CONSTRAINED TRACE-OPTIMIZATION OF POLYNOMIALS IN FREELY NONCOMMUTING VARIABLES

IGOR KLEP[1] AND JANEZ POVH[2]

ABSTRACT. The study of matrix inequalities in a dimension-free setting is in the realm of free real algebraic geometry (RAG). In this paper we investigate constrained trace and eigenvalue optimization of noncommutative polynomials. We present Lasserre's relaxation scheme for trace optimization based on semidefinite programming (SDP) and demonstrate its convergence properties. Finite convergence of this relaxation scheme is governed by flatness, i.e., a rank-preserving property for associated dual SDPs. If flatness is observed, then optimizers can be extracted using the Gelfand-Naimark-Segal construction and the Artin-Wedderburn theory verifying exactness of the relaxation. To enforce flatness we employ a noncommutative version of the randomization technique championed by Nie.

The implementation of these procedures in our computer algebra system NCSOStools is presented and several examples are given to illustrate our results.

1. INTRODUCTION

Free real algebraic geometry (RAG) is a branch of the booming area of free analysis that studies positivity of polynomials in freely noncommuting (nc) variables. In recent years free RAG has found many applications of which we mention only three. In [HMdOP08] the authors survey applications and connections to control, and systems engineering. Pironio, Navascués, Acín [PNA10] give applications to quantum physics and also consider computational aspects of nc sum of squares. Cimprič [Cim10] uses nc sum of squares to investigate PDEs and eigenvalues of polynomial partial differential operators.

We developed NCSOStools [CKP11] as a consequence of this recent interest in free RAG. NCSOStools is an open source Matlab toolbox for solving nc sum of squares problems using *semidefinite programming* (SDP). As a side product our toolbox implements symbolic computation with nc variables in Matlab. Readers interested in sums of squares problems for commuting polynomials are referred to one of the many existing excellent packages, such as GloptiPoly [HLL09], SOSTOOLS [PPSP05], SparsePOP [WKK+09], or YALMIP [Löf04].

Date: December 19, 2014.

2010 *Mathematics Subject Classification.* Primary 90C22, 14P10; Secondary 13J30, 47A57.

Key words and phrases. noncommutative polynomial, optimization, sum of squares, semidefinite programming, moment problem, Hankel matrix, flat extension, Matlab toolbox, real algebraic geometry, free positivity.

[1]Supported by the Marsden Fund Council of the Royal Society of New Zealand. Partially supported by the Slovenian Research Agency grants P1-0222, L1-4292 and L1-6722. Part of this research was done while the author was on leave from the University of Maribor.

[2]Supported by the Slovenian Research Agency via program P1-0383 and project L7-4119. Supported by the Creative Core FISNM-3330-13-500033 'Simulations' project funded by the European Union.

1.1. **Contribution and reader's guide.** In this article we focus on constrained *trace* optimization of nc polynomials. We also touch upon *eigenvalue* optimization of nc polynomials in Section 3, and refer the reader to [PNA10] for further details on this important topic.

We give Lasserre's relaxation scheme [Las01] for trace optimization based on nc sum of squares and semidefinite programming (SDP), and demonstrate its convergence properties (see Section 5). Finite convergence of this relaxation scheme is governed by *flatness*, i.e., a rank-preserving property for associated dual SDPs. If flatness is observed, then optimizers can be extracted using the Gelfand-Naimark-Segal (GNS) construction and the Artin-Wedderburn theory verifying and proving exactness of the relaxation. To enforce flatness we employ a noncommutative version of the randomization technique championed by Nie [Nie14]. All this is presented in Section 6.

The implementation of these procedures in our computer algebra system `NCSOStools` is presented and several examples are given to illustrate our results.

2. Preliminaries

In this section we introduce notation and terminology used throughout the paper. For unexplained terminology we refer the reader to [CKP12].

2.1. **Notation.** We shall use our standard notation for noncommutative polynomials: $\mathbb{R}\langle \underline{X}\rangle$ will denote the free algebra consisting of noncommutative polynomials on $\underline{X} = (X_1, \dots, X_n)$ endowed with the involution fixing the X_j pointwise. The space of all degree $\leq d$ polynomials is $\mathbb{R}\langle \underline{X}\rangle_d$. The free monoid on $\underline{X}$ is $\langle \underline{X}\rangle$, and $\langle \underline{X}\rangle_d := \mathbb{R}\langle \underline{X}\rangle_d \cap \langle \underline{X}\rangle$. We denote the number of words from $\langle \underline{X}\rangle_d$ by $\sigma(d)$. The column vector obtained by stacking the words from $\langle \underline{X}\rangle_d$ using the graded lexicographic order will be denoted by $W_{\sigma(d)}$. Note that $\sigma(d) = \frac{n^{d+1}-1}{n-1}$.

2.2. **Nc semialgebraic sets.** We let $\mathbb{S}_k$ denote the set of all real symmetric $k \times k$ matrices, and let $\mathbb{S}_k^+$ be the set of all positive semidefinite $k \times k$ matrices. That is, $A \in \mathbb{S}_k$ is in $\mathbb{S}_k^+$ iff all its eigenalues are nonnegative. In this case we also write $A \succeq 0$.

Definition 2.1. Fix a subset $S = \{g_1, g_2, \dots\} \subseteq \operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle$. The *matricial semialgebraic set* $\mathcal{D}_S$ associated to S is the set of all tuples $\underline{A} = (A_1, \dots, A_n) \in \mathbb{S}_k^n$ of symmetric $k \times k$ matrices for $k \in \mathbb{N}$ making $g_i(\underline{A})$ positive semidefinite for every $g_i \in S$. When considering tuples of symmetric matrices of a fixed size $k \in \mathbb{N}$, we shall use $\mathcal{D}_S(k) := \mathcal{D}_S \cap \mathbb{S}_k^n$. Likewise, the *operator semialgebraic set* $\mathcal{D}_S^\infty$ associated to S is the set of tuples $\underline{A} = (A_1, \dots, A_n) \in B(\mathscr{H})$ of bounded self-adjoint operators on a separable infinite-dimensional Hilbert space $\mathscr{H}$ (e.g. $\mathscr{H} = \ell^2(\mathbb{N})$) making $s(\underline{A})$ positive semidefinite for every $g_i \in S$.

Remark 2.2. Clearly, $\mathcal{D}_S \subseteq \mathcal{D}_S^\infty$. On the other hand, there are examples of finite $S \subseteq \operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle$ with

$$\varnothing = \mathcal{D}_S \subsetneq \mathcal{D}_S^\infty,$$

and archimedean quadratic module M_S (for a definition of (archimedean) quadratic module see e.g. [CKP12]). For concrete examples, one can start with finitely presented groups that do not admit finite-dimensional representations, and encode the defining relations of such groups. Alternately, employ the generalized Clifford algebras that admit infinite dimensional $*$-representations but no finite-dimensional representations, e.g. algebras associated to Brändén's Vamos polynomial [Brä11, NT14].

2.3. Archimedean quadratic modules and a Positivstellensatz.

The main existing result in the literature concerning nc polynomials (strictly) positive on $\mathcal{D}_S^\infty$ is due to Helton and McCullough [HM04]. It is a perfect generalization of Putinar's Positivstellensatz [Put93] for commutative polynomials.

Theorem 2.3 (Helton & McCullough [HM04, Theorem 1.2]). *Let $S \cup \{f\} \subseteq \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle$ and suppose that M_S is archimedean. If $f(A) \succ 0$ for all $A \in \mathcal{D}_S^\infty$, then $f \in M_S$.*

Remark 2.4. In general it does not suffice to test for positive definiteness of f on $\mathcal{D}_S$ (as opposed to $\mathcal{D}_S^\infty$) in Theorem 2.3; cf. Remark 2.2 above. However, if $\mathcal{D}_S$ is convex [HM04, §2], then it is by [HM12] an LMI (linear matrix inequality) domain $\mathcal{D}_L$. In this case every polynomial positive semidefinite on $\mathcal{D}_L$ admits a weighted sum of squares certificate with optimal degree bounds [HKM12].

2.4. Flatness.

Let $A \in \mathbb{R}^{s\times s}$ be a symmetric matrix. An extension of A is a symmetric matrix $\tilde{A} \in \mathbb{R}^{(s+\Delta)\times(s+\Delta)}$ of the form

$$\tilde{A} = \begin{bmatrix} A & B \\ B^t & C \end{bmatrix}$$

for some $B \in \mathbb{R}^{s\times\Delta}$ and $C \in \mathbb{R}^{\Delta\times\Delta}$.

Using Schur complements, $\tilde{A} \succeq 0$ if and only if $A \succeq 0$, and there is some Z with

$$B = AZ \quad \text{and} \quad C \succeq Z^t AZ. \tag{1}$$

An extension $\tilde{A}$ of A is *flat* if $\operatorname{rank} A = \operatorname{rank} \tilde{A}$, or, equivalently, if $B = AZ$ and $C = Z^t AZ$ for some matrix Z. For a comprehensive study of flatness in functional analysis we refer the reader to [CF96, CF98].

If $\tilde{A} \succeq 0$ we can express its deviation from flatness by computing

$$\texttt{err}_{\texttt{flat}} = \frac{\|C - Z^t AZ\|_F}{1 + \|C\|_F + \|Z^t AZ\|_F}$$

using the Frobenius norm. Here Z is as in (1); it is easy to see that $\texttt{err}_{\texttt{flat}}$ is independent of the choice of Z.

Suppose $L : \mathbb{R}\langle \underline{X} \rangle_{2d+2\delta} \to \mathbb{R}$ is a linear functional and let $\check{L} : \mathbb{R}\langle \underline{X} \rangle_{2d} \to \mathbb{R}$ denote its restriction. We can associate to L and $\check{L}$ the Hankel matrices H_L and $H_{\check{L}}$, respectively (see e.g. Definition 4.7 for tracial Hankel matrices or [CKP12] for more details). In block form,

$$H_L = \begin{bmatrix} H_{\check{L}} & B \\ B^t & C \end{bmatrix}. \tag{2}$$

If H_L is flat over $H_{\check{L}}$, we say that H_L is δ-flat. Similarly we say that L is *δ-flat*.

2.5. **Semidefinite programming (SDP).** Semidefinite programming (SDP) is a subfield of convex optimization dealing with optimization of a linear objective function subject to linear matrix inequality (LMI) constraints [WSV00, AL12].

