad G = \begin{bmatrix} 1 & -1 & -1 & 0 & \alpha \\ -1 & 2 & 0 & -\alpha & 0 \\ -1 & 0 & 3 & -1 & 9 \\ 0 & -\alpha & -1 & 6 & -27 \\ \alpha & 0 & 9 & -27 & 142 \end{bmatrix}, \quad \alpha \in \mathbb{R},$$

and consider

$$(5.2) \quad \begin{aligned} f &= vGv^* \\ &= X_1^2 - X_1X_2 - X_2X_1 + 3X_2^2 - 2X_1X_2X_1 + 2X_1X_2^2X_1 \\ &\quad - X_2X_3 - X_3X_2 + 6X_3^2 + 9X_2^2X_3 + 9X_3X_2^2 - 54X_3X_2X_3 + 142X_3X_2^2X_3. \end{aligned}$$

The polynomial f is clearly sparse w.r.t. $I_1 = \{x_1, x_2\}$ and $I_2 = \{x_2, x_3\}$. Note that the matrix G is positive semidefinite iff $0.270615 \leq \alpha \leq 1.1075$, whence f is a sparse polynomial that is an SOHS.

We claim that $f \notin \Sigma\langle X(I_1) \rangle + \Sigma\langle X(I_2) \rangle$, i.e., f is not a sum of sparse hermitian squares. By the Newton chip method [BKP16, Section 2.3] only monomials in v can appear in a sum of squares decomposition of f . Further, every Gram matrix of f (with border vector v) is of the form (5.1). However, the matrix G with $\alpha = 0$ is not positive semidefinite, hence $f \notin \Sigma\langle X(I_1) \rangle + \Sigma\langle X(I_2) \rangle$.

5.1. Unconstrained Eigenvalue Optimization with Sparsity.

Let \mathbf{I} stands for the identity matrix. Given $f \in \text{Sym } \mathbb{R}\langle X \rangle$ of degree $2d$, the smallest eigenvalue of f is obtained by solving the following optimization problem

$$(5.3) \quad \lambda_{\min}(f) := \inf\{\langle f(\underline{A})\mathbf{v}, \mathbf{v} \rangle : \underline{A} \in \mathbb{S}^n, \|\mathbf{v}\| = 1\}.$$

The optimal value $\lambda_{\min}(f)$ of Problem (5.3) is the greatest lower bound on the eigenvalues of $f(\underline{A})$ over all n -tuples \underline{A} of real symmetric matrices. Problem (5.3) can be rewritten as follows:

$$(5.4) \quad \begin{aligned} \lambda_{\min}(f) &= \sup_{\lambda} \lambda \\ \text{s.t.} \quad & f(\underline{A}) - \lambda \mathbf{I} \succeq 0, \quad \forall \underline{A} \in \mathbb{S}^n, \end{aligned}$$

which is in turn equivalent to

$$(5.5) \quad \begin{aligned} \lambda_{\min,d}(f) &= \sup_{\lambda} \lambda \\ \text{s.t.} \quad & f(\underline{X}) - \lambda \in \Sigma\langle \underline{X} \rangle_d, \end{aligned}$$

as a consequence of Theorem 5.1.

The dual of SDP (5.5) is

$$(5.6) \quad \begin{aligned} L_{\text{sohs},d}(f) &= \inf_{H_L} \langle H_L, G_f \rangle \\ \text{s.t.} \quad & (H_L)_{u,v} = (H_L)_{w,z}, \quad \text{for all } u^*v = w^*z, \\ & (H_L)_{1,1} = 1, \quad H_L \succeq 0, \\ & H_L \in \mathbb{S}_{\sigma(n,d)}, \end{aligned}$$

where G_f is a Gram matrix for f (see Proposition 2.1).

One can compute $\lambda_{\min}(f)$ by solving a single SDP, either SDP (5.6) or SDP (5.5), since there is no duality gap between these two programs (see, e.g., [BKP16, Theorem 4.1]), that is, one has $L_{\text{sohs},d}(f) = \lambda_{\min,d}(f) = \lambda_{\min}(f)$.

Now, we address eigenvalue optimization for a given sparse nc polynomial $f = f_1 + \dots + f_p$ of degree $2d$, with $f_k \in \text{Sym} \mathbb{R}\langle \underline{X}(I_k) \rangle_{2d}$, for all $k = 1, \dots, p$. For all $k = 1, \dots, p$, let G_{f_k} be a Gram matrix associated to f_k . The sparse variant of SDP (5.6) is

$$(5.7) \quad \begin{aligned} L_{\text{sohs},d}^{\text{sparse}}(f) &= \inf_{H_L} \sum_{k=1}^p \langle H_L(I_k), G_{f_k} \rangle \\ \text{s.t.} \quad & (H_L(I_k))_{u,v} = (H_L(I_k))_{w,z}, \quad \text{for all } u^*v = w^*z, \\ & (H_L(I_k))_{1,1} = 1, \quad H_L(I_k) \succeq 0, \\ & H_L(I_k) \in \mathbb{S}_{\sigma(n_k,d)}, \quad k = 1, \dots, p, \end{aligned}$$

whose dual is the sparse variant of SDP (5.5):

$$(5.8) \quad \begin{aligned} \lambda_{\min,d}^{\text{sparse}}(f) &= \sup_{\lambda} \lambda \\ \text{s.t.} \quad & f - \lambda \in \Sigma\langle \underline{X}(I_1) \rangle_{2d} + \dots + \Sigma\langle \underline{X}(I_p) \rangle_{2d}, \end{aligned}$$

To prove that there is no duality gap between SDP (5.7) and SDP (5.8), we need a sparse variant of [MP05, Proposition 3.4], which says that $\Sigma\langle \underline{X} \rangle_d$ is closed in $\mathbb{R}\langle \underline{X} \rangle_{2d}$:

Proposition 5.3. *The set $\Sigma\langle \underline{X} \rangle_d^{\text{sparse}}$ is a closed convex subset of $\mathbb{R}\langle \underline{X}(I_1) \rangle_{2d} + \dots + \mathbb{R}\langle \underline{X}(I_p) \rangle_{2d}$.*

Proof. For each $k \in \{1, \dots, p\}$, we endow each $\mathbb{R}\langle \underline{X}(I_k) \rangle_{2d}$ with a norm $\|\cdot\|_k$. For each $f \in \mathbb{R}\langle \underline{X}(I_1) \rangle_{2d} + \dots + \mathbb{R}\langle \underline{X}(I_p) \rangle_{2d}$, we set

$$\|f\| := \inf \{ \|f_1\|_1 + \dots + \|f_p\|_2 : f = f_1 + \dots + f_p, f_k \in \mathbb{R}\langle \underline{X}(I_k) \rangle_{2d} \}$$

Let us consider an element $h = h_1 + \dots + h_p \in \Sigma\langle \underline{X} \rangle_d^{\text{sparse}}$, with $h_k \in \Sigma\langle \underline{X}(I_k) \rangle_d$. For each $k \in \{1, \dots, p\}$, h_k can be written as a sum of at most $\sigma(n_k, d)$ hermitian squares of degree at most $2d$ by Proposition 2.1. Define the mapping

$$\begin{aligned} \phi_k : (\mathbb{R}\langle \underline{X}(I_k) \rangle_d)^{\sigma(n_k, d)} &\rightarrow \mathbb{R}\langle \underline{X}(I_k) \rangle_{2d} \\ (h_{kj})_{j=1}^{\sigma(n_k, d)} &\mapsto \sum_{j=1}^{\sigma(n_k, d)} h_{kj}^* h_{kj}, \end{aligned}$$

and let us denote $\mathbf{h}_k = (h_{kj})_{j=1}^{\sigma(n_k, d)}$. Then, the image of the map ϕ , defined by

$$\begin{aligned} \phi : \prod_{k=1}^p (\mathbb{R}\langle \underline{X}(I_k) \rangle_d)^{\sigma(n_k, d)} &\rightarrow \mathbb{R}\langle \underline{X}(I_1) \rangle_{2d} + \dots + \mathbb{R}\langle \underline{X}(I_p) \rangle_{2d} \\ (\mathbf{h}_1, \dots, \mathbf{h}_p) &\mapsto \phi_1(\mathbf{h}_1) + \dots + \phi_p(\mathbf{h}_p), \end{aligned}$$

is equal to $\Sigma\langle \underline{X} \rangle_d^{\text{sparse}}$.

Let us define the subset $\mathcal{V} \subset \prod_{k=1}^p (\mathbb{R}\langle \underline{X}(I_k) \rangle_d)^{\sigma(n_k, d)}$ by

$$\mathcal{V} := \left\{ \mathbf{h} = \mathbf{h}_1 + \dots + \mathbf{h}_p : \sum_{k=1}^p \left\| \sum_{j=1}^{\sigma(n_k, d)} h_{kj}^* h_{kj} \right\|_k = 1 \right\}.$$

Since \mathcal{V} is compact, then $\phi(\mathcal{V})$ is also compact. Note that $0 \notin \mathcal{V}$, implying that $0 \notin \phi(\mathcal{V})$. Next, let us consider a sequence $(f^\ell)_{\ell \geq 1}$ of elements of $\Sigma\langle \underline{X} \rangle_d^{\text{sparse}}$, converging to $f \in \mathbb{R}\langle \underline{X}(I_1) \rangle_{2d} + \dots + \mathbb{R}\langle \underline{X}(I_p) \rangle_{2d}$. One can write $f^\ell = \lambda^\ell v^\ell$ for $\lambda^\ell \in \mathbb{R}^{\geq 0}$ and $v^\ell \in \phi(\mathcal{V})$. By compactness of $\phi(\mathcal{V})$, there exists a subsequence of $(v^\ell)_\ell$, (also denoted $(v^\ell)_\ell$), which converges to $v \in \phi(\mathcal{V}) \subset \Sigma\langle \underline{X} \rangle_d^{\text{sparse}}$. By definition of $\|\cdot\|$ and \mathcal{V} , one has $\|v^\ell\| \leq 1$, for all $\ell \geq 1$. Since $0 \notin \phi(\mathcal{V})$ and $\phi(\mathcal{V})$ is compact, there exists an $\varepsilon > 0$ such that $\|v^\ell\| > \varepsilon$, for all $\ell \geq 1$. Therefore, $\lambda^\ell = \frac{\|f^\ell\|}{\|v^\ell\|}$ converges to $\frac{\|f\|}{\|v\|}$, as ℓ goes to infinity. From this, we deduce that f^ℓ converges to $f = \frac{\|f\|}{\|v\|} v \in \Sigma\langle \underline{X} \rangle_d^{\text{sparse}}$, yielding the desired result. \square

From Proposition 5.3, we obtain the following theorem which does not require Assumption 2.4.

Theorem 5.4. *Let $f \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ of degree $2d$, with $f = f_1 + \dots + f_p$, $f_k \in \text{Sym } \mathbb{R}\langle \underline{X}(I_k) \rangle_{2d}$, for all $k = 1, \dots, p$. Then, one has $\lambda_{\min, d}^{\text{sparse}}(f) = L_{\text{sohs}, d}^{\text{sparse}}(f)$, i.e., there is no duality gap between SDP (5.7) and SDP (5.8).*

Proof. The strong duality is obtained exactly as for the dense case [BKP16, Theorem 4.1], and relies on the closedness of $\Sigma\langle \underline{X} \rangle_d^{\text{sparse}}$, stated in Proposition 5.3. \square

Remark 5.5. By contrast with the dense case, it is not enough to compute the solution of SDP (5.7) to obtain the optimal value $\lambda_{\min}(f)$ of the unconstrained optimization problem (5.3). However, one can still compute a certified lower bound $\lambda_{\min}^{\text{sparse}, d}(f)$ by solving a single SDP, either in the primal form (5.7) or in the dual form (5.8). Note that the related computational cost is potentially much less expensive. Indeed, SDP (5.8) involves $\sum_{k=1}^p \sigma(n_k, 2d)$ equality constraints and $\sum_{k=1}^p \sigma(n_k, d) + 1$ variables. This is in contrast with the dense version (5.5), which involves $\sigma(n, 2d)$ equality constraints and $1 + \sigma(n, d)$ variables.

5.2. Constrained Eigenvalue Optimization with Sparsity.

Here, we focus on providing lower bounds for the constrained eigenvalue optimization of nc polynomials. Given $f \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and $S := \{g_1, \dots, g_m\} \subset \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ as in (2.2), let us define $\lambda_{\min}(f, S)$ as follows:

$$(5.9) \quad \lambda_{\min}(f, S) := \inf \{ \langle f(\underline{A})\mathbf{v}, \mathbf{v} \rangle : \underline{A} \in \mathcal{D}_S^\infty, \|\mathbf{v}\| = 1 \},$$

which is, as for the unconstrained case, equivalent to

$$(5.10) \quad \begin{aligned} \lambda_{\min}(f, S) = \sup_{\lambda} \quad & \lambda \\ \text{s.t.} \quad & f(\underline{A}) - \lambda \mathbf{I} \succeq 0, \quad \forall \underline{A} \in \mathcal{D}_S^\infty. \end{aligned}$$

Let $d_j := \lceil \deg g_j / 2 \rceil$, for each $j = 1, \dots, m$ and $d := \max\{\lceil \deg f / 2 \rceil, d_1, \dots, d_m\}$. As shown in [PNA10, CKP12] (see also [HM04]), one can approximate $\lambda_{\min}(f, S)$ from below via the following hierarchy of SDP programs, indexed by $s \geq d$:

$$(5.11) \quad \begin{aligned} \lambda_s(f, S) := \sup_{\lambda} \quad & \lambda \\ \text{s.t.} \quad & f - \lambda \in M_{S, 2s}. \end{aligned}$$

The dual of SDP (5.11) is

$$(5.12) \quad \begin{aligned} L_s(f, S) := \inf_{H_L, H_{L, g_j}^\uparrow} \quad & \langle H_L, G_f \rangle \\ \text{s.t.} \quad & (H_L)_{u,v} = (H_L)_{w,z}, \quad \text{for all } u^*v = w^*z, \\ & (H_L)_{1,1} = 1, \\ & H_L \succeq 0, \quad H_L \in \mathbb{S}_{\sigma(n,s)}, \\ & H_{L, g_j}^\uparrow \succeq 0, \quad H_{L, g_j}^\uparrow \in \mathbb{S}_{\sigma(n, s-d_j)}, \quad j = 1, \dots, m. \end{aligned}$$

Under additional assumptions, this hierarchy of primal-dual SDP (5.11)-(5.12) converges to the value of the constrained eigenvalue problem. Recall that an ε -neighborhood of 0 is the set \mathcal{N}_ε defined for a given $\varepsilon > 0$ by:

$$\mathcal{N}_\varepsilon := \bigcup_{k \in \mathbb{N}} \left\{ \underline{A} := (A_1, \dots, A_n) \in \mathbb{S}_k^n : \varepsilon^2 - \sum_{i=1}^n A_i^2 \succeq 0 \right\}.$$

Lemma 5.6. *If $h \in \mathbb{R}\langle \underline{X} \rangle$ vanishes on an ε -neighborhood of 0, then $h = 0$.*

Proof. See [BKP16, Lemma 1.35]. □

Corollary 5.7. *Assume that \mathcal{D}_S contains an ε -neighborhood of 0 and that the quadratic module M_S is archimedean. Then the following holds for each $f \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle$:*

$$(5.13) \quad \lim_{s \rightarrow \infty} L_s(f, S) = \lim_{s \rightarrow \infty} \lambda_s(f, S) = \lambda_{\min}(f, S).$$

The main ingredient of the proof (see, e.g., [BKP16, Corollary 4.11]) is the nc analog of Putinar's Positivstellensatz, stated in Theorem 2.2.

Let $S \cup \{f\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and let \mathcal{D}_S be as in (2.6) with the additional quadratic constraints (2.5). Let M_S^{sparse} be as in (3.1) and let us define $M_{S,s}^{\text{sparse}}$ in the same way as the truncated quadratic module $M_{S,s}$ in (2.4). Now, let us state the sparse variant of the primal-dual hierarchy (5.11)-(5.12) of lower bounds for $\lambda_{\min}(f, S)$.

For all $s \geq d$, the sparse variant of SDP (5.12) is
(5.14)

$$\begin{aligned} L_s^{\text{sparse}}(f, S) := & \inf_{H_L, H_{L, g_j}^\uparrow} \sum_{k=1}^p \langle H_L(I_k), G_{f_k} \rangle \\ \text{s.t. } & (H_L(I_k))_{u,v} = (H_L(I_k))_{w,z}, \quad \text{for all } u^*v = w^*z, \\ & (H_L(I_k))_{1,1} = 1, \\ & H_L(I_k) \succeq 0, \quad H_L \in \mathbb{S}_{\sigma(n_k, s)}, \quad k = 1, \dots, p, \\ & H_{L, g_j}^\uparrow(I_k) \succeq 0, \\ & H_{L, g_j}^\uparrow \in \mathbb{S}_{\sigma(n_k, s-d_j)}, \quad j \in J_k, k = 1, \dots, p, \end{aligned}$$

whose dual is the sparse variant of SDP (5.11):

$$(5.15) \quad \begin{aligned} \lambda_s^{\text{sparse}}(f, S) := & \sup_{\lambda} \lambda \\ \text{s.t. } & f - \lambda \in M_{S, s}^{\text{sparse}}. \end{aligned}$$

Proposition 5.8. *Let $S \cup \{f\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$, assume that \mathcal{D}_S contains an ε -neighborhood of 0 and that \mathcal{D}_S is as in (2.6) with the additional quadratic constraints (2.5). Then SDP (5.14) admits strictly feasible solutions.*

Proof. This proof being almost the same as the one of [BKP16, Proposition 4.9] is presented for the sake of completeness. By Lemma 2.5, it is enough to build a linear map $L : \text{Sym } \mathbb{R}\langle \underline{X} \rangle_{2s} \rightarrow \mathbb{R}$ such that for all $k = 1, \dots, p$ one has:

- $L(h^*h) > 0$, for all nonzero $h \in \mathbb{R}\langle \underline{X}(I_k) \rangle_s$;
- for all $j \in J_k$, one has $L(h^*g_jh) > 0$, for all nonzero $h \in \mathbb{R}\langle \underline{X}(I_k) \rangle_{s-\lceil \deg g_j/2 \rceil}$.

Let us pick $N > s$ and let \mathcal{U} stands for the set of all $N \times N$ matrices from \mathcal{D}_S with rational entries:

$$\mathcal{U} := \{ \underline{A}^{(k)} := (A_1^{(k)}, \dots, A_1^{(r)}) : r \in \mathbb{N}, \underline{A}^{(r)} \in \mathcal{D}_S(N) \}$$

Note that this set \mathcal{U} contains a dense subset of \mathcal{N}_ε . Let us associate to $\underline{A} \in \mathcal{U}$ the linear map $L_{\underline{A}} : \text{Sym } \mathbb{R}\langle \underline{X} \rangle_{2d} \rightarrow \mathbb{R}$ defined by $L_{\underline{A}}(h) := \text{tr}(h(\underline{A}))$. From this, we define L as follows:

$$L := \sum_{r=1}^{\infty} 2^{-r} \frac{L_{\underline{A}^{(r)}}}{\|L_{\underline{A}^{(r)}}\|}.$$

Now let us fix $k \in \{1, \dots, p\}$. Obviously, one has $L(h^*h) \geq 0$, for all nonzero $h \in \mathbb{R}\langle \underline{X}(I_k) \rangle_d$. Let us suppose that $L(h^*h) = 0$ for some $h \in \mathbb{R}\langle \underline{X}(I_k) \rangle_d$. Then, one has $L_{\underline{A}^{(r)}}(h^*h) = 0 = \text{tr}(h^*(\underline{A}^{(r)})h(\underline{A}^{(r)}))$, for all $r \in \mathbb{N}$. This implies that for all $r \in \mathbb{N}$, one has $h^*(\underline{A}^{(r)})h(\underline{A}^{(r)}) = 0$, which in turn yields $h(\underline{A}^{(r)}) = 0$. Since \mathcal{U} contains a dense subset of \mathcal{N}_ε , this implies that h vanishes on a ε -neighborhood of 0. As a consequence of Lemma 5.6, one has $h = 0$.