In the last two decades SDP found widespread in combinatorial optimization [AL12], in control theory [XL08] or in (commutative and noncommutative) polynomial optimization [Las09, Lau09, PNA10] where we can construct approximation hierarchies for these problems based on semidefinite programming. Furthermore, efficient methods for solving SDPs have been developed, cf. [WSV00]. There exist several open source packages which can solve SDP problems numerically (i.e., can find solution that is sufficiently close to an optimal solution). If the problem is of medium size (i.e., the matrix variable has order less than 1000 and there are fewer than 10000 linear constraints) then these packages are based on interior point methods, while packages for larger semidefinite programs use some variant of the first order methods (see e.g. [WSV00, MPRW09]). We suggest the reader to visit web page `http://plato.asu.edu/bench.html` for a comprehensive benchmarks of optimization solvers, including SDP solvers, see also [Mit03].

3. Eigenvalue Optimization of Noncommutative Polynomials

Pironio, Navascues and Acin [PNA10] have employed Theorem 2.3 to present a noncommutative version of Lasserre's [Las01, Las09] relaxation scheme for eigenvalue optimization of nc polynomials. As in the classical case, flatness governs exactness of this scheme. That is, if the solution to the dual semidefinite program (SDP) is flat, then the obtained optimal value is indeed the minimum and one can construct optimizers [PNA10, Theorem 2]. In [CKP12, HKM12] it was shown that optimization over a convex nc semialgebraic set is equivalent to a *single* SDP, and flatness can be enforced [CKP12]. Here we revisit this theme motivated by Nie's [Nie14] fundamental results on randomization methods forcing flatness in polynomial optimization. We apply this to noncommutative optimization, and present theoretical (Theorem 3.1) and numerical evidence (Subsection 3.2) to support its effectiveness in this context.

The main problem in eigenvalue optimization of nc polynomials can be stated as follows. Given $f \in \operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle$ and a subset $S = \{g_1, g_2, \ldots\} \subseteq \operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle$, compute

$$f_\star := \inf\big\{\langle f(\underline{A})\xi, \xi\rangle \mid \underline{A} \in \mathcal{D}_S^\infty,\ \xi \text{ a unit vector}\big\}. \tag{3}$$

Hence $f_\star$ is the greatest lower bound on the eigenvalues of $f(\underline{A})$ taken over all tuples $\underline{A}$ of bounded self-adjoint operators on a separable infinite-dimensional Hilbert space (i.e., ℓ^2). That is, $(f - f_\star)(\underline{A}) \succeq 0$ for all $\underline{A} \in \mathcal{D}_S^\infty$, and $f_\star$ is the largest real number with this property.

Following [PNA10] and [CKP12] we recall the hierarchy of primal lower bounds for $f_\star$:

$$\begin{array}{rcrcl} f_\star & \geq & f_{\text{sohs}}^{(s)} & := & \sup\ \lambda \\ & & \text{s. t.} & & f - \lambda \in M_{S,s}, \end{array} \tag{SPSDP$_{\text{eig}-\min}$}$$

for $s \geq d$. Here $\deg f \leq 2d$. The corresponding hierarchy of dual problems is

$$\begin{array}{rll} L^{(s)}_{\mathrm{sohs}} & = & \inf L(f) \\ \text{s. t.} & & L : \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle_{2s} \to \mathbb{R} \quad \text{is linear} \\ & & L(1) = 1 \\ & & L(q^* q) \geq 0 \quad \text{for all } q \in \mathbb{R}\langle \underline{X} \rangle_s \\ & & L(h^* g_i h) \geq 0 \quad \text{for all } \ h \in \mathbb{R}\langle \underline{X} \rangle_{d_i}, \ g_i \in S, \end{array} \qquad (\mathrm{DSDP}_{\mathrm{eig-min}})_s$$

where we use $d_i = \lfloor s - \deg(g_i)/2 \rfloor$. Letting G_f denote a Gram matrix for f, $(\mathrm{DSDP}_{\mathrm{eig-min}})_s$ can be equivalently represented as a semidefinite programming problem (SDP):

$$\begin{array}{rll} L^{(s)}_{\mathrm{sohs}} & = & \inf \ \langle H_L, G_f \rangle \\ \text{s. t.} & & (H_L)_{u,v} = L(u^* v), \quad \text{for all } u, v \in \langle \underline{X} \rangle_s \\ & & (H_L)_{1,1} = 1, \ H_L \in \mathbb{S}^+_{\sigma(s)}, \ H_L^i \in \mathbb{S}^+_{\sigma(d_i)}, \ \forall i \\ & & (H_L^i)_{u,v} = L(u^* g_i v), \ \text{for all } u, v \in \langle \underline{X} \rangle_{d_i} \\ & & L \ \text{ linear functional on } \mathbb{R}\langle \underline{X} \rangle_{2s}. \end{array} \qquad (\mathrm{DSDP'}_{\mathrm{eig-min}})_s$$

We can prove by a standard technique that the dual problems have Slater points hence $L^{(s)}_{\mathrm{sohs}} = f^{(s)}_{\mathrm{sohs}}$, for all $s \geq d$ (cf. [CKP12, Proposition 4.4]). By Theorem 2.3,

$$\lim_{s \to \infty} f^{(s)}_{\mathrm{sohs}} = f_\star \tag{4}$$

whenever the quadratic module M_S is archimedean. We refer to [HM04, PNA10] for further details.

3.1. **Randomized algorithm.** A natural question is whether the convergence in (4) is finite. That is, does

$$f^{(s)}_{\mathrm{sohs}} = f_\star \tag{5}$$

for some s? A sufficient condition (close to being necessary) for (5) to hold is flatness of the optimizer for $(\mathrm{DSDP'}_{\mathrm{eig-min}})_s$, which also enables us to extract the optimizer, cf. [PNA10, CKP12].

Recently Nie [Nie14] presented a hierarchy of semidefinite programming problems, similar to $(\mathrm{DSDP}_{\mathrm{eig-min}})_s$, with a random objective function that under mild conditions converges to a flat solution. Motivated by his ideas we present the following algorithm:

INPUT: $f \in \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle$ *with* $\deg f \leq 2d$, $S = \{g_1, \ldots, g_r\}$, $\delta = \lceil \max_i \deg(g_i)/2 \rceil$, $\delta_{\max}$.

FOR $s = d + \delta, d + \delta + 1, \ldots, d + \delta + \delta_{\max}$,

STEP 1: *Compute* $L^{(s)}$ *– the optimal solution for* $(\mathrm{DSDP}_{\mathrm{eig-min}})_s$.

STEP 2: IF $L^{(s)}$ *is* δ*-flat* THEN *STOP.*

STEP 3: ELSE *compute* $L^{(s)}_{\mathrm{rand}}$ *– the optimal solution for* $(\mathrm{DSDP}_{\mathrm{rand}})_s$.

IF $L^{(s)}_{\mathrm{rand}}$ *is* δ*-flat* THEN *STOP.*

OUTPUT: $L^{(s)}$ *or* $L^{(s)}_{\mathrm{rand}}$.

Algorithm 1: Randomized algorithm to find flat solutions for $(\mathrm{DSDP'}_{\mathrm{eig-min}})_s$

In STEP 3 we are solving the following semidefinite program

$$\begin{array}{ll} & \inf\ \langle H_L,\, R\rangle \\ \text{s. t.} & (H_L)_{u,v} = L(u^*v), \quad \text{for all } u,v \in \langle \underline{X}\rangle_s \\ & (H_L)_{u,v} = L^{(s)}(u^*v), \quad \text{for all } u,v \in \langle \underline{X}\rangle_{s-\delta} \\ & H_L \in \mathbb{S}^+_{\sigma(s)},\ H_L^i \in \mathbb{S}^+_{\sigma(\deg(g_i))},\ \forall i \\ & (H_L^i)_{u,v} = L(u^*g_iv),\ \text{for all } u,v \in \langle \underline{X}\rangle_{\deg(g_i)} \\ & L \ \text{ linear functional on } \mathbb{R}\langle \underline{X}\rangle_{2s}. \end{array} \qquad (\mathrm{DSDP}_{\mathrm{rand}})_s$$

The objective function is *random*: we use R which is a random positive definite Gram matrix (corresponding to a random sum of hermitian squares polynomial). In our NCSOStools implementation we actually repeat the STEP 3 several times since it is cheaper to compute $(\mathrm{DSDP}_{\mathrm{rand}})_s$ multiple times than going to the next value of s. The second constraint in $(\mathrm{DSDP}_{\mathrm{rand}})_s$ implies that the solution L of this problem must coincide with $L^{(s)}$ on $\langle \underline{X}\rangle_{2(s-\delta)}$. Following Nie we expect that Algorithm 1 will often find a δ-flat extension.

Theorem 3.1. *If S is the nc ball $\{1-\sum_j X_j^2\}$ or the nc polydisc $\{1-X_1^2,\dots,1-X_n^2\}$, then Algorithm 1 always finds a 1-flat solution in the first iteration of the for loop.*

Proof. In this case we have $\delta = 1$. From [CKP12] it follows that for $s = d+1$ the optimal value of $(\mathrm{DSDP}_{\mathrm{eig-min}})_s$ equals $f_\star$ and that we can transform $L^{(d+1)}$ into a 1-flat solution. However it is not necessary that $L^{(d+1)}$ from STEP 1 is flat. If $L^{(d+1)}$ is not 1-flat, then Algorithm 1 comes to STEP 3 and computes $L^{(d+1)}_{\mathrm{rand}}$. We claim that it is always 1-flat. Let

$$H = \begin{bmatrix} \check{H} & B \\ B^t & C \end{bmatrix} \tag{6}$$

be the (Hankel) matrix, corresponding to $L^{(d+1)}_{\mathrm{rand}}$ (note that every feasible L has such a matrix representation via $H(p,q) = L(p^*q)$, see also Definition 4.7). Rows of $\check{H}$ and B are labeled by words of length $\le d$ and the rows of B^t and C by words of length $d+1$. Since $H \succeq 0$, we have $B = \check{H}W$ for some matrix W, i.e., the columns of B are in the range of the columns of $\check{H}$. Likewise, $C \succeq W^t\check{H}W$.