In a similar way, we prove that if $L(h^*g_jh) = 0$ for some $h \in \mathbb{R}\langle \underline{X}(I_k) \rangle_{s-\lceil \deg g_j/2 \rceil}$, then one necessarily has $h = 0$. \square

Corollary 5.9. *Let $S \cup \{f\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$, assume that \mathcal{D}_S contains an ε -neighborhood of 0 and that \mathcal{D}_S is as in (2.6) with the additional quadratic constraints (2.5). Let Assumption 2.4 hold. Then, one has*

$$(5.16) \quad \lim_{s \rightarrow \infty} L_s^{\text{sparse}}(f, S) = \lim_{s \rightarrow \infty} \lambda_s^{\text{sparse}}(f, S) = \lambda_{\min}(f, S).$$

Proof. The proof is similar to the one in the dense case. Let us take $\lambda := \lambda_{\min}(f, S) - \varepsilon$. Then, one has $f - \lambda \succ 0$ on \mathcal{D}_S^∞ , so Theorem 3.3 implies that $f - \lambda \in M_S^{\text{sparse}}$. Hence, there exists d such that $f - \lambda \in M_{S,s}^{\text{sparse}}$, yielding a feasible solution for SDP (5.11). By Proposition 5.8, SDP (5.14) admits strictly feasible solutions. This implies in particular that there is no duality gap between SDP (5.14) and SDP (5.15), thus $L_s^{\text{sparse}}(f, S) = \lambda_s^{\text{sparse}}(f, S)$. Therefore, one has $\lambda_{\min}(f, S) \geq L_s^{\text{sparse}}(f, S) = \lambda_s^{\text{sparse}}(f, S) \geq \lambda_{\min}(f, S) - \varepsilon$, yielding the desired result. \square

As for the unconstrained case, there is no sparse variant of the ‘‘perfect’’ Positivstellensatz stated in [BKP16, §4.4] or [HKM12], for constrained eigenvalue optimization over convex nc semialgebraic sets, such as those associated either to the sparse nc ball $\mathbb{B}^{\text{sparse}} := \{1 - \sum_{i \in I_1} X_i^2, \dots, 1 - \sum_{i \in I_p} X_i^2\}$ or the nc polydisc $\mathbb{D} := \{1 - X_1^2, \dots, 1 - X_n^2\}$. Namely, for an nc polynomial f of degree $2d + 1$, computing only SDP (5.7) with optimal value $\lambda_{\min, d+1}^{\text{sparse}}(f, S)$ when $S = \mathbb{B}^{\text{sparse}}$ or $S = \mathbb{D}^{\text{sparse}}$ does not suffice to obtain the value of $\lambda_{\min}(f, S)$. This is explained in Example 5.10 below, which implies that there is no sparse variant of [BKP16, Corollary 4.18] when $S = \mathbb{B}^{\text{sparse}}$.

Example 5.10. Let us consider a randomly generated cubic polynomial $f = f_1 + f_2$ with

$$\begin{aligned} f_1 &= 4 - X_1 + 3X_2 - 3X_3 - 3X_1^2 - 7X_1X_2 + 6X_1X_3 - X_2X_1 - 5X_3X_1 + 5X_3X_2 \\ &\quad - 5X_1^3 - 3X_1^2X_3 + 4X_1X_2X_1 - 6X_1X_2X_3 + 7X_1X_3X_1 + 2X_1X_3X_2 - X_1X_3^2 \\ &\quad - X_2X_1^2 + 3X_2X_1X_2 - X_2X_1X_3 - 2X_2^3 - 5X_2^2X_3 - 4X_2X_3^2 - 5X_3X_1^2 \\ &\quad + 7X_3X_1X_2 + 6X_3X_2X_1 - 4X_3X_2X_2 - X_3^2X_1 - 2X_3^2X_2 + 7X_3^3, \\ f_2 &= -1 + 6X_2 + 5X_3 + 3X_4 - 5X_2^2 + 2X_2X_3 + 4X_2X_4 - 4X_3X_2 + X_3^2 - X_3X_4 \\ &\quad + X_4X_2 - X_4X_3 + 2X_4^2 - 7X_2^3 + 4X_2X_3^2 + 5X_2X_3X_4 - 7X_2X_4X_3 - 7X_2X_4^2 \\ &\quad + X_3X_2^2 + 6X_3X_2X_3 - 6X_3X_2X_4 - 3X_3^2X_2 - 7X_3^2X_4 + 6X_3X_4X_2 \\ &\quad - 3X_3X_4X_3 - 7X_3X_4^2 + 3X_4X_2^2 - 7X_4X_2X_3 - X_4X_2X_4 - 5X_4X_3^2 \\ &\quad + 7X_4X_3X_4 + 6X_4^2X_2 - 4X_4^3, \end{aligned}$$

and the nc polyball $S = \mathbb{B}^{\text{sparse}} = \{1 - X_1^2 - X_2^2 - X_3^2, 1 - X_2^2 - X_3^2 - X_4^2\}$ corresponding to $I_1 = \{1, 2, 3\}$ and $I_2 = \{2, 3, 4\}$. Then, one has $\lambda_2^{\text{sparse}}(f, S) \simeq -27.536 < \lambda_3^{\text{sparse}}(f, S) \simeq -27.467 \simeq \lambda_{\min, 2}(f, S) = \lambda_{\min}(f, S)$.

5.3. Extracting Optimizers. Here, we explain how to extract a pair of optimizers $(\underline{A}, \mathbf{v})$ for the eigenvalue optimization problems. when the flatness and irreducibility conditions of Theorem 4.2 hold. We apply the **SparseGNS** procedure described in Algorithm 4.6 on the optimal solution of SDP (5.7) in the unconstrained case or SDP (5.14) in the constrained case. In the unconstrained case, we have the following sparse variant of [BKP16, Proposition 4.4].

Proposition 5.11. *Given f as in Theorem 5.4, let us assume that SDP (5.7) yields an optimal solution H_L associated to $L_{\text{sohs}, d+1}^{\text{sparse}}(f)$. If the linear functional L underlying H_L satisfies the flatness (H1) and irreducibility (H2) conditions stated in Theorem 4.2, then one has*

$$\lambda_{\min}(f) = L_{\text{sohs}, d+1}^{\text{sparse}}(f) = \sum_{k=1}^p \langle H_L(I_k), G_{f_k} \rangle.$$

Proof. The first equality comes from Theorem 5.4. Let us assume that each H_L satisfies the assumptions of Theorem 4.2. Then, we obtain a tuple \underline{A} of symmetric matrices and a unit vector \mathbf{v} such that $L(f) = \langle f(\underline{A})\mathbf{v} \mid \mathbf{v} \rangle$. Since one has $L(f) = \sum_{k=1}^p \langle H_L(I_k), G_{f_k} \rangle = \lambda_{\min}(f)$, the desired result holds. \square

We can extract optimizers for the unconstrained minimal eigenvalue problem (5.3) thanks to the following algorithm.

Algorithm 5.12. *SparseEigGNS*

Input: $f \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle_{2d}$ satisfying Assumption 2.4.

1: Compute $L_{\text{sohs}, d+1}^{\text{sparse}}(f)$ by solving SDP (5.7)

2: **if** SDP (5.7) is unbounded or its optimum is not attained **then**

3: Stop

4: **end if**

5: Let H_L be an optimizer of SDP (5.7). Compute $\underline{A}, \mathbf{v} := \text{SparseGNS}(H_L)$.

Output: \underline{A} and \mathbf{v} .

In the constrained case, the next result is the sparse variant of [BKP16, Theorem 4.12] and is a direct corollary of Theorem 4.2.

Corollary 5.13. *Let $S \cup \{f\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$, assume that \mathcal{D}_S contains an ε -neighborhood of 0 and that \mathcal{D}_S is as in (2.6) with the additional quadratic constraints (2.5). Suppose Assumptions 2.4(i)-(ii) hold. Let H_L be an optimal solution of SDP (5.14) with value $L_s(f, S)$, for $s \geq d + \delta$, such that L satisfies the assumptions of Theorem 4.2. Then, there exist $r \in \mathbb{N}$, $\underline{A} \in \mathcal{D}_S(r)$ and a unit vector \mathbf{v} such that*

$$\lambda_{\min}(f, S) = \langle f(\underline{A})\mathbf{v}, \mathbf{v} \rangle = L_s(f, S).$$

Remark 5.14. As in the dense case [BKP16, Algorithm 4.2], one can provide a randomized algorithm to look for flat optimal solutions for the constrained eigenvalue problem (5.9). The underlying reason which motivates this randomized approach is work by Nie, who derives in [Nie14] a hierarchy of SDP programs, with a random objective function, that converges to a flat solution (under mild assumptions).

Example 5.15. Consider the sparse polynomial $f = f_1 + f_2$ from Example 5.10. The Hankel matrix H_L obtained when computing $\lambda_3^{\text{sparse}}$ by solving (5.14) for $s = 3$ satisfies the flatness (H1) and irreducibility (H2) conditions of Theorem 4.2. We can thus apply the SparseGNS algorithm yielding

$$A_1 = \begin{bmatrix} 0.0059 & 0.0481 & 0.1638 & 0.4570 \\ 0.0481 & -0.2583 & 0.5629 & -0.2624 \\ 0.1638 & 0.5629 & 0.3265 & -0.3734 \\ 0.4570 & -0.2624 & -0.3734 & -0.2337 \end{bmatrix}$$

$$A_2 = \begin{bmatrix} -0.3502 & 0.0080 & 0.1411 & 0.0865 \\ 0.0080 & -0.4053 & 0.2404 & -0.1649 \\ 0.1411 & 0.2404 & -0.0959 & 0.3652 \\ 0.0865 & -0.1649 & 0.3652 & 0.4117 \end{bmatrix}$$

$$A_3 = \begin{bmatrix} -0.7669 & -0.0074 & -0.1313 & -0.0805 \\ -0.0074 & -0.4715 & -0.2238 & 0.1535 \\ -0.1313 & -0.2238 & 0.0848 & -0.3400 \\ -0.0805 & 0.1535 & -0.3400 & -0.2126 \end{bmatrix}$$

$$A_4 = \begin{bmatrix} 0.3302 & -0.1839 & 0.1811 & -0.0404 \\ -0.1839 & -0.1069 & 0.5114 & -0.0570 \\ 0.1811 & 0.5114 & 0.1311 & -0.3664 \\ -0.0404 & -0.0570 & -0.3664 & 0.4440 \end{bmatrix}$$

where

$$f(\underline{A}) = \begin{bmatrix} -10.3144 & 3.9233 & -5.0836 & -7.7828 \\ 3.9233 & 1.8363 & 4.5078 & -7.5905 \\ -5.0836 & 4.5078 & -19.5827 & 13.9157 \\ -7.7828 & -7.5905 & 13.9157 & 8.3381 \end{bmatrix}$$

has minimal eigenvalue -27.4665 with unit eigenvector

$$\mathbf{v} = [0.1546 \quad -0.2507 \quad 0.8840 \quad -0.3631]^T.$$

In this case all the ranks involved were equal to four. So A_2 and A_3 were computed already from $H_L(I_1 \cap I_2)$, after an appropriate basis change A_1 (and the same A_2, A_3) was obtained from $H_L(I_1)$, and finally A_4 was computed from $H_L(I_2)$.

6. TRACE OPTIMIZATION OF NONCOMMUTATIVE SPARSE POLYNOMIALS

The aim of this section is to provide SDP relaxations allowing one to under-approximate the smallest trace of an nc polynomial on a semialgebraic set. In Section 6.1, we provide a sparse tracial representation for tracial linear functionals. In Section 6.2, we address the unconstrained trace minimization problem. As in Section 5.1, we compute a lower bound on the smallest trace via SDP. The constrained case is handled in Section 6.3, where we derive a hierarchy of lower bounds converging to the minimal trace, assuming that the quadratic module is archimedean and that RIP holds (Assumption 2.4). Most proofs are similar to the ones of eigenvalue problems addressed in Section 5, so our treatment here is more concise.

We start this section by introducing useful notations about commutators and trace zero polynomials. Given $g, h \in \mathbb{R}\langle \underline{X} \rangle$, the nc polynomial $[g, h] := gh - hg$ is called a *commutator*. Two nc polynomials $g, h \in \mathbb{R}\langle \underline{X} \rangle$ are called *cyclically equivalent* ($g \stackrel{\text{cyc}}{\sim} h$) if $g - h$ is a sum of commutators. Given $S \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ with corresponding quadratic module M_S and truncated variant $M_{S,d}$, one defines $\Theta_{S,d} := \{g \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle_{2d} : g \stackrel{\text{cyc}}{\sim} h \text{ for some } h \in M_{S,d}\}$ and $\Theta_S := \bigcup_{d \in \mathbb{N}} \Theta_{S,d}$. In this case, Θ_S stands for the *cyclic quadratic module* generated by S and $\Theta_{S,d}$ stands for the *truncated cyclic quadratic module* generated by S .

For $S \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and \mathcal{D}_S as in (2.6) with the additional quadratic constraints (2.5), let us define $\Theta_{S,d}^k := \{g \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle_{2d} : g \stackrel{\text{cyc}}{\sim} h \text{ for some } h \in M_{S,d}^k\}$, $\Theta_S^k := \bigcup_{d \in \mathbb{N}} \Theta_{S,d}^k$, for all $k = 1, \dots, p$ and the sum

$$(6.1) \quad \Theta_{S,d}^{\text{sparse}} := \Theta_{S,d}^1 + \dots + \Theta_{S,d}^p,$$

as well as $\Theta_S^{\text{sparse}} := \bigcup_{d \in \mathbb{N}} \Theta_{S,d}^{\text{sparse}}$. If S is empty, we drop the S in the above notations.

The normalized trace of a matrix $A \in \mathbb{S}_n$ is given by $\text{tr } A = \frac{1}{n} \sum_{i=1}^n a_{i,i}$. An nc polynomial $g \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ is called a *trace zero* nc polynomial if $\text{tr}(g(\underline{A})) = 0$, for all $\underline{A} \in \mathbb{S}^n$. This is equivalent to $g \stackrel{\text{cyc}}{\sim} 0$ (see e.g. [KS08, Proposition 2.3]).

We also recall the definition of the cyclic degree, which generalizes the degree of an nc polynomial. Given $\underline{\alpha} = (\alpha_1, \dots, \alpha_n) \in \mathbb{R}^n$, the $\underline{\alpha}$ -degree of a word $w \in$

\underline{X} is defined by $\deg_{\underline{\alpha}} := \sum_{i=1}^n \alpha_i d_{w,i}$, where $d_{w,i}$ denotes how many times X_i appears in w . For a given nc polynomial $g = \sum g_w w \in \mathbb{R}\langle \underline{X} \rangle$, one has $\deg_{\underline{\alpha}} g := \max_{g_w \neq 0} \deg_{\underline{\alpha}} w$. The *cyclic- $\underline{\alpha}$ -degree* of a polynomial g is given by $\text{cdeg}_{\underline{\alpha}} g := \min_{h \stackrel{\text{cyc}}{\sim} g} \deg_{\underline{\alpha}} h$. Finally, the cyclic degree is then defined by $\text{cdeg } g := \text{cdeg}_{(1,\dots,1)} g$.

6.1. Sparse Tracial Representations.

The next theorem allows one to obtain a sparse tracial representation of a tracial linear functional, under the same flatness and irreducibility conditions stated in Theorem 4.2. This is a sparse variant of [BKP16, Theorem 1.71].

Theorem 6.1. *Let $S \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle_{2d}$, and assume that the semialgebraic set \mathcal{D}_S is as in (2.6) with the additional quadratic constraints (2.5). Let Assumption 2.4(i) hold. Set $\delta := \max\{\lceil \deg(g)/2 \rceil : g \in S \cup 1\}$. Let $L : \mathbb{R}\langle \underline{X} \rangle_{2d+2\delta} \rightarrow \mathbb{R}$ be a unital tracial linear functional satisfying $L(\Theta_{S,d}^{\text{sparse}}) \subseteq \mathbb{R}^{\geq 0}$. Assume that the flatness (H1) and irreducibility (H2) conditions of Theorem 4.2 hold. Then there are finitely many n -tuples $\underline{A}^{(j)}$ of symmetric matrices in $\mathcal{D}_S(r)$ for some $r \in \mathbb{N}$, and positive scalars λ_j with $\sum_j \lambda_j = 1$, such that for all $f \in \mathbb{R}\langle \underline{X}(I_1) \rangle_{2d} + \dots + \mathbb{R}\langle \underline{X}(I_p) \rangle_{2d}$, one has:*

$$(6.2) \quad L(f) = \sum_j \lambda_j \text{tr } f(\underline{A}^{(j)}).$$

Proof. As in Theorem 4.2, we perform the finite-dimensional GNS construction to obtain a tuple $\underline{A} \in \mathcal{D}_S(r)$, for some $r \in \mathbb{N}$, and unit vector \mathbf{v} such that (6.2) holds. To obtain the tracial representation, the proof is essentially the same as the one of [BKP16, Theorem 1.71] and relies on the Wedderburn theorem, see e.g. [Lam13, Chapter 1] for more details. \square

6.2. Unconstrained Trace Optimization with Sparsity.

Given $f \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle$, the *trace-minimum* of f is obtained by solving the following optimization problem

$$(6.3) \quad \text{tr}_{\min}(f) := \inf\{\text{tr } f(\underline{A}) : \underline{A} \in \mathbb{S}^n\},$$

which is equivalent to

$$(6.4) \quad \text{tr}_{\min}(f) = \sup\{a : \text{tr}(f - a)(\underline{A}) \geq 0, \forall \underline{A} \in \mathbb{S}^n\},$$

If the cyclic degree of f is odd, then $\text{tr}_{\min}(f) = -\infty$, thus let us assume that $2d = \text{cdeg } f$. To approximate $\text{tr}_{\min}(f)$ from below, one considers the following relaxation:

$$(6.5) \quad \text{tr}_{\Theta}(f) = \sup\{a : f - a \in \Theta_d\},$$

whose dual is

$$(6.6) \quad \begin{aligned} L_{\Theta}(f) &:= \inf_{H_L} \langle H_L, G_f \rangle \\ \text{s.t. } & (H_L)_{u,v} = (H_L)_{w,z}, \quad \text{for all } u^*v \stackrel{\text{cyc}}{\sim} w^*z, \\ & (H_L)_{1,1} = 1, \\ & H_L \succeq 0, \quad H_L \in \mathbb{S}_{\sigma(n_k, d)}. \end{aligned}$$

One has $\text{tr}_{\Theta}(f) = L_{\Theta}(f) \leq \text{tr}_{\min}(f)$, where the inequality comes from [BKP16, Lemma 5.2] and the equality results from the strong duality between SDP (6.6) and SDP (6.5), see e.g. [BKP16, Theorem 5.3] for a proof. In addition, if the optimizer H_L^{opt} of SDP (6.6) satisfies the flatness condition, i.e., the linear functional

underlying H_L^{opt} is 1-flat (see Definition 2.6), then the above relaxations are exact and one has $\text{tr}_\Theta(f) = L_\Theta(f) = \text{tr}_{\min}(f)$. This exactness result is stated in [BKP16, Theorem 5.4].