Write

$$H = \begin{bmatrix} \check{H} & \check{H}W \\ W^t\check{H}^t & W^t\check{H}W \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & C - W^t\check{H}W \end{bmatrix}. \tag{7}$$

The first matrix is obviously feasible for almost all constraints in $(\mathrm{DSDP}_{\mathrm{rand}})_s$ and the second is positive semidefinite. The only constraint that is not obvious is $H_L^i \in \mathbb{S}^+_{\sigma(d_i)}$, which is equivalent to $L(p^*(1-\sum_i X_i^2)p) \ge 0$ for all $p \in \mathbb{R}\langle \underline{X}\rangle_d$ in the nc ball case, and to $L(p^*(1-X_j^2)p) \ge 0$ for all $p \in \mathbb{R}\langle \underline{X}\rangle_d$ in the nc polydisc case.

Let $\tilde{L}$ be the linear functional corresponding to the first matrix on the right-hand side of (7). Let H_Δ denote the second matrix on the right-hand side of (7). Observe that $\tilde{L}$ and L coincide on words of length at most $2d$. Then for $p \in \mathbb{R}\langle \underline{X}\rangle_d$

$$\begin{aligned} \tilde{L}\big(p^*(1-X_i^2)p\big) &= \tilde{L}(p^*p) - \tilde{L}(p^*X_i^2p) \\ &= L(p^*p) - \big(L(p^*X_i^2p) - (H_\Delta)_{pX_i,pX_i}\big) \\ &= L\big(p^*(1-X_i^2)p\big) + (H_\Delta)_{pX_i,pX_i} \ge 0, \end{aligned}$$

whence $H_L^i \in \mathbb{S}^+_{\sigma(d_i)}$. (We used that $H_\Delta \succeq 0$, a consequence of $C \succeq W^t \check{H} W$.) Similar reasoning works for $1 - \sum_i X_i^2$ (or any quadratic polynomial of the form $g = 1 - \sum_i q_i^* q_i$). Therefore $\tilde{L}$ is feasible for $(\text{DSDP}_{\text{rand}})_s$.

In $(\text{DSDP}_{\text{rand}})_s$ we minimize

$$\begin{aligned} \langle H_L, R\rangle &= \langle H_{\tilde{L}}, R\rangle + \langle C - W^t \check{H} W, \hat{R}\rangle \\ &\geq \langle H_{\tilde{L}}, R\rangle. \end{aligned} \tag{8}$$

Here $\hat{R}$ is the diagonal block of R corresponding to words of length $d+1$ – the bottom right part.

Since $\tilde{L}$ is feasible for $(\text{DSDP}_{\text{rand}})_s$, the minimum of (8) is attained where the second summand is zero. Since R is positive definite this happens iff $C = W^t \check{H} W$, i.e., $L = \tilde{L}$ is 1-flat. ∎

Remark 3.2. We implemented Algorithm 1 in our open source Matlab package `NCSOStools` [CKP11] and numerical evidence corroborates Theorem 3.1. If flatness is checked by computing ranks with accuracy up to 10^{-6} then we get flat solutions in all the examples we tested. Furthermore, Algorithm 1 works very well in practice. It often returns flat solutions when S is archimedean even if it is not the nc ball or nc polydisc; however, see also Example 3.3 below. The question which archimedean S admit flat extensions is difficult.

For trace minimization we propose a similar algorithm in Section 5, but there the theoretical and practical performance is weaker.

Example 3.3. Consider an S as in Remark 2.2. Then $\mathcal{D}_S$ is empty, M_S is archimedean and $\mathcal{D}_S^\infty \neq \varnothing$. None of the dual solutions can be flat, as each flat linear functional would yield a point in $\mathcal{D}_S$.

Example 3.4. Let us consider $f = XYX$ and $S = \{1 - X^2 - Y^2\}$. We can write it as

$$f = -1 + \frac{X^2}{2} + Y^2 + (1 - X^2 - Y^2) + \frac{1}{2} X(1+Y)^2 X + \frac{1}{2} X(1 - X^2 - Y^2)X,$$

hence $f_\star \geq -1$. We use `NCSOStools`

```
>> NCvars X Y
>> f = X*Y*X;
>> [A,fA,eig_val,eig_vec]=NCeigOptRand(f,{1-X^2-Y^2},6);
```

to obtain $f_\star \geq f_{\text{sohs}}^{(2)} = -0.3849 \approx -\frac{2\sqrt{3}}{9}$. By some manual rounding we see

$$f = -\frac{2\sqrt{3}}{9} + (\sqrt[4]{\frac{4}{27}} - \sqrt[4]{\frac{3}{4}} X^2)^2 + q^* q + \frac{\sqrt{3}}{2} X(1 - X^2 - Y^2)X,$$

where $q = \sqrt[4]{\frac{1}{12}} X - \sqrt[4]{\frac{3}{4}} YX$. Therefore $f_\star = -\frac{2\sqrt{3}}{9}$ which follows also from Theorem 3.1 (the solution underlying $L_{\text{sohs}}^{(3)}$ is 1-flat).

We point out that optimum of f (considered as polynomial in commutative variables) over the unit ball in $\mathbb{R}^2$ is also $-\frac{2\sqrt{3}}{9}$.

3.2. **Numerical experiments on random polynomials.** In this section we report numerical results obtained by running Algorithm 1 on random polynomials. Random polynomials were generated using a sparse random symmetric matrix (with elements coming from a standard normal distribution) of order $\sigma(d)$ with proportion of non-zero elements 0.2, for $n = 2, 3$ and $2d = 2, 4, 6$. We called in Matlab

```
>>R=sprandn(length(W),length(W),0.2);
>>R=R+R';
>>poly = W'*R*W;
```

Here W is the vector with all monomials of order $\leq d$. We considered the TV screens $S = \{1 - X^4 - Y^4\}$ (for $n = 2$) and $S = \{1 - X^4 - Y^4 - Z^4\}$ (when $n = 3$).

For every random polynomial we run Algorithm 1 for $s = d + \delta, \ldots, 6$ as otherwise the complexity exceeds the capability of our computer (we used a laptop with four 2.4 GHz cores and 4GB RAM). This means that if $d = 3$ (we have an nc polynomial of degree $2d = 6$) then we do only one iteration of the for loop in Algorithm 1 (recall S contains an nc polynomial of degree 4, i.e., $\delta = 2$).

For every s we compute in Step 3 the functional $L^{(s)}_{\mathrm{rand}}$ and thus its associated tracial Hankel matrix

$$\hat{H} = \begin{bmatrix} H_L & B \\ B^t & C \end{bmatrix}. \tag{9}$$

We test it for δ-flatness by comparing the rank of $\hat{H}$ with the rank of its top left part H_L. We compute rank in three different ways: using Matlab functions `rank, rref` and by SVD decomposition. In all three cases we take the tolerance to be 10^{-3}. With this tolerance we noticed that in all tested (random) cases Algorithm 1 returned a flat optimal solution already after the first step, i.e., for $s = d+\delta$. Even if we set the tolerance to be $\min\{30 \cdot \mathtt{err_{flat}}, 10^{-3}\}$, we still obtain flat solutions for all random instances.

In the following table we report numerical results:

n	$2d$	# of rand inst.	% of flat sol.	average $\mathtt{err_{flat}}$
2	2	100	100 %	$1,1 \cdot 10^{-5}$
2	4	100	100 %	$< 10^{-6}$
2	6	100	100 %	$< 10^{-6}$
3	2	100	100 %	$1,6 \cdot 10^{-5}$
3	4	100	100 %	$9,9 \cdot 10^{-06}$

Table 1. Numerical results obtained by running Algorithm 1 on random nc polynomials in n variables of degree $2d$. For every n and d we generated 100 instances and computed the percentage of 2-flat solutions obtained by Algorithm 1. We see that for all generated random instances we found a 2-flat solution. The last column contains the average of $\mathtt{err_{flat}}$ over all 100 tested random instances.

4. Noncommutative Polynomials and the Trace

We next turn our attention to trace optimization. In this section we present main technical ingredients needed for the tracial Lasserre relaxation scheme presented in Section 5 below.

4.1. **Notation and terminology.** To facilitate our considerations of the trace, we need to consider a distinguished subset of $\mathcal{D}_S^\infty$ obtained by restricting our attention from the algebra of all bounded operators $B(\mathscr{H})$ on a Hilbert space $\mathscr{H}$ (which does not admit a trace if $\mathscr{H}$ is infinite dimensional) to finite von Neumann algebras [Tak03].

Let $\mathcal{F}$ be a type II_1-von Neumann algebra [Tak03, Chapter 5], and let $\mathcal{D}_S^{\mathcal{F}}$ be the $\mathcal{F}$-*semialgebraic set* generated by S; that is, $\mathcal{D}_S^{\mathcal{F}}$ consists of all tuples $\underline{A} = (A_1, \dots, A_n) \in \mathcal{F}^n$ making $s(\underline{A})$ a positive semidefinite operator for every $s \in S$. Then

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

where the union is over all type II_1-von Neumann algebras $\mathcal{F}$ with separable predual, is called the *von Neumann (vN) semialgebraic set* generated by S.

Remark 4.1. There are inclusions

$$\mathcal{D}_S \subseteq \mathcal{D}_S^{\mathrm{II}_1} \subseteq \mathcal{D}_S^\infty; \tag{10}$$

here the first is obtained via embedding matrix algebras in the hyperfinite II_1-factor $\mathcal{R}$, and for the second inclusion simply consider a separable II_1-factor as a subalgebra of $B(\mathscr{H})$.

Whether the first inclusion in (10) is "dense" in the sense that a polynomial $f \in \mathbb{R}\langle \underline{X} \rangle$ is trace-positive on $\mathcal{D}_S$ iff f is trace-positive on $\mathcal{D}_S^{\mathrm{II}_1}$ is closely related to Connes' embedding conjecture [Con76, KS08], a deep and important open problem in operator algebras. To sidestep this problem, we shall focus on values of nc polynomials on $\mathcal{D}_S^{\mathrm{II}_1}$ instead of $\mathcal{D}_S$.

4.2. **A tracial Positivstellensatz.** We next give the tracial version of Theorem 2.3. It provides the theoretical underpinning for the tracial version of Lasserre's relaxation scheme (presented in Section 5 below) used to minimize the trace of an nc polynomial.