For a given nc polynomial $f = f_1 + \cdots + f_p$, with $f_k \in \text{Sym } \mathbb{R}\langle \underline{X}(I_k) \rangle_{2d}$, for all $k = 1, \dots, p$, we consider the following sparse variant of SDP (6.6):

$$(6.7) \quad \begin{aligned} L_\Theta^{\text{sparse}}(f) &= \inf_{H_L} \sum_{k=1}^p \langle H_L(I_k), G_{f_k} \rangle \\ \text{s.t.} \quad & (H_L(I_k))_{u,v} = (H_L(I_k))_{w,z}, \quad \text{for all } u^*v \stackrel{\text{cyc}}{\sim} w^*z, \\ & (H_L(I_k))_{1,1} = 1, \quad H_L(I_k) \succeq 0, \\ & H_L(I_k) \in \mathbb{S}_{\sigma(n_k, d)}, \quad k = 1, \dots, p, \end{aligned}$$

whose dual is the sparse variant of SDP (6.5):

$$(6.8) \quad \begin{aligned} \text{tr}_\Theta^{\text{sparse}}(f) &= \sup_{\lambda} \lambda \\ \text{s.t.} \quad & f - \lambda \in \Theta_d^{\text{sparse}}. \end{aligned}$$

Now, we are ready to state the sparse variant of [BKP16, Theorem 5.3].

Theorem 6.2. *Let $f \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ of degree $2d$, with $f = f_1 + \cdots + f_p$, $f_k \in \text{Sym } \mathbb{R}\langle \underline{X}(I_k) \rangle_{2d}$, for all $k = 1, \dots, p$. There is no duality gap between SDP (6.7) and SDP (6.8), namely $\text{tr}_\Theta^{\text{sparse}}(f) = L_\Theta^{\text{sparse}}(f)$.*

Proof. The proof of strong duality is essentially the same as the one of Theorem 5.4. It relies on the closedness of the convex cone Θ_d^{sparse} which comes from the closedness of Σ_d^{sparse} , proved in Proposition 5.3. \square

As for unconstrained eigenvalue optimization, one can retrieve the solution of the initial trace minimization problem under the same assumptions as Theorem 4.2. This is stated in the next proposition, which is the sparse variant of [BKP16, Theorem 5.4].

Proposition 6.3. *Let f be as in Theorem 6.2, and assume that SDP (6.7) admits an optimal solution H_L . If the linear functional L underlying H_L satisfies the flatness (H1) and irreducibility (H2) conditions stated in Theorem 4.2, then*

$$\text{tr}_\Theta^{\text{sparse}}(f) = L_\Theta^{\text{sparse}}(f) = \text{tr}_{\min}(f).$$

Proof. The first equality comes from Theorem 6.2. By Theorem 6.1, there exist finitely many n -tuples of symmetric matrices $\underline{A}^{(j)}$ and positive scalars λ_j with $\sum_j \lambda_j = 1$ such that $L(f) = \sum_j \lambda_j \text{tr } f(\underline{A}^{(j)})$. Since $L(f) = \sum_{k=1}^p \langle H_L(I_k), G_{f_k} \rangle = L_\Theta^{\text{sparse}}(f)$ and $\text{tr}_{\min}(f) = \sum_j (\lambda_j \text{tr}_{\min}(f)) \leq \sum_j \lambda_j \text{tr } f(\underline{A}^{(j)}) = L(f)$, one has $\text{tr}_{\min}(f) \leq L_\Theta^{\text{sparse}}(f)$. The desired result then follows from weak duality between SDP (6.7) and SDP (6.8). \square

In practice, Proposition 6.3 allows one to derive an algorithm similar to the **SparseEigGNS** procedure (described in Algorithm 5.12) to find flat optimal solutions for the unconstrained trace problem.

6.3. Constrained Trace Optimization with Sparsity.

In this subsection, we provide the sparse tracial version of Lasserre's hierarchy to minimize the trace of a noncommutative polynomial on a semialgebraic set. Given $f \in \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and $S := \{g_1, \dots, g_m\} \subset \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ as in (2.2), let us define $\text{tr}_{\min}(f, S)$ as follows:

$$(6.9) \quad \text{tr}_{\min}(f, S) := \inf\{\text{tr } f(\underline{A}) : \underline{A} \in \mathcal{D}_S\}.$$

Since an infinite-dimensional Hilbert space does not admit a trace, we obtain lower bounds on the minimal trace by considering a particular subset of \mathcal{D}_S^∞ . This subset is obtained by restricting from the algebra of all bounded operators $\mathcal{B}(\mathcal{H})$ on a Hilbert space \mathcal{H} to finite von Neumann algebras [Tak03] of type I and type II. We introduce $\text{tr}_{\min}(f, S)^{\text{II}_1}$ as the trace-minimum of f on $\mathcal{D}_S^{\text{II}_1}$. This latter set is defined as follows (see [BKP16, Definition 1.59]):

Definition 6.4. *Let \mathcal{F} be a type-II₁-von Neumann algebra [Tak03, Chapter 5]. Let us define $\mathcal{D}_S^\mathcal{F}$ as the set of all tuples $\underline{A} = (A_1, \dots, A_n) \in \mathcal{F}^n$ making $s(\underline{A})$ a positive semidefinite operator for every $s \in S$. The von Neumann semialgebraic set $\mathcal{D}_S^{\text{II}_1}$ generated by S is defined as*

$$\mathcal{D}_S^{\text{II}_1} := \bigcup_{\mathcal{F}} \mathcal{D}_S^\mathcal{F},$$

where the union is over all type-II₁-von Neumann algebras with separable predual.

By [BKP16, Proposition 1.62], if $f \in \Theta_S$, then $\text{tr } f(\underline{A}) \geq 0$, for all $A \in \mathcal{D}_S$ and $A \in \mathcal{D}_S^{\text{II}_1}$. Since \mathcal{D}_S can be modeled by $\mathcal{D}_S^{\text{II}_1}$, one has $\text{tr}_{\min}(f, S)^{\text{II}_1} \leq \text{tr}_{\min}(f, S)$. With d being defined as in Section 5.2, one can approximate $\text{tr}_{\min}(f, S)^{\text{II}_1}$ from below via the following hierarchy of SDP programs, indexed by $s \geq d$:

$$(6.10) \quad \text{tr}_{\Theta, s}(f, S) = \sup\{a : f - a \in \Theta_{S, d}\},$$

whose dual is

$$(6.11) \quad \begin{aligned} L_{\Theta, s}(f, S) &:= \inf_{H_L} \langle H_L, G_f \rangle \\ \text{s.t. } & (H_L)_{u,v} = (H_L)_{w,z}, \quad \text{for all } u^*v \stackrel{\text{cyc}}{\sim} w^*z, \\ & (H_L)_{1,1} = 1, \\ & H_L \succeq 0, \quad H_L \in \mathbb{S}_{\sigma(n,s)}. \\ & H_{L, g_j}^\uparrow \succeq 0, \quad H_{L, g_j}^\uparrow \in \mathbb{S}_{\sigma(n, s-d_j)}, \quad j = 1, \dots, m. \end{aligned}$$

If the quadratic module M_S is archimedean, the resulting hierarchy of SDP programs provides a sequence of lower bounds $\text{tr}_{\Theta, s}(f, S)$ monotonically converging to $\text{tr}_{\min}(f, S)^{\text{II}_1}$, see e.g. [BKP16, Corollary 3.5].

Next, we present a sparse variant hierarchy of SDP programs providing a sequence of lower bounds $\text{tr}_{\Theta, s}^{\text{sparse}}(f, S)$ monotonically converging to $\text{tr}_{\min}(f, S)^{\text{II}_1}$. Let $S \cup \{f\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and let \mathcal{D}_S be as in (2.6) with the additional quadratic

constraints (2.5). Let us define the sparse variant of SDP (6.11), indexed by $s \geq d$:

$$(6.12) \quad \begin{aligned} L_{\Theta, s}^{\text{sparse}}(f, S) &= \inf_{H_L} \sum_{k=1}^p \langle H_L(I_k), G_{f_k} \rangle \\ \text{s.t.} \quad &(H_L(I_k))_{u,v} = (H_L(I_k))_{w,z}, \quad \text{for all } u^*v \stackrel{\text{cyc}}{\sim} w^*z, \\ &(H_L(I_k))_{1,1} = 1, \quad H_L(I_k) \succeq 0, \\ &H_L(I_k) \in \mathbb{S}_{\sigma(n_k, s)}, \quad k = 1, \dots, p, \\ &H_{L, g_j}^{\uparrow} \succeq 0, \quad H_{L, g_j}^{\uparrow} \in \mathbb{S}_{\sigma(n, s-d_j)}, \quad j = 1, \dots, m. \end{aligned}$$

whose dual is the sparse variant of SDP (6.10):

$$(6.13) \quad \text{tr}_{\Theta, s}^{\text{sparse}}(f, S) = \sup\{a : f - a \in \Theta_{S, d}^{\text{sparse}}\},$$

With the same conditions as the ones assumed in Proposition 5.8 for constrained eigenvalue optimization, SDP (6.12) admits strictly feasible solutions, so there is no duality gap between SDP (6.12) and SDP (6.13). The proof is the same since the constructed linear functional in Proposition 5.8 is tracial. In order to prove convergence of the hierarchy of bounds given by the SDP (6.12)-(6.13), we need the following proposition, which is the sparse variant of [BKP16, Proposition 1.63].

Proposition 6.5. *Let $S \cup \{f\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and let \mathcal{D}_S be as in (2.6) with the additional quadratic constraints (2.5). Let Assumption 2.4 hold. Then the following are equivalent:*

- (i) $\text{tr } f(\underline{A}) \geq 0$ for all $\underline{A} \in \mathcal{D}_S^{\text{II}_1}$;
- (ii) for all $\varepsilon > 0$, there exists $g \in M_S^{\text{sparse}}$ with $f + \varepsilon \stackrel{\text{cyc}}{\sim} g$.

Proof. The implication (ii) \implies (i) is trivial. For the converse implication, let us fix $\varepsilon > 0$ such that the conclusion of (ii) does not hold. By the Hahn-Banach separation theorem, there exists a linear functional $L : \text{Sym } \mathbb{R}\langle \underline{X} \rangle \rightarrow \mathbb{R}$ with $L(f + \varepsilon) \leq 0$ and $L(M_S^{\text{sparse}}) \subseteq \mathbb{R}^{\geq 0}$. As in Theorem 3.3, the GNS construction leads to operator algebras $\mathcal{A}_k, \mathcal{A}_{jk}$ for $j, k = 1, \dots, p$ and $j \neq k$, with $\mathcal{A}_{jk} \subseteq \mathcal{A}_j, \mathcal{A}_k$. However, in this case the GNS construction yields tracial states on these, whence they are all finite von Neumann algebras. Now amalgamate in the category of von Neumann algebras (cf. [VDN92]) to obtain a finite von Neumann algebra \mathcal{A} with trace τ so that $\tau(f) \leq -\varepsilon < 0$. \square

Proposition 6.5 implies the following convergence property.

Corollary 6.6. *Let $S \cup \{f\} \subseteq \text{Sym } \mathbb{R}\langle \underline{X} \rangle$ and let \mathcal{D}_S be as in (2.6) with the additional quadratic constraints (2.5). Let Assumption 2.4 hold. Then*

$$\lim_{s \rightarrow \infty} \text{tr}_{\Theta, s}^{\text{sparse}}(f, S) = \lim_{s \rightarrow \infty} L_{\Theta, s}^{\text{sparse}}(f, S) = \text{tr}_{\min}(f, S)^{\text{II}_1}.$$

Proof. By weak duality, one has $\text{tr}_{\Theta, s}^{\text{sparse}}(f, S) \leq L_{\Theta, s}^{\text{sparse}}(f, S) \leq \text{tr}_{\min}(f, S)^{\text{II}_1}$. In addition, Proposition 6.5 implies that for each $m \in \mathbb{N}$, there exists $s(m) \in \mathbb{N}$ such that $f - \text{tr}_{\min}(f, S)^{\text{II}_1} + \frac{1}{m} \in \Theta_{S, s(m)}^{\text{sparse}}$. This implies that

$$\text{tr}_{\min}(f, S)^{\text{II}_1} - \frac{1}{m} \leq \text{tr}_{\Theta, s(m)}^{\text{sparse}}(f, S),$$

yielding the desired conclusion. \square

To extract solutions of constrained trace minimization problems, we rely on the following variant of Theorem 6.1. It is, in turn, the tracial analog of Theorem 4.2.

Proposition 6.7. *Let $S \subseteq \text{Sym } \mathbb{R}\langle X \rangle_{2d}$, and assume that the semialgebraic set \mathcal{D}_S is as in (2.6) with the additional quadratic constraints (2.5). Let Assumption 2.4(i) hold. Set $\delta := \max\{\lceil \deg(g)/2 \rceil : g \in S \cup 1\}$. Let H_L be an optimal solution of SDP (6.7) with value $L_{\Theta,s}^{\text{sparse}}(f, S)$, for $s \geq d + \delta$, such that L satisfies the flatness (H1) and irreducibility (H2) conditions of Theorem 4.2. Then there are finitely many n -tuples $\underline{A}^{(j)}$ of symmetric matrices in $\mathcal{D}_S(r)$ for some $r \in \mathbb{N}$, and positive scalars λ_j with $\sum_j \lambda_j = 1$ such that*

$$L(f) = \sum_j \lambda_j \text{tr } f(\underline{A}^{(j)}).$$

In particular, one has $\text{tr}_{\min}(f, S) = \text{tr}_{\min}(f, S)^{\text{II}_1} = L_{\Theta,s}^{\text{sparse}}(f, S) = \text{tr}_{\Theta,s}^{\text{sparse}}(f, S)$.

As in the dense case [BKP16, Algorithm 5.1], one can rely on Proposition 6.7 to provide a randomized algorithm to look for flat optimal solutions for the constrained trace problem (6.9).

7. NUMERICAL EXPERIMENTS

The aim of this section is to provide experimental comparison between the bounds given by the dense relaxations (using `NCEigMin` under `NCSOStools`) and the ones produced by our sparse variants. For the sake of conciseness, we focus on minimal eigenvalue computation.

In Section 7.1 we focus on the unconstrained case. For a given nc polynomial f of degree $2d$, we compare the smallest eigenvalue $\lambda_{\min}(f) = \lambda_{\min,d}(f) = L_{\text{sohs},d}(f)$ computed via SDP (5.6) (or equivalently SDP (5.5)) with $\lambda_{\min,d}^{\text{sparse}}(f) = L_{\text{sohs},d}^{\text{sparse}}(f)$, computed via SDP (5.7) (or equivalently SDP (5.8)).

In Section 7.2 we focus on the unconstrained case. We compare the values of $\lambda_{\min,s}(f, S) = L_{\text{sohs},s}(f, S)$, obtained in the dense setting via SDP (5.12) (or equivalently SDP (5.11)), with the values of $\lambda_{\min,d}^{\text{sparse}}(f) = L_{\text{sohs},d}^{\text{sparse}}(f)$, obtained in the sparse setting via SDP (5.14) (or equivalently SDP (5.15)), for various sets of constraints S and increasing values of the relaxation order s .

The resulting algorithm, denoted by `NCEigMinSparse`, is currently implemented in `NCSOStools` [CKP11]. This software library is available within Matlab and interfaced with the SDP solver Mosek 8.1 [Mos], which turned out to yield better performance than SeDuMi 1.3 [Stu99]. All numerical results were obtained using a cluster available at the Faculty of mechanical engineering, University of Ljubljana, which has 30 TFlops computing performance. For our computations we used only one computing node which consisted of 2 Intel Xeon X5670 2,93GHz processors, each with 6 computing cores; 48 GB DDR3 memory; 500 GB hard drive. We ran Matlab in a plain (sequential) mode, without imposing any paralelization.

7.1. Unconstrained Optimization. In Table 1, we report results obtained for minimizing the eigenvalue of the nc variants of the following functions:

- The chained singular function [CGT88]:

$$f_{\text{cs}} := \sum_{i \in J} ((X_i + 10X_{i+1})^2 + 5(X_{i+2} - X_{i+3})^2 + (X_{i+1} - 2X_{i+2})^4 + 10(X_i - X_{i+3})^4),$$

TABLE 1. **NCeigMin** vs **NCeigMinSparse** for unconstrained minimal eigenvalues of the chained singular and generalized Rosenbrock functions.

f	n	NCeigMin				NCeigMinSparse			
		m_{sdp}	n_{sdp}	$\lambda_{\min,2}(f)$	time (s)	m_{sdp}	n_{sdp}	$\lambda_{\min,2}^{\text{sparse}}(f)$	time (s)
f_{cs}	4	78	169	0	0.42	78	169	0	0.37
	8	398	841	0	1.33	165	1323	0	3.69
	12	974	2025	0	4.35	298	2205	0	6.28
	16	1806	3721	0	14.29	413	3087	0	9.18
	20	2894	5929	0	52.47	537	3969	0	12.78
	24	4238	8649	0	152.17	661	4851	0	17.65
f_{gR}	10	200	400	0	0.56	95	441	0	1.39
	12	288	576	0	0.81	117	539	0	1.78
	14	392	784	0	1.12	139	637	0	2.20
	16	512	1024	0	1.46	161	735	0	2.67
	18	648	1296	0	2.15	183	833	0	3.26
	20	800	1600	0	2.92	205	931	0	4.10

where $J = \{1, 3, 4, \dots, n-3\}$ and n is a multiple of 4. In this case, one can choose $I_k = \{k, k+1, k+2, k+3\}$ for all $k = 1, \dots, n-3$ so that the associated sparsity pattern satisfies (2.7).

- The generalized Rosenbrock function [Nas84]:

$$f_{\text{gR}} := 1 + \sum_{i=1}^{n-1} \left(100(X_{i+1} - X_i^2)^2 + (1 - X_{i+1})^2 \right).$$

In this case, one can choose $I_k = \{k, k+1\}$ for all $k = 1, \dots, n-1$ so that the associated sparsity pattern satisfies (2.7).