Proposition 4.2. *Let $S \cup \{f\} \subseteq \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle$ and suppose that M_S is archimedean. Then the following are equivalent:*

(i) $\operatorname{tr} f(\underline{A}) \geq 0$ *for all* $\underline{A} \in \mathcal{D}_S^{\mathrm{II}_1}$*;*
(ii) *for all* $\varepsilon > 0$ *there exists* $g \in M_S$ *with* $f + \varepsilon \overset{\mathrm{cyc}}{\sim} g$.

Proof. Since the argument is standard, we only present a sketch of the proof. The implication (ii) $\Rightarrow$ (i) is obvious. For the converse, assume $\varepsilon > 0$ is such that the conclusion of (ii) fails. By archimedeanity of M_S, there is a tracial linear form $L : \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle \to \mathbb{R}$ with $L(f + \varepsilon) \leq 0$, $L(M_S) \subseteq \mathbb{R}_{\geq 0}$. The usual Gelfand-Naimark-Segal (GNS) construction yields bounded self-adjoint operators A_j and a tracial linear form on the algebra generated by the A_j. Its double commutant is thus a finite von Neumann algebra, where $\operatorname{tr} f(\underline{A}) \leq -\varepsilon < 0$, contradicting (i). (Note that assumption (i) implies trace positivity of f on the hyperfinite II_1-factor $\mathcal{R}$ [Tak03] and hence on all finite type I von Neumann algebras, in particular matrices.) ∎

4.3. **Truncated cyclic quadratic modules.** For notational convenience, we recall cyclic equivalence [KS08] of nc polynomials. Polynomials $f, g \in \mathbb{R}\langle \underline{X} \rangle$ are called *cyclically equivalent* ($f \overset{\text{cyc}}{\sim} g$) if $f - g$ is a sum of commutators:

$$f - g = \sum_{i=1}^{k} [p_i, q_i] = \sum_{i=1}^{k} (p_i q_i - q_i p_i) \text{ for some } k \in \mathbb{N} \text{ and } p_i, q_i \in \mathbb{R}\langle \underline{X} \rangle.$$

This notion is of interest to us because trace zero nc polynomials are exactly sums of commutators [KS08, BK09]; see also Lemma 4.4.

Given a subset $S \subseteq \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle$, we define

$$\begin{aligned} \Theta^2_{S,d} &= \{ f \in \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle \mid \exists g \in M_{S,d} : \; f \overset{\text{cyc}}{\sim} g \} \\ &= \Big\{ f \in \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle \mid f \overset{\text{cyc}}{\sim} \sum_i h_{ij}^* g_i h_{ij} \text{ for some } h_{ij} \in \mathbb{R}\langle \underline{X} \rangle, \, g_i \in S \cup \{1\}, \, \deg(h_{ij}^* g_i h_{ij}) \leq 2d \Big\}, \\ \Theta^2_S &= \bigcup_{d \in \mathbb{N}} \Theta^2_{S,d}, \end{aligned} \tag{11}$$

and call Θ^2_S the *cyclic quadratic module* generated by S, and $\Theta^2_{S,d}$ the *truncated cyclic quadratic module* generated by S. Here, $M_{S,d}$ is the truncated quadratic module generated by S, cf. [CKP12]. In the special case when $S = \{g_1, \ldots, g_r\}$ is finite, every element f of $\Theta^2_{S,d}$ is cyclically equivalent to an element of the form

$$\sum_{k=1}^{N} a_k^* a_k + \sum_{i=1}^{r} \sum_{j=1}^{N_i} b_{ij}^* \, g_i \, b_{ij} \in M_{S,d} \tag{12}$$

for some $a_k, b_{ij} \in \mathbb{R}\langle \underline{X} \rangle$ with $\deg(a_k) \leq d$ and $\deg(b_{ij}^* g_i b_{ij}) \leq 2d$. By Caratheodory's theorem on convex hulls [Bar02, Theorem I.2.3] it is possible to give the uniform bounds $N, N_i \leq 1 + \sigma(2d) = 1 + \dim \mathbb{R}\langle \underline{X} \rangle_{2d}$.

For $\varepsilon > 0$ we also introduce

$$\begin{aligned} \mathcal{N}_\varepsilon &= \bigcup_{k \in \mathbb{N}} \Big\{ \underline{A} = (A_1, \ldots, A_n) \in \mathbb{S}_k^n \mid \varepsilon^2 - \sum_{i=1}^{n} A_i^2 \succeq 0 \Big\} \\ &= \bigcup_{k \in \mathbb{N}} \Big\{ \underline{A} = (A_1, \ldots, A_n) \in \mathbb{S}_k^n \mid \Big\| \begin{bmatrix} A_1 & \cdots & A_n \end{bmatrix}^t \Big\| \leq \varepsilon \Big\}, \end{aligned} \tag{13}$$

the *nc ε-neighborhood of* 0. (Unless mentioned otherwise, all our norms are assumed to be operator norms, i.e., $\|A\| = \sup \big\{ \|Ax\| \mid \|x\| = 1 \big\}$.) We will also refer in the sequel to $\mathcal{N}_\varepsilon(N) = \mathbb{S}_N^n \bigcap \mathcal{N}_\varepsilon$.

4.4. **Vanishing nc polynomials.** The following results are consequences of the standard theory of polynomial identities, cf. [Row80]. They all essentially boil down to the well–known fact that there are no nonzero polynomial identities that hold for all sizes of (symmetric) matrices. In fact, it is enough to test on an ε-neighborhood of 0.

Lemma 4.3. *If $f \in \mathbb{R}\langle \underline{X} \rangle$ is zero on $\mathcal{N}_\varepsilon$ for some $\varepsilon > 0$, then $f = 0$.*

Proof. This follows from the following: an nc polynomial of degree $< 2d$ that vanishes on all n-tuples of symmetric matrices $\underline{A} \in \mathcal{N}_\varepsilon(N)$, for some $N \geq d$, is zero (this uses the standard multilinearization trick together with e.g. [Row80, §2.5, §1.4]). ∎

Lemma 4.4. *If $f \in \operatorname{Sym}\mathbb{R}\langle \underline{X} \rangle$ has zero trace on $\mathcal{N}_\varepsilon$ for some $\varepsilon > 0$, then f is a sum of commutators, i.e., $f \overset{\text{cyc}}{\sim} 0$.*

Proof. This is [KS08, Theorem 2.1]. Alternately, for a more algebraic approach see [BK09]. ∎

Lemma 4.5. *Suppose $f \in \mathbb{R}\langle \underline{X} \rangle$ and let $\varepsilon > 0$. If $f(\underline{A})$ is singular for all $\underline{A} \in \mathcal{N}_\varepsilon$, then $f = 0$.*

Proof. Let $\underline{A} \in \mathbb{S}_k^n$ for some $k \in \mathbb{N}$ be arbitrary. Then $p(t) = \det f(t\underline{A})$ is a real polynomial in t. By assumption it vanishes on all small enough $t > 0$. Hence $p = 0$ as every polynomial of finite degree in one real variable has only finitely many zeros. This implies $f(\underline{A})$ is singular for all $k \in \mathbb{N}$ and all $\underline{A} \in \mathbb{S}_k^n$.

Now consider the ring $\mathrm{GM}_{2^\ell}(n)$ of n symmetric $2^\ell \times 2^\ell$ generic matrices. It is a PI ring and a domain, so admits a skew field of fractions $\mathrm{UD}_{2^\ell}(n)$ [Pro76, PS76]. However, by the Cayley-Hamilton theorem, the image $\check{f}$ of f in $\mathrm{UD}_{2^\ell}(n)$ is a zero divisor, so $\check{f} = 0$, i.e., f is a polynomial identity for symmetric $2^\ell \times 2^\ell$ matrices. Since ℓ was arbitrary, this yields $f = 0$. ∎

In our subsequent analysis, we will need to deal with neighborhoods of non-scalar points $\underline{A}$. Given $\underline{A} \in \mathbb{S}_k^n$, let

$$\mathcal{B}(\underline{A}, \varepsilon) = \bigcup_{\ell \in \mathbb{N}} \{\underline{B} \in \mathbb{S}_{k\ell}^n \mid \|\underline{B} - I_\ell \otimes \underline{A}\| \leq \varepsilon\}$$

denote the *nc neighborhood* of $\underline{A}$. These are used to define topologies in free analysis [KVV14].

Proposition 4.6. *Suppose $f \in \mathbb{R}\langle \underline{X} \rangle$, $\varepsilon > 0$, and let $\underline{A} \in \mathbb{S}_{2^k}^n$. If $f(\underline{B})$ is singular for all $\ell \in \mathbb{N}$ and all $\underline{B} \in \mathcal{B}(\underline{A}, \varepsilon)(2^{k+\ell})$, then $f = 0$.*

Proof. For $\ell \in \mathbb{N}$ and $\underline{B} \in \mathbb{S}_{2^{k+\ell}}^n$ consider the univariate polynomial $\Phi_{\underline{B}}$ defined by

$$t \mapsto \det f(I_{2^\ell} \otimes \underline{A} + t\underline{B}).$$

By assumption, $\Phi_{\underline{B}}$ vanishes for all t of small absolute value. Hence by analyticity it vanishes everywhere. We can now proceed as in the proof of Lemma 4.5 to deduce f is a polynomial identity for symmetric matrices of all sizes, whence $f = 0$. ∎

4.5. **Tracial Hankel matrices.** We call a linear functional L on $\mathbb{R}\langle \underline{X} \rangle_d$ or $\mathbb{R}\langle \underline{X} \rangle$ *symmetric* if $L(f^*) = L(f)$ for all f in the domain of L.

Definition 4.7. To each symmetric linear functional $L : \mathbb{R}\langle \underline{X} \rangle_{2d} \to \mathbb{R}$ we associate a matrix H_L (called an *nc Hankel matrix*) indexed by words $u, v \in \langle \underline{X} \rangle_d$, with

$$(H_L)_{u,v} = L(u^* v). \tag{14}$$

It is easy to see that L is *positive* (i.e., $L(p^*p) \geq 0$ for all $p \in \mathbb{R}\langle \underline{X} \rangle_d$) iff $H_L \succeq 0$.

As we are interested in the trace, a crucial notion is invariance of L under cyclic equivalence, i.e., $L(f) = L(g)$ if $f \overset{\text{cyc}}{\sim} g$. Equivalently, $(H_L)_{u,v} = (H_L)_{w,z}$ whenever $u^*v \overset{\text{cyc}}{\sim} w^*z$ for $u, v, w, z \in \langle \underline{X} \rangle_d$. In this case we call L *tracial*, and H_L is a *tracial Hankel matrix*.

Definition 4.8. Given $g \in \operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle_{2d}$, we associate to L the *localizing matrix* $H_{L,g}^{\Uparrow}$ indexed by words $u, v \in \langle \underline{X}\rangle_{d-\lceil \deg(g)/2\rceil}$ with

$$(H_{L,g}^{\Uparrow})_{u,v} = L(u^* g v). \tag{15}$$

As before, $L(h^*gh) \geq 0$ for all h with $h^*gh \in \mathbb{R}\langle \underline{X}\rangle_{2d}$ iff $H_{L,g}^{\Uparrow} \succeq 0$.

We say that L is *unital* if $L(1) = 1$.

Remark 4.9. Note that a matrix H indexed by words of length $\leq d$ satisfying the *nc Hankel condition* $H_{u_1,v_1} = H_{u_2,v_2}$ whenever $u_1^* v_1 = u_2^* v_2$, gives rise to a linear functional L on $\mathbb{R}\langle \underline{X}\rangle_{2d}$ as in (14). If $H \succeq 0$, then L is symmetric and positive. Furthermore, if H is invariant under the cyclic equivalence, i.e., $H_{u,v} = H_{w,z}$ whenever $u^*v \overset{\text{cyc}}{\sim} w^*z$ for $u, v, w, z \in \langle \underline{X}\rangle_d$, then the obtained L is tracial.

4.6. **Closedness of the truncated cyclic quadratic module and a separation argument.** The following technical proposition is a variant of a Powers-Scheiderer result [PS01, §2].

Proposition 4.10. *Suppose $S = \{g_1, \dots, g_r\} \subseteq \operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle$ is such that $\mathcal{D}_S$ contains an ε-neighborhood of* 0*. Then $M_{S,d}$ is a closed convex cone in the finite dimensional real vector space* $\operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle_{2d}$*.*

For the proof of this proposition we need to isolate a (possibly) non-scalar point and its neighborhood where all the g_j are positive definite:

Lemma 4.11. *Suppose $0 \notin S = \{g_1, \dots, g_r\} \subseteq \operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle$ is such that $\mathcal{D}_S$ contains an ε-neighborhood of* 0*. Then there is an $\underline{A} \in \mathbb{S}_{2^k}^n$ and $\bar{\varepsilon} > 0$ such that all g_j are positive definite on $\mathcal{B}(\underline{A}, \bar{\varepsilon})$.*

Proof. By Proposition 4.6, we find a $\delta_1 > 0$ and $\underline{A}^1 \in \mathcal{N}_{\delta_1}(2^{k_1})$ such that $g_1(\underline{A}^1) \succ 0$. Then there is an $\varepsilon_1 > 0$ such that $g_1(\underline{B}) \succ 0$ for all $\underline{B} \in \mathcal{B}(\underline{A}^1, \varepsilon_1)$.

Now g_2 is not singular everywhere on $\mathcal{B}(\underline{A}^1, \varepsilon_1)$ by Proposition 4.6. Hence we find $\underline{A}^2 \in \mathcal{B}(\underline{A}^1, \varepsilon_1)(2^{k_2})$ with $g_2(\underline{A}^2) \succ 0$, and a corresponding $\varepsilon_2 > 0$ with $g_2|_{\mathcal{B}(\underline{A}^2,\varepsilon_2)} \succ 0$. Without loss of generality, $\mathcal{B}(\underline{A}^2, \varepsilon_2) \subseteq \mathcal{B}(\underline{A}^1, \varepsilon_1)$. We repeat this procedure for $g_3, \dots, g_r$. Finally, setting $\underline{A} = \underline{A}^r$, $\bar{\varepsilon} = \varepsilon_r$ yields the desired conclusion. ∎

Proof of Proposition 4.10. By Lemma 4.11, we find an $\bar{\varepsilon} > 0$ and $\underline{A} \in \mathbb{S}_k^n$ such that $g_j(\underline{B}) \succ 0$ for all j and all $\underline{B} \in \mathcal{B}(\underline{A}, \bar{\varepsilon})$. Using $\mathcal{B}(\underline{A}, \bar{\varepsilon})$ we norm $\mathbb{R}\langle \underline{X}\rangle_{2d}$ by

$$|||p||| := \sup\big\{\|p(\underline{B})\| \mid \underline{B} \in \mathcal{B}(\underline{A}, \bar{\varepsilon})\big\}. \tag{16}$$

Let $\delta > 0$ be a lower bound on all the $g_j(\underline{B})$ for $\underline{B} \in \mathcal{B}(\underline{A}, \bar{\varepsilon})$, i.e., $g_j(\underline{B}) - \delta I \succeq 0$ for all $\underline{B} \in \mathcal{B}(\underline{A}, \bar{\varepsilon})$.

Now the proof of the proposition follows a standard argument, and is essentially a consequence of Caratheodory's theorem on convex hulls [Bar02, Theorem I.2.3]. Suppose $(p_m)_m$ is a sequence from $M_{S,d}$ which converges to some $p \in \mathbb{R}\langle \underline{X}\rangle$ of degree at most $2d$. By Caratheodory's theorem, there is an M (at most the dimension of $\mathbb{R}\langle \underline{X}\rangle_{2d}$ plus one) such that for each m there exist nc polynomials $r_{m,i} \in \mathbb{R}\langle \underline{X}\rangle_d$ and $t_{m,i,j} \in \mathbb{R}\langle \underline{X}\rangle_d$ such that

$$p_m = \sum_{i=1}^{M} r_{m,i}^* r_{m,i} + \sum_{j=1}^{r}\sum_{i=1}^{M} t_{m,i,j}^* g_i t_{m,i,j}.$$

Since $|||p_m||| \leq N^2$ for some $N > 0$, it follows that $|||r_{m,i}||| \leq N$ and likewise $\left|\left|\left| t_{m,i,j}^* g_j t_{m,i,j} \right|\right|\right| \leq N^2$. In view of the choice of ε, δ, we obtain $|||t_{m,i,j}||| \leq \frac{1}{\sqrt{\delta}} N$ for all i, m, j. Hence for each i, j, the sequences $(r_{m,i})$ and $(t_{m,i,j})$ are bounded in m. They thus have convergent subsequences. Tracking down these subsequential limits finishes the proof. ■

Proposition 4.10 allows us to deduce the following separation result:

Corollary 4.12. *Assume $\mathcal{D}_S$ contains an ε-neighborhood of 0, and $f \in \operatorname{Sym} \mathbb{R}\langle \underline{X}\rangle_{2d} \setminus M_{S,d}$. Then there exists a linear functional $L : \mathbb{R}\langle \underline{X}\rangle_{2d} \to \mathbb{R}$ which is nonnegative on $M_{S,d}$, strictly positive on nonzero elements of $\Sigma_d^2 = M_{\varnothing,d}$ with $L(f) < 0$.*

Proof. The existence of a separating linear functional L follows from Proposition 4.10. If necessary, add a small multiple of a linear functional strictly positive on $\Sigma_d^2 \setminus \{0\}$, and the proof is complete. ■

As a consequence of Proposition 4.10, the cone $\Theta_{S,d}^2$ is closed as well. For the proof we need a preliminary result:

Lemma 4.13. *Assume $\mathcal{D}_S$ contains an ε-neighborhood of 0, and*

$$\sum_j h_j^* h_j + \sum_{i,j} r_{ij}^* g_i r_{ij} \overset{\text{cyc}}{\sim} 0. \tag{17}$$

Then $h_j = r_{ij} = 0$ for all i, j.

Proof. Let $\underline{A}$, ε be such that $g_i \succ 0$ on $\mathcal{B}(\underline{A}, \varepsilon)$ for all i. For each $\underline{B} \in \mathcal{B}(\underline{A}, \varepsilon)$ we have

$$\sum_j \operatorname{tr}\left(h_j(\underline{B})^* h_j(\underline{B})\right) + \sum_{i,j} \operatorname{tr}\left(r_{ij}(\underline{B})^* g_i(\underline{B}) r_{ij}(\underline{B})\right) = 0$$

by (17). Hence $h_j(\underline{B}) = r_{ij}(\underline{B}) = 0$. Now apply Proposition 4.6. ■

Corollary 4.14. *Suppose $S = \{g_1, \ldots, g_r\} \subseteq \operatorname{Sym} \mathbb{R}\langle \underline{X}\rangle$ and assume $\mathcal{D}_S$ contains an ε-neighborhood of 0. Then $\Theta_{S,d}^2$ is a closed convex cone in the finite dimensional real vector space $\mathbb{R}\langle \underline{X}\rangle_{2d}$. In particular, if $f \in \operatorname{Sym} \mathbb{R}\langle \underline{X}\rangle_{2d} \setminus \Theta_{S,d}^2$ then there exists a tracial linear functional $L : \mathbb{R}\langle \underline{X}\rangle_{2d} \to \mathbb{R}$ which is nonnegative on $\Theta_{S,d}^2$, positive on $\Sigma_d^2 \setminus \{0\}$ with $L(f) < 0$.*

With Lemma 4.13 at hand, the proof of this corollary is the same as that of [BK12, Lemma 4.5] so is omitted.

5. The Lasserre Relaxation Scheme for Trace-Optimization of Noncommutative Polynomials

In this section we present the tracial version of Lasserre's relaxation scheme to minimize the trace of an nc polynomial.

5.1. **Trace optimization.** Let $S \subseteq \operatorname{Sym}\mathbb{R}\langle\underline{X}\rangle$ be finite and let $f \in \operatorname{Sym}\mathbb{R}\langle\underline{X}\rangle$. We are interested in the smallest trace $f_\star \in \mathbb{R}$ the polynomial f attains on $\mathcal{D}_S$, i.e.,

$$f_\star := \inf\left\{ \operatorname{tr} f(\underline{A}) \mid \underline{A} \in \mathcal{D}_S \right\}. \tag{18}$$

Hence $f_\star$ is the greatest lower bound on the trace of $f(\underline{A})$ for tuples of symmetric matrices $\underline{A} \in \mathcal{D}_S$, i.e., $\operatorname{tr}\big((f - f_\star)(\underline{A})\big) \geq 0$ for all $\underline{A} \in \mathcal{D}_S$, and $f_\star$ is the largest real number with this property.