We compute bounds on the minimal eigenvalues of $f = f_{\text{cs}}$ for each $n \in \{4, \dots, 24\}$ being a multiple of 4, and f_{gR} for even values of $n \in \{2, \dots, 20\}$. For both functions, the minimal eigenvalue is 0. We indicate in Table 1 the data related to the semi-definite programs solved by Mosek. For each value of n , m_{sdp} stands for the total number of constraints and n_{sdp} stands for the total number of variables either of the SDP program (5.6) solved to compute $\lambda_{\min}(f)$ or the SDP program (5.7) solved to compute $\lambda_{\min,2}^{\text{sparse}}(f)$. As emphasized in the columns corresponding to m_{sdp} , the size of the SDP programs can be significantly reduced after exploiting sparsity, which is consistent with Remark 5.5. While the procedure **NCeigMin** does not take sparsity into account, it relies on the Newton chip method [BKP16, §2.3] to reduce the number of variables involved in the Hankel matrix from SDP (5.6). This explains why n_{sdp} is smaller for some values of n (e.g. $n = 8$ for f_{cs}) when running **NCeigMin**. However, the sparse procedure **NCeigMinSparse** turns out to be very often more efficient to compute the minimal eigenvalue. So far, our **NCeigMinSparse** procedure is limited by the computational abilities of current SDP solvers (such as Mosek) to handle matrices with more constraints and variables than the ones obtained e.g. for the chained singular function at $n = 24$ (see the related values of m_{sdp} and n_{sdp} in the corresponding column). It turns out that exploiting the sparsity pattern yields SDP programs with significantly fewer variables than the ones obtained after

TABLE 2. NCeigMin vs NCeigMinSparse for minimal eigenvalue of the chained singular function on the nc polydisc S_{cs} .

n	NCeigMin				NCeigMinSparse			
	m_{sdp}	n_{sdp}	$\lambda_2(f_{cs}, S_{cs})$	time (s)	m_{sdp}	n_{sdp}	$\lambda_2^{\text{sparse}}(f_{cs}, S_{cs})$	time (s)
4	161	641	315.21	3.25	161	641	315.21	2.95
8	1009	6625	965.48	146.99	525	1923	965.48	4.66
12	3121	28705	1615.7	7891.6	889	3205	1615.7	7.43
16			—		1253	4487	2266.05	13.20
20			—		1617	5769	2916.32	18.50
24			—		1981	7051	3566.56	26.38

running the Newton chip method.

In the column reporting timings, we indicate the time needed to prepare *and* solve the SDP relaxation. For values of $n, d \gtrsim 8$, our current implementation in (interpreted) Matlab happens to be rather inefficient to construct the SDP problem itself, mainly because we rely on a naive nc polynomial arithmetic. To overcome this computational burden, we plan to interface NCSOSTools with a C library implementing a more sophisticated monomial arithmetic. We also emphasize that for these unconstrained problems, each function is a sum of sparse hermitian squares, thus the sparse procedure NCeigMinSparse always retrieves the same optimal value as the dense procedure NCeigMin. However, the bound computed via the sparse procedure can be a strict lower bound of the minimal eigenvalue, as shown in Example 5.2.

7.2. Constrained Optimization. In Table 2, we report results obtained for minimizing the eigenvalue of the nc chained singular function on the semialgebraic set $S_{cs} := \{1 - X_1^2, \dots, 1 - X_n^2, X_1 - 1/3, \dots, X_n - 1/3\}$ for $n \in \{4, 8, 12, 16, 20, 24\}$. Since f has degree 4, it follows from [BKP16, Corollary 4.18] that it is enough to solve SDP (5.7) with optimal value $\lambda_2(f, S_{cs})$ to compute the minimal eigenvalue $\lambda_{\min}(f, S_{cs})$. For the experiments described in Table 2, we cannot rely on the Newton chip method as in the unconstrained case. Thus the dense procedure NCeigMin suffers from a severe computational burden for $n > 10$; the symbol “—” in a column entry indicates that the calculation did not finish in a couple of hours. As already observed before for the unconstrained case, the sparse procedure NCeigMinSparse performs much better than NCeigMin. Surprisingly, NCeigMinSparse yields the same bounds as NCeigMin at the minimal relaxation order $s = 2$, for all values of $n \leq 10$.

As shown in Example 5.10, there is no guarantee to obtain the above mentioned convergence behavior in a systematic way. We consider randomly generated cubic n -variate polynomials f_{rand} satisfying Assumption 2.4 with $I_k = \{k, k+1, k+2\}$, for all $k = 1, \dots, n-2$. The corresponding nc polyball is given by $\mathbb{B}^{\text{sparse}} := \{1 - X_1^2 - X_2^2 - X_3^2, \dots, 1 - X_{n-2}^2 - X_{n-1}^2 - X_n^2\}$. In Table 3, we report results obtained for minimizing the eigenvalue of f_{rand} on $\mathbb{B}^{\text{sparse}}$, for each value of $n \in \{4, \dots, 10\}$. Here again, the sparse procedure NCeigMinSparse yields better performance than NCeigMin. Moreover, the sparse bound obtained for each $n \leq 10$ at minimal relaxation order $s = 2$ already gives an accurate approximation of the optimal bound provided by the dense procedure. We emphasize that the value of

TABLE 3. NCeigMin vs NCeigMinSparse for minimal eigenvalue of random cubic polynomials on the nc polyball $S = \mathbb{B}^{\text{sparse}}$.

n	NCeigMin				NCeigMinSparse				
	m_{sdp}	n_{sdp}	$\lambda_2(f_{\text{rand}}, S)$	time (s)	s	m_{sdp}	n_{sdp}	$\lambda_s^{\text{sparse}}(f_{\text{rand}}, S)$	time (s)
4	71	491	-53.64	3.31	2	79	370	-53.72	1.18
					3	729	3538	-53.64	12.64
6	239	2045	-142.52	26.79	2	179	740	-142.62	2.33
					3	1535	7076	-142.52	29.52
8	559	5815	-165.89	171.30	2	279	1110	-166.32	3.73
					3	2341	10614	-165.91	62.70
10	1079	13289	-199.62	857.95	2	379	1480	-200.51	5.43
					3	3147	14152	-199.66	139.22
11	1429	18985	-180.39	2111.26	2	429	1665	-180.93	6.58
					3	3550	15921	-180.40	209.73
12	-	-	-	-	2	479	1850	-385.89	7.82
					3	3953	17690	-384.87	289.12
16	-	-	-	-	2	679	2590	-344.31	15.46
					3	5565	24766	-342.15	975.43
20	-	-	-	-	2	879	3330	-504.36	31.41
					3	7177	31842	-503.02	2587.61

the third order relaxation obtained with the sparse procedure is almost equal to the optimal bound. In addition, the dense procedure cannot handle to solve the minimal order relaxation for $n > 10$, while we can always obtain a lower bound of the eigenvalue with NCeigMinSparse.

8. CONCLUSION AND PERSPECTIVES

We have presented a sparse variant of Putinar's Positivstellensatz for positive noncommutative polynomials, yielding a converging hierarchy of semidefinite relaxations for eigenvalue and trace optimization. We also designed a general algorithm to extract solutions of such sparse problems, thanks to a sparse variant of the Gelfand-Naimark-Segal construction and amalgamation properties of operator algebras. Experimental results obtained with NCSOSTools prove that one can obtain accurate lower bounds via these semidefinite relaxations in an efficient way.

An obvious direction of further research is to investigate whether and how one can benefit from sparsity exploitation in other application fields, for instance to compute certified approximations of quantum graph parameters or maximum violation bounds of Bell inequalities in quantum information theory.

We have proved that there is no sparse analog of the Helton-McCullough Sums of Squares theorem. Thus, another interesting track of research is to look for alternative representations of sparse positive polynomials, e.g., representations involving noncommutative rational functions.

Apart from sparsity, we also intend to pursue research efforts to take into account other properties of structured noncommutative polynomials, such as symmetry.

REFERENCES

- [Bar02] Alexander Barvinok. *A course in convexity*, volume 54 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2002. 7
- [BCKP13] Sabine Burgdorf, Kristijan Cafuta, Igor Klep, and Janez Povh. The tracial moment problem and trace-optimization of polynomials. *Math. Program.*, 137(1-2, Ser. A):557–578, 2013. 3
- [BEGFB94] Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*, volume 15 of *Studies in Applied Mathematics*. SIAM, Philadelphia, PA, June 1994. 2
- [BKP16] Sabine Burgdorf, Igor Klep, and Janez Povh. *Optimization of polynomials in non-commuting variables*. SpringerBriefs in Mathematics. Springer, [Cham], 2016. 3, 6, 7, 11, 16, 17, 18, 19, 20, 21, 22, 24, 25, 26, 27, 28, 29, 30
- [Bla78] Bruce E. Blackadar. Weak expectations and nuclear C^* -algebras. *Indiana Univ. Math. J.*, 27(6):1021–1026, 1978. 7
- [BMV75] Daniel Bessis, Pierre Moussa, and Matteo Villani. Monotonic converging variational approximations to the functional integrals in quantum statistical mechanics. *J. Math. Phys.*, 16(11):2318–2325, 1975. 1
- [BO13] Grégory Berhuy and Frédérique Oggier. *An introduction to central simple algebras and their applications to wireless communication*, volume 191 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2013. 12
- [BP93] Jean R. S. Blair and Barry Peyton. An introduction to chordal graphs and clique trees. In *Graph theory and sparse matrix computation*, volume 56 of *IMA Vol. Math. Appl.*, pages 1–29. Springer, New York, 1993. 6
- [CGT88] Andrew R. Conn, Nicholas I. M. Gould, and Philippe L. Toint. Testing a class of methods for solving minimization problems with simple bounds on the variables. *Math. Comp.*, 50(182):399–430, 1988. 28
- [CKP11] Kristijan Cafuta, Igor Klep, and Janez Povh. NCSOSTools: a computer algebra system for symbolic and numerical computation with noncommutative polynomials. *Optim. Methods Softw.*, 26(3):363–380, 2011. 3, 15, 28
- [CKP12] Kristijan Cafuta, Igor Klep, and Janez Povh. Constrained polynomial optimization problems with noncommuting variables. *SIAM J. Optim.*, 22(2):363–383, 2012. 3, 19
- [FKMN01] Mitsuhiro Fukuda, Masakazu Kojima, Kazuo Murota, and Kazuhide Nakata. Exploiting sparsity in semidefinite programming via matrix completion. I. General framework. *SIAM J. Optim.*, 11(3):647–674, 2000/01. 3
- [GdLL18] Sander Gribling, David de Laat, and Monique Laurent. Bounds on entanglement dimensions and quantum graph parameters via noncommutative polynomial optimization. *Math. Program.*, 170(1, Ser. B):5–42, 2018. 1
- [GdLL19] Sander Gribling, David de Laat, and Monique Laurent. Lower bounds on matrix factorization ranks via noncommutative polynomial optimization. *Found. Comput. Math.*, to appear 2019. 1
- [GLV09] Nebojsa Gvozdenovic, Monique Laurent, and Frank Vallentin. Block-diagonal semidefinite programming hierarchies for 0/1 programming. *Oper. Res. Lett.*, 37(1):27–31, 2009. 2
- [Hel02] J. William Helton. “Positive” noncommutative polynomials are sums of squares. *Ann. of Math. (2)*, 156(2):675–694, 2002. 3, 4, 16
- [HKM12] J. William Helton, Igor Klep, and Scott McCullough. The convex Positivstellensatz in a free algebra. *Adv. Math.*, 231(1):516–534, 2012. 21
- [HLL09] Didier Henrion, Jean-Bernard Lasserre, and Johan Löfberg. GloptiPoly 3: moments, optimization and semidefinite programming. *Optim. Methods Softw.*, 24(4-5):761–779, 2009. 3
- [HLS09] Didier Henrion, Jean-Bernard Lasserre, and Carlo Savorgnan. Approximate volume and integration for basic semialgebraic sets. *SIAM Rev.*, 51(4):722–743, 2009. 2
- [HM04] J. William Helton and Scott A. McCullough. A Positivstellensatz for non-commutative polynomials. *Trans. Amer. Math. Soc.*, 356(9):3721–3737, 2004. 5, 19
- [Jos16] Cédric Jozz. *Application of polynomial optimization to electricity transmission networks*. Theses, Université Pierre et Marie Curie - Paris VI, July 2016. 2
- [Kri64] Jean-Louis Krivine. Anneaux préordonnés. *J. Analyse Math.*, 12:307–326, 1964. 2