5.2. Θ^2_S**-relaxation.** We introduce $f_\star^{\mathrm{II}_1} \in \mathbb{R}$ as the trace-minimum of f on $\mathcal{D}_S^{\mathrm{II}_1}$. Since $\mathcal{D}_S^{\mathrm{II}_1} \supseteq \mathcal{D}_S$, we have $f_\star^{\mathrm{II}_1} \leq f_\star$. As mentioned in Remark 4.1 (see also Proposition 4.2), $f_\star^{\mathrm{II}_1}$ is more approachable than $f_\star$. In fact, in this section we shall present Lasserre's relaxation scheme producing a sequence of computable lower bounds $f_{\Theta^2}^{(s)}$ monotonically converging to $f_\star^{\mathrm{II}_1}$. Here, as always, the constraint set S is assumed to produce an archimedean quadratic module M_S.

Proposition 5.1. *Let $S \subseteq \operatorname{Sym}\mathbb{R}\langle\underline{X}\rangle$. If $f \in \Theta^2_{S,d}$, then $\operatorname{tr} f|_{\mathcal{D}_S^{\mathrm{II}_1}} \succeq 0$.*

From Proposition 5.1 we can bound $f_\star^{\mathrm{II}_1}$ from below as follows

$$f_\star^{\mathrm{II}_1} \;\geq\; f_{\Theta^2}^{(s)} \;:=\; \begin{array}{l} \sup\ \lambda \\ \text{s. t. } f - \lambda \in \Theta^2_{S,s}, \end{array} \tag{$\mathrm{SPSDP}_{\mathrm{tr-min}})_s$$

for $s \geq d$. For $s < d$, $(\mathrm{SPSDP}_{\mathrm{tr-min}})_s$ boils down to

$$f_{\Theta^2}^{(s)} \;:=\; \begin{array}{l} \sup\ \lambda \\ \text{s. t. } f - \lambda \in \Theta^2_{\varnothing,s}. \end{array}$$

which is usually infeasible (it might be feasible if the parts of f of degrees $> s$ are sums of commutators). As global trace optimization is fairly well understood [BCKP13], we shall restrict our attention to $s \geq d$. For each *fixed* s, $(\mathrm{SPSDP}_{\mathrm{tr-min}})_s$ is an SDP (see Proposition 5.4 below) and leads to the tracial version of the Lasserre relaxation scheme.

Corollary 5.2. *Let $S \subseteq \operatorname{Sym}\mathbb{R}\langle\underline{X}\rangle$, and let $f \in \operatorname{Sym}\mathbb{R}\langle\underline{X}\rangle$. If M_S is archimedean, then*

$$f_{\Theta^2}^{(s)} \underset{s\to\infty}{\longrightarrow} f_\star^{\mathrm{II}_1}. \tag{19}$$

The sequence $f_{\Theta^2}^{(s)}$ is monotone and bounded above, but the convergence in (19) is not finite in general.

Proof. This follows from Proposition 4.2. For each $m \in \mathbb{N}$, there is $s(m) \in \mathbb{N}$ with

$$f - f_\star^{\mathrm{II}_1} + \frac{1}{m} \in \Theta^2_{S,s(m)}.$$

In particular,

$$f_{\Theta^2}^{(s(m))} \geq f_\star^{\mathrm{II}_1} - \frac{1}{m}.$$

Since also

$$f_{\Theta^2}^{(s(m))} \leq f_\star^{\mathrm{II}_1},$$

we obtain

$$\lim_{s\to\infty} f_{\Theta^2}^{(s)} = \lim_{m\to\infty} f_{\Theta^2}^{(s(m))} \leq f_\star^{\mathrm{II}_1}. \qquad \blacksquare$$

Example 5.3. For a simple example with non-finite convergence, consider

$$p = (1 - X^2)(1 - Y^2) + (1 - Y^2)(1 - X^2),$$

and

$$S = \{1 - X^2, 1 - Y^2\}.$$

Then $\operatorname{tr} p|_{\mathcal{D}_S^{\mathrm{II}_1}} \geq 0$, but $p \notin \Theta_S^2$ [KS08, Example 4.3].

The advantage of $f_{\Theta^2}^{(s)}$ over $f_\star^{\mathrm{II}_1}$ is that although we are interested in the latter, there is no good procedure or algorithm for computing it. The former provides an easier accessible approximation; its computational feasibility comes from the fact that verifying whether $f \in \Theta_{S,s}^2$ is a semidefinite programming (SDP) feasibility problem, when S is finite.

5.3. **Interpreting Θ_S^2-relaxations as SDPs.**

Proposition 5.4. *Let $f = \sum_{w \in \langle \underline{X} \rangle_{2d}} f_w w \in \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle$ and $S = \{g_1, \ldots, g_r\} \subseteq \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle$ with $g_i = \sum_{w \in \langle \underline{X} \rangle_{\deg(g_i)}} g_w^i w$. Then $f \in \Theta_{S,d}^2$ if and only if there exists a positive semidefinite matrix A of order $\sigma(d)$ and positive semidefinite matrices B^i of order $\sigma(d_i)$ (recall that $d_i = \lfloor d - \deg(g_i)/2 \rfloor$) such that for all $w \in \langle \underline{X} \rangle_{2d}$,*

$$f_w = \sum_{\substack{u,v \in \langle \underline{X} \rangle_d \\ u^* v \overset{\mathrm{cyc}}{\sim} w}} A_{u,v} + \sum_i \sum_{\substack{u,v \in \langle \underline{X} \rangle_{d_i}, z \in \langle \underline{X} \rangle_{\deg(g_i)} \\ u^* z v \overset{\mathrm{cyc}}{\sim} w}} g_z^i B_{u,v}^i. \tag{20}$$

Proof. We start with the "only if" part. Suppose $f \in \Theta_{S,s}^2$, hence there exist nc polynomials $a_i = \sum_{w \in \langle \underline{X} \rangle_d} a_w^i w$ and $b_{i,j} = \sum_{w \in \langle \underline{X} \rangle_{d_i}} b_w^{i,j} w$ such that $f \overset{\mathrm{cyc}}{\sim} \sum_i a_i^* a_i + \sum_{i,j} b_{i,j}^* g_i b_{i,j}$. In particular this means that for every $w \in \langle \underline{X} \rangle_{2d}$ the following must hold:

$$\begin{aligned} f_w &= \sum_i \sum_{\substack{u,v \in \langle \underline{X} \rangle_d \\ u^* v \overset{\mathrm{cyc}}{\sim} w}} a_u^i a_v^i u^* v + \sum_{i,j} \sum_{\substack{u,v \in \langle \underline{X} \rangle_{d_i}, z \in \langle \underline{X} \rangle_{\deg(g_i)} \\ u^* z v \overset{\mathrm{cyc}}{\sim} w}} b_u^{i,j} b_v^{i,j} g_z^i u^* z v \\ &= \sum_{\substack{u,v \in \langle \underline{X} \rangle_d \\ u^* v \overset{\mathrm{cyc}}{\sim} w}} u^* v \sum_i a_u^i a_v^i + \sum_i \sum_{\substack{u,v \in \langle \underline{X} \rangle_{d_i}, z \in \langle \underline{X} \rangle_{\deg(g_i)} \\ u^* z v \overset{\mathrm{cyc}}{\sim} w}} g_z^i u^* z v \sum_j b_u^{i,j} b_v^{i,j}. \end{aligned}$$

If we define matrix A of order $\sigma(d)$ and matrices B^i of order $\sigma(d_i)$ by $A_{u,v} = \sum_i a_u^i a_v^i$ and $B_{u,v}^i = \sum_j b_u^{i,j} b_v^{i,j}$, then these matrices are positive semidefinite and satisfy (20).

To prove the "if" part we use that A and B^i are positive semidefinite, therefore we can find (column) vectors A_i and $B_{i,j}$ such that $A = \sum_i A_i A_i^t$ and $B^i = \sum_j B_{i,j} B_{i,j}^t$. These vectors yield nc polynomials $a_i = A_i^t W_{\sigma(d)}$ and $b_{i,j} = B_{i,j}^t W_{\sigma(d_i)}$, which give a certificate for $f \in \Theta_{S,s}^2$. ∎

Remark 5.5. The last part of the proof of Proposition 5.4 explains how to construct the certificate for $f \in \Theta_{S,d}^2$. First we solve semidefinite feasibility problem in the variables $A \in \mathbb{S}_{\sigma(d)}^+$, $B^i \in \mathbb{S}_{\sigma(d_i)}^+$ subject to constraints (20). Then we compute by Cholesky or eigenvalue decomposition column vectors $A_i \in \mathbb{R}^{\sigma(d)}$ and $B_{i,j} \in \mathbb{R}^{\sigma(d_i)}$ which yield desired polynomial certificates $a_i \in \mathbb{R}\langle \underline{X} \rangle_d$ and $b_{i,j} \in \mathbb{R}\langle \underline{X} \rangle_{d_i}$.

By Proposition 5.4, $(\mathrm{SPSDP}_{\mathrm{tr-min}})_s$ is an SDP. It can be explicitly presented as

$$
\begin{array}{rl}
f^{(s)}_{\Theta^2} = & \sup\ f_1 - A_{1,1} - \sum_i g_1^i B^i_{1,1} \\
\text{s. t.} & f_w = \sum_{\substack{u,v\in\langle \underline{X}\rangle_s \\ u^*v \overset{\text{cyc}}{\sim} w}} A_{u,v} + \sum_i \sum_{\substack{u,v\in\langle \underline{X}\rangle_{d_i}, z\in\langle \underline{X}\rangle_{\deg(g_i)} \\ u^* z v \overset{\text{cyc}}{\sim} w}} g^i_z B^i_{u,v} \\
& \quad \text{for all } 1\neq w\in\langle \underline{X}\rangle_{2d}, \\
& A \in \mathbb{S}^+_{\sigma(d)},\ B^i \in \mathbb{S}^+_{\sigma(d_i)},
\end{array}
\qquad (\mathrm{SPSDP'}_{\mathrm{tr-min}})_s
$$

where we use $d_i = \lfloor s - \deg(g_i)/2 \rfloor$.

5.4. The dual SDP.

Lemma 5.6. *The dual semidefinite program to* $(\mathrm{SPSDP}_{\mathrm{tr-min}})_s$ *and* $(\mathrm{SPSDP'}_{\mathrm{tr-min}})_s$ *is:*

$$
\begin{array}{rl}
L^{(s)}_{\Theta^2} = & \inf L(f) \\
\textit{s. t.} & L : \operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle_{2s} \to \mathbb{R} \quad \textit{is linear} \\
& L(1) = 1 \\
& L(pq-qp) = 0 \quad \textit{for all } p,q\in\mathbb{R}\langle \underline{X}\rangle_s \\
& \qquad\qquad \textit{with } \ pq-pq\in \operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle \\
& L(q^*q) \geq 0 \quad \textit{for all } q\in\mathbb{R}\langle \underline{X}\rangle_s \\
& L(h^*g_i h) \geq 0 \quad \textit{for all } \ i \ \textit{and all } \ h\in\mathbb{R}\langle \underline{X}\rangle_{d_i} \\
& \qquad \textit{where } d_i = \lfloor s-\deg(g_i)/2 \rfloor
\end{array}
\qquad (\mathrm{DSDP}_{\mathrm{tr-min}})_s
$$

Proof. For this proof it is beneficial to adopt a functional analytic viewpoint of $(\mathrm{SPSDP}_{\mathrm{tr-min}})_s$ and $(\mathrm{SPSDP'}_{\mathrm{tr-min}})_s$. The primal SDP is of the form

$$
\begin{array}{rl}
f^{(s)}_{\Theta^2} := & \sup\ \lambda \\
& \text{s. t. } f-\lambda\in\Theta^2_{S,s}.
\end{array}
\tag{21}
$$

We have the following chain of reasoning (recall $s \geq d$):

$$
\begin{aligned}
\sup\{\lambda \mid f-\lambda\in\Theta^2_{S,s}\} &= \sup\{\lambda \mid f-\lambda \in \overline{\Theta^2_{S,s}}\} = \\
&= \sup\{\lambda \mid \forall L\in \left(\Theta^2_{S,s}\right)^\vee :\ L(f-\lambda)\geq 0\} && (22)\\
&= \sup\{\lambda \mid \forall L\in \left(\Theta^2_{S,s}\right)^\vee \text{ with } \ L(1)=1 :\ L(f)\geq\lambda\} && (23)\\
&= \inf\{L(f) \mid L\in\left(\Theta^2_{S,s}\right)^\vee \text{ with } \ L(1)=1\}. && (24)
\end{aligned}
$$

(Here we used $\left(\Theta^2_{S,s}\right)^\vee$ to denote the set of all linear functionals $\mathbb{R}\langle \underline{X}\rangle_{2s}\to\mathbb{R}$ nonnegative on $\Theta^2_{S,s}$.) The last equality is trivial. We next give the reasoning behind the third equality. Clearly, " $\leq$ " holds since every λ feasible for the right-hand side of (22) is also feasible for the right-hand side of (23). To see the reverse inequality we consider an arbitrary λ feasible for (23). Note that $\lambda \leq f_1 = \tilde{L}(f)$, where $\tilde{L}\in\left(\Theta^2_{S,s}\right)^\vee$ maps every polynomial into its constant term. We shall prove that $L(f-\lambda)\geq 0$ for every $L\in\left(\Theta^2_{S,s}\right)^\vee$. Consider an arbitrary $L\in\left(\Theta^2_{S,s}\right)^\vee$ and define $\hat{L}=\frac{L+\varepsilon}{L(1)+\varepsilon}$ for some $\varepsilon>0$. Then $\hat{L}(1)=1$ and $\hat{L}\in\left(\Theta^2_{S,s}\right)^\vee$, therefore $\hat{L}(f-\lambda)\geq 0$, whence $L(f-\lambda)\geq\varepsilon(\lambda-1)$. Since ε was arbitrary we get $L(f-\lambda)\geq 0$.

The problem $\inf\{L(f) \mid L \in \left(\Theta^2_{S,s}\right)^\vee$ with $L(1)=1\}$ is an SDP, and this is easily seen to be equivalent to the form $(\mathrm{DSDP}_{\mathrm{tr-min}})_s$ given above. Indeed, if $L \in \left(\Theta^2_{S,s}\right)^\vee$, $L(1)=1$, then L must be nonnegative on the terms (12) and on every commutator, therefore L is feasible for the constraints in $(\mathrm{DSDP}_{\mathrm{tr-min}})_s$. ∎

Proposition 5.7. $(\mathrm{DSDP}_{\mathrm{tr-min}})_s$ *admits Slater points.*

Proof. For this it suffices to find a tracial linear map $L : \operatorname{Sym}\mathbb{R}\langle\underline{X}\rangle_{2s} \to \mathbb{R}$ satisfying $L(p^*p) > 0$ for all nonzero $p \in \mathbb{R}\langle\underline{X}\rangle_s$, and $L(h^*g_jh) > 0$ for all j and nonzero $h \in \mathbb{R}\langle\underline{X}\rangle_{d_j}$. We again exploit a variant of the fact that there are no nonzero polynomial identities that hold for all sizes of matrices, as given in Proposition 4.6.

Let $\varepsilon > 0$ and $\underline{A} \in \mathbb{S}^n_k$ be as in Proposition 4.6, and choose a countable dense subset $\mathcal{U} = \{\underline{A}^{(j)} \mid j \in \mathbb{N}\}$ of $\mathcal{B}(\underline{A},\varepsilon)$ (for instance, take all matrices from $\mathcal{B}(\underline{A},\varepsilon)$ with entries in $\mathbb{Q}$). To each $\underline{B} \in \mathcal{U}$ we associate the linear map

$$L_{\underline{B}} : \operatorname{Sym}\mathbb{R}\langle\underline{X}\rangle_{2s} \to \mathbb{R}, \qquad f \mapsto \operatorname{tr} f(\underline{B}).$$

Form

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

We claim that L is the desired linear functional.

Obviously, $L(p^*p) \geq 0$ for all $p \in \mathbb{R}\langle\underline{X}\rangle_s$. Suppose $L(p^*p) = 0$ for some $p \in \mathbb{R}\langle\underline{X}\rangle_s$. Then $L_{\underline{A}^{(j)}}(p^*p) = 0$ for all $j \in \mathbb{N}$, i.e., for all j we have $\operatorname{tr}\left(p^*(\underline{A}^{(j)})p(\underline{A}^{(j)})\right) = 0$, hence $p^*(\underline{A}^{(j)})p(\underline{A}^{(j)}) = 0$. Since $\mathcal{U}$ was dense in $\mathcal{B}(\underline{A},\varepsilon)$, by continuity it follows that p^*p vanishes on all $\mathcal{B}(\underline{A},\varepsilon)$. Proposition 4.6 implies that $p = 0$. Similarly, $L(h^*g_jh) = 0$ implies $h = 0$ for all $h \in \mathbb{R}\langle\underline{X}\rangle_{s-\lceil \deg g_j/2\rceil}$. ∎

Remark 5.8. Having Slater points for $(\mathrm{DSDP}_{\mathrm{tr-min}})_s$ is important for the clean duality theory of SDP to kick in [VB96, dK02]. In particular, there is no duality gap, so

$$L^{(s)}_{\Theta^2} = f^{(s)}_{\Theta^2}$$

and

$$L_{\Theta^2} := \lim_{s\to\infty} L^{(s)}_{\Theta^2} = f^{\mathrm{II}_1}_\star.$$

We have implemented algorithms to compute the lower bound $f^{(s)}_{\Theta^2} = L^{(s)}_{\Theta^2}$ for $f^{\mathrm{II}_1}_\star$ and $f_\star$ in our open source toolbox `NCSOStools`. We are solving the semidefinite program $(\mathrm{DSDP}_{\mathrm{tr-min}})_s$ using one of the following SDP solvers: SDPA [YFK03], SDPT3 [TTT99] or SeDuMi [Stu99]). We demonstrate it on a few examples.

Example 5.9. We firstly demonstrate our software for set $S = \{1-X^2, 1-Y^2\}$ with the polynomial $p = (1-X^2)(1-Y^2)+(1-Y^2)(1-X^2)$ from Example 5.3, and the noncommutative version of the Motzkin polynomial,

$$q = XY^4X + YX^4Y - 3XY^2X + 1.$$

It is obvious (see Example 5.3 and [KS08, Example 4.3]) that $p^{\mathrm{II}_1}_\star = p_\star = 0$. Similarly, $q^{\mathrm{II}_1}_\star = q_\star = 0$ (see [KS08, Example 4.4]). We define these polynomials as follows.

```
>> NCvars X Y
>> S = {1 - X^2, 1 - Y^2};
>> p = (1-X^2)*(1-Y^2)+(1-Y^2)*(1-X^2);
>> q = X*Y^4*X+Y*X^4*Y-3*X*Y^2*X+1;
```

To compute the sequence of lower bounds $p_{\Theta^2}^{(s)}$ for $p_\star^{\mathrm{II}_1}$ we call

```
>> [opt,decom_sohs,decom_S,base,SDP_data,Z,Zg,H,Hg] = NCtraceOpt(p,S,2*s);
```

with $s = 2, 3, 4, 5$. Similarly we obtain bounds for q. Results are reported in Table 2.

s	$p_{\Theta^2}^{(s)}$	$q_{\Theta^2}^{(s)}$
2	-0.2500	n.d.
3	-0.0178	0
4	-0.0031	0
5	-0.0010	0

Table 2. Lower bounds $f_{\Theta^2}^{(s)}$ for p and q over $S = \{1 - X^2, 1 - Y^2\}$

We can see that the sequence of bounds for p increases and does not reach the limit $p_\star^{\mathrm{II}_1}$ when $s \leq 5$. Actually, it never reaches $p_\star^{\mathrm{II}_1}$; see Example 5.3. On the other hand, the sequence of bounds for q is finite and reaches the optimal value already for $s = 3$ ($q_{\Theta^2}^{(2)}$ is not defined).