- [KS08] Igor Klep and Markus Schweighofer. Sums of Hermitian squares and the BMV conjecture. *J. Stat. Phys.*, 133(4):739–760, 2008. 1, 23
- [Lam13] Tsit-Yuen Lam. *A first course in noncommutative rings*, volume 131. Springer Science & Business Media, 2013. 24
- [Las01] Jean-Bernard Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2000/01. 2
- [Las06] Jean-Bernard Lasserre. Convergent SDP-relaxations in polynomial optimization with sparsity. *SIAM J. Optim.*, 17(3):822–843, 2006. 2, 3, 14
- [Lau09a] Monique Laurent. Matrix completion problems. In Christodoulos A. Floudas and Panos M. Pardalos, editors, *Encyclopedia of Optimization*, pages 1967–1975. Springer, 2009. 2
- [Lau09b] Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging applications of algebraic geometry*, volume 149 of *IMA Vol. Math. Appl.*, pages 157–270. Springer, New York, 2009. 2
- [Lax58] Peter D. Lax. Differential equations, difference equations and matrix theory. *Comm. Pure Appl. Math.*, 11:175–194, 1958. 1
- [LS04] Elliott H. Lieb and Robert Seiringer. Equivalent forms of the Bessis-Moussa-Villani conjecture. *J. Statist. Phys.*, 115(1-2):185–190, 2004. 1
- [LTY17] Jean-Bernard Lasserre, Kim-Chuan Toh, and Shouguang Yang. A bounded degree SOS hierarchy for polynomial optimization. *EURO J. Comput. Optim.*, 5(1-2):87–117, 2017. 2
- [Mag18] Victor Magron. Interval Enclosures of Upper Bounds of Roundoff Errors Using Semidefinite Programming. *ACM Trans. Math. Softw.*, 44(4):41:1–41:18, June 2018. 2
- [McC01] Scott McCullough. Factorization of operator-valued polynomials in several non-commuting variables. *Linear Algebra Appl.*, 326(1-3):193–203, 2001. 3, 16
- [MCD17] Victor Magron, George Constantinides, and Alastair Donaldson. Certified round-off error bounds using semidefinite programming. *ACM Trans. Math. Software*, 43(4):Art. 34, 31, 2017. 2
- [MKKK10] Kazuo Murota, Yoshihiro Kanno, Masakazu Kojima, and Sadayoshi Kojima. A numerical algorithm for block-diagonal decomposition of matrix *-algebras with application to semidefinite programming. *Jpn. J. Ind. Appl. Math.*, 27(1):125–160, 2010. 15
- [Mos] The MOSEK optimization software. <http://www.mosek.com/>. 2, 28
- [MP05] Scott McCullough and Mihai Putinar. Noncommutative sums of squares. *Pacific J. Math.*, 218(1):167–171, 2005. 17
- [Nas84] Stephen G. Nash. Newton-type minimization via the Lánczos method. *SIAM J. Numer. Anal.*, 21(4):770–788, 1984. 29
- [NFF+03] Kazuhide Nakata, Katsuki Fujisawa, Mitsuhiro Fukuda, Masakazu Kojima, and Kazuo Murota. Exploiting sparsity in semidefinite programming via matrix completion. II. Implementation and numerical results. *Math. Program.*, 95(2, Ser. B):303–327, 2003. 3
- [Nie14] Jiawang Nie. The \mathcal{A} -truncated K -moment problem. *Found. Comput. Math.*, 14(6):1243–1276, 2014. 22
- [NPA08] Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New J. Phys.*, 10(7):073013, 2008. 1
- [NT14] Tim Netzer and Andreas Thom. Hyperbolic polynomials and generalized Clifford algebras. *Discrete Comput. Geom.*, 51(4):802–814, 2014. 1
- [PNA10] Stefano Pironio, Miguel Navascués, and Antonio Acín. Convergent relaxations of polynomial optimization problems with noncommuting variables. *SIAM J. Optim.*, 20(5):2157–2180, 2010. 3, 19
- [Put93] Mihai Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana Univ. Math. J.*, 42(3):969–984, 1993. 2, 5
- [PV09] Károly F. Pál and Tamás Vértesi. Quantum bounds on Bell inequalities. *Phys. Rev. A (3)*, 79(2):022120, 12, 2009. 1, 3
- [Rez78] Bruce Reznick. Extremal PSD forms with few terms. *Duke Math. J.*, 45(2):363–374, 1978. 2

- [RTAL13] Cordian Riener, Thorsten Theobald, Lina Jansson Andrén, and Jean-Bernard Lasserre. Exploiting symmetries in SDP-relaxations for polynomial optimization. *Math. Oper. Res.*, 38(1):122–141, 2013. 2
- [SIG98] Robert E. Skelton, Tetsuya Iwasaki, and Karolos M. Grigoriadis. *A unified algebraic approach to linear control design*. The Taylor & Francis Systems and Control Book Series. Taylor & Francis, Ltd., London, 1998. 1
- [Sta13] Herbert R. Stahl. Proof of the BMV conjecture. *Acta Math.*, 211(2):255–290, 2013. 1
- [Stu99] Jos F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optim. Methods Softw.*, 11/12(1-4):625–653, 1999. 2, 28
- [Tak03] Masamichi Takesaki. *Theory of operator algebras. III*, volume 127 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin, 2003. Operator Algebras and Non-commutative Geometry, 8. 26
- [TTT03] Reha H. Tütüncü, Kim-Chuan Toh, and Michael J. Todd. Solving semidefinite-quadratic-linear programs using SDPT3. *Math. Program.*, 95(2, Ser. B):189–217, 2003. 2
- [TWLH19] Matteo Tacchi, Tillmann Weisser, Jean-Bernard Lasserre, and Didier Henrion. Exploiting sparsity for semi-algebraic set volume computation. *preprint arXiv:1902.02976*, 2019. 2
- [VDN92] Dan-Virgil Voiculescu, Kenneth J. Dykema, and Alexandru Nica. *Free random variables*, volume 1 of *CRM Monograph Series*. American Mathematical Society, Providence, RI, 1992. 27
- [Voi85] Dan-Virgil Voiculescu. Symmetries of some reduced free product C^* -algebras. In *Operator algebras and their connections with topology and ergodic theory (Busteni, 1983)*, volume 1132 of *Lecture Notes in Math.*, pages 556–588. Springer, Berlin, 1985. 7
- [Wit15] Peter Wittek. Algorithm 950: Ncpol2sdpa-sparse semidefinite programming relaxations for polynomial optimization problems of noncommuting variables. *ACM Trans. Math. Software*, 41(3):Art. 21, 12, 2015. 3
- [WKK⁺09] Hayato Waki, Sunyoung Kim, Masakazu Kojima, Masakazu Muramatsu, and Hiroshi Sugimoto. Algorithm 883: sparsePOP—a sparse semidefinite programming relaxation of polynomial optimization problems. *ACM Trans. Math. Software*, 35(2):Art. 15, 13, 2009. 2
- [WKKM06] Hayato Waki, Sunyoung Kim, Masakazu Kojima, and Masakazu Muramatsu. Sums of squares and semidefinite program relaxations for polynomial optimization problems with structured sparsity. *SIAM J. Optim.*, 17(1):218–242, 2006. 2, 3, 6
- [WLT18] Tillmann Weisser, Jean-Bernard Lasserre, and Kim-Chuan Toh. Sparse-BSOS: a bounded degree SOS hierarchy for large scale polynomial optimization with sparsity. *Math. Program. Comput.*, 10(1):1–32, 2018. 3
- [YFK03] Makoto Yamashita, Katsuki Fujisawa, and Masakazu Kojima. Implementation and evaluation of SDPA 6.0 (semidefinite programming algorithm 6.0). volume 18, pages 491–505. 2003. The Second Japanese-Sino Optimization Meeting, Part II (Kyoto, 2002). 2

IGOR KLEP: DEPARTMENT OF MATHEMATICS, UNIVERSITY OF LJUBLJANA, SLOVENIA
Email address: igor.klep@fmf.uni-lj.si

VICTOR MAGRON: LAAS-CNRS, TOULOUSE, FRANCE
Email address: vmagron@laas.fr

JANEZ POVH: FACULTY OF MECHANICAL ENGINEERING, UNIVERSITY OF LJUBLJANA, SLOVENIA
Email address: janez.povh@lecad.fs.uni-lj.si