Example 5.10. Let p, q be as in the previous example, and let $r = XYX$. Let us define $S = \{1 - X, 1 - Y, 1 + X, 1 + Y\}$. The resulting sequences from the relaxation are in Table 3 and show that there is again no convergence in the first four steps for p, while for q we get convergence at $s = 4$ and for r we get the optimal value immediately (at $s = 2$).

s	$p_{\Theta^2}^{(s)}$	$q_{\Theta^2}^{(s)}$	$r_{\Theta^2}^{(s)}$
2	-2.0000	n.d.	-1.0000
3	-0.2500	-0.0261	-1.0000
4	-0.0178	0.0000	-1.0000
5	-0.0031	0.0000	-1.0000

Table 3. Lower bounds $p_{\Theta^2}^{(s)}$, $q_{\Theta^2}^{(s)}$ and $r_{\Theta^2}^{(s)}$ over $S = \{1 - X, 1 - Y, 1 + X, 1 + Y\}$

(Note that the bounds $p_{\Theta^2}^{(s)}$ are equal to bound $p_{\Theta^2}^{(s)}$ from Table 2.) To compute e.g. $p_{\Theta^2}^{(5)}$ we need to solve $(\mathrm{DSDP}_{\mathrm{tr-min}})_s$ which has 3739 linear constraints and 5 positive semidefinite constraints with matrix variables of sizes 63, 31, 31, 31, 31.

Example 5.11. Consider $f = p^*q + q^*p$, where $p = XY$ and $q = 1 + X(Y - 2) + Y(X - 2)$, and $S = \{4 - X^2, 4 - Y^2\}$. If we use NCSOStools and call

```
>> NCvars X Y
>> p = X*Y;Q=1+X*(Y-2)+Y*(X-2);f=p'*q+q'*p;
>> [opt_2,decom_sohs,decom_S,base,SDP_data,Z,Zg,H,Hg,decom_err] = NCtraceOpt(f,{4-X^2,4-Y^2},4);
>> [opt_3,decom_sohs,decom_S,base,SDP_data,Z,Zg,H,Hg,decom_err] = NCtraceOpt(f,{4-X^2,4-Y^2},6);
>> [opt_4,decom_sohs,decom_S,base,SDP_data,Z,Zg,H,Hg,decom_err] = NCtraceOpt(f,{4-X^2,4-Y^2},8);
```

we obtain `opt_2` $= f_{\Theta^2}^{(2)} = -8$ and `opt_3` $= f_{\Theta^2}^{(3)} = -5.2165$ and later $f_{\Theta^2}^{(s)} = -5.2165$ for $s \geq 3$ (this follows from the fact that the optimal solution underlying $L_{\Theta^2}^{(3)}$ is 1-flat, see Section 6). It is easy to see that the minimum of f on $\mathcal{D}_S \cap \mathbb{R}^2 = [-2, 2]^2$ is -4.5.

6. Flatness and Extracting the Optimizers

In this section we assume $S \subseteq \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle_{2\delta}$ is finite, and $f \in \operatorname{Sym} \mathbb{R}\langle \underline{X} \rangle_{2d}$. Let M_S be archimedean. In this case $\mathcal{D}_S^{\mathrm{II}_1}$ is bounded and hence $f_\star^{\mathrm{II}_1} > -\infty$. Since M_S is archimedean, for s big enough, $(\mathrm{SPSDP}_{\mathrm{tr-min}})_s$ will be feasible.

Like in constrained eigenvalue optimization (cf. [PNA10, CKP12]) flatness is a sufficient condition for finite convergence of the bounds $f_{\Theta^2}^{(s)} = L_{\Theta^2}^{(s)}$, i.e., exactness of trace optimization, and also enables extraction of the minimizers.

6.1. Extract the optimizers.

We first recall a variant of the flatness theorem adapted to the tracial setting.

Proposition 6.1. *Suppose $L^{(d+k)} : \mathbb{R}\langle \underline{X} \rangle_{2d+2k} \to \mathbb{R}$ is an optimal solution of $(\mathrm{DSDP}_{\mathrm{tr-min}})_s$ for $s = d + k$ that is δ-flat for some $k \geq \delta$. Then there are finitely many n-tuples $\underline{A}^{(j)}$ of symmetric matrices in $\mathcal{D}_S(t)$ for some $t < 4\sigma(d)$ and positive scalars $\lambda_j > 0$ with $\sum_j \lambda_j = 1$ such that*

$$L^{(d+k)}(p) = \sum_j \lambda_j \operatorname{tr} p(\underline{A}^{(j)}) \tag{25}$$

for all $p \in \mathbb{R}\langle \underline{X} \rangle_{2d}$. In particular, $f_\star^{\mathrm{II}_1} = f_{\Theta^2}^{(d+k)}$.

Proof. We show this is a consequence of the Gelfand-Naimark-Segal (GNS) construction and the Artin-Wedderburn theory.

First of all, consider the truncated tracial moment sequence $(y_w)_{|w| \leq 2d+2k}$ defined by

$$y_w = L^{(d+k)}(w)$$

for $w \in \langle \underline{X} \rangle_{2d+2k}$. By assumption, the sequence is flat over $(y_w)_{|w| \leq 2d+2k-2\delta}$, hence in particular over $(y_w)_{|w| \leq 2d+2k-2}$. By [BK12, Theorem 3.18], $(y_w)_{|w| \leq 2d+2k}$ extends to a full tracial moment sequence $(y_w)_{|w| < \infty}$ which is flat over the original truncated sequence $(y_w)_{|w| \leq 2d+2k}$. This tracial sequence yields a tracial linear functional $L : \mathbb{R}\langle \underline{X} \rangle \to \mathbb{R}$ by $L(w) = y_w$, extending $L^{(d+k)}$. We note L is symmetric, i.e., $L(w^*) = L(w)$, by construction.

Now apply the usual GNS construction to L. More precisely, L induces the semidefinite sesquilinear form

$$(p, q) \mapsto L(q^* p) \tag{26}$$

on $\mathbb{R}\langle \underline{X} \rangle$. Letting $\mathcal{N} = \{x \in \mathbb{R}\langle \underline{X} \rangle \mid (x, x) = 0\}$ denote the nullvectors, $\mathcal{H} := \mathbb{R}\langle \underline{X} \rangle / \mathcal{N}$ is a finite-dimensional vector space (this is where flatness enters) spanned by the images $\bar{w}$ for $w \in \langle \underline{X} \rangle_d$. The sesquilinear form (26) induces an inner product $\langle \_, \_ \rangle$ on $\mathcal{H}$ by

$$\langle \bar{u}, \bar{w} \rangle = (u, w).$$

Let A_j denote the left multiplication by X_j on $\mathcal{H}$, i.e.,

$$A_j(\bar{w}) = \overline{X_j w}. \tag{27}$$

It is easy to see these A_j are (well-defined and) self-adjoint, that is, $A_j^* = A_j$. Now

$$L(p) = \langle p(\underline{A})\bar{1}, \bar{1} \rangle = \langle \bar{p}, \bar{1} \rangle. \tag{28}$$

We claim that $\underline{A} \in \mathcal{D}_S$. Let $g \in S$ and consider an arbitrary $v \in \mathcal{H}$. There is a $u \in \mathbb{R}\langle \underline{X} \rangle_d$ with $\bar{u} = u(\underline{A})\bar{1} = v$. Now

$$\begin{aligned} \langle g(\underline{A})v, v \rangle &= \langle g(\underline{A})u(\underline{A})\bar{1}, u(\underline{A})\bar{1} \rangle \\ &= \langle (gu)(\underline{A})\bar{1}, u(\underline{A})\bar{1} \rangle = L(u^* g u). \end{aligned}$$

Since $u^* g u \in M_{S,d+k}$,

$$L(u^* g u) = L^{(d+k)}(u^* g u) \geq 0$$

by assumption.

Now the A_j are matrices of size $\dim(\mathcal{H}) \leq \sigma(d)$.

Let $\mathcal{A}$ be the subalgebra generated by the symmetric matrices A_j. Since the Hermitian square of a nonzero matrix is not nilpotent, $\mathcal{A}$ is semisimple. By the Artin-Wedderburn theorem, $\mathcal{A}$ can be (orthogonally) block diagonalized into

$$\mathcal{A} = \oplus_{j=1}^{r} \mathcal{A}_j. \tag{29}$$

Here $\mathcal{A}_i$ are simple algebras (with involution), and thus $*$-isomorphic to full matrix algebras over $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$. With respect to the decomposition (29), $A_j = \oplus_{\ell=1}^{r} A_j^\ell$. Each A_j^ℓ is a self-adjoint matrix, and the tuple $\underline{A}^\ell \in \mathcal{D}_S$. Without loss of generality each A_j^ℓ is a real matrix; if one of the blocks $\mathcal{A}_j$ is a matrix algebra over $\mathbb{C}$ or $\mathbb{H}$, we embed it into the real matrix algebra (twice the size for $\mathbb{C}$ and four times the size for $\mathbb{H}$).

The tracial linear functional L induces tracial $\mathbb{R}$-linear functionals L_j on the simple $*$-algebras $\mathcal{A}_j$. If $p(\underline{A}) = \underline{B} = \oplus_{j=1}^{r} \underline{B}^j$, then

$$L(p) = \sum_j L_j(\underline{B}^j).$$

Each L_j is a positive multiple, say λ_j, of the usual trace [BK12, Lemma 3.11]. Thus

$$L(p) = \sum_j L_j(\underline{B}^j) = \sum_j \lambda_j \operatorname{tr} p(\underline{A}^j).$$

Since $L(1) = 1$, $\sum_j \lambda_j = 1$. ∎

We propose Algorithm 2 to find solutions of $(\text{DSDP}_{\text{tr}-\min})_s$ for $s \geq \deg(f) + \delta$ which are δ-flat enabling us to extract a minimizer of $(\text{SPSDP}_{\text{tr}-\min})_s$. It is a variant of Algorithm 1 and performs surprisingly well; actually it finds flat solutions in all tested situations where finite convergence was detected.

INPUT: $f \in \operatorname{Sym}\mathbb{R}\langle \underline{X}\rangle$ *with* $\deg f = 2d$, $S = \{g_1, \dots, g_r\}$, $\delta = \lceil \max_i \deg(g_i)/2\rceil$, $\delta_{\max}$.

FOR $s = d+\delta, d+\delta+1, ..., d+\delta+d_{\max}$,

STEP 1: *Compute* $L^{(s)}$ – *the optimal solution for* $(\text{DSDP}_{\text{tr}-\min})_s$.

STEP 2: IF $L^{(s)}$ *is* δ*-flat* THEN *STOP.*

STEP 3: ELSE *compute* $L^{(s)}_{\text{rand}}$. IF $L^{(s)}_{\text{rand}}$ *is* δ*-flat* THEN *STOP.*

OUTPUT: $L^{(s)}_{\text{rand}}$

Algorithm 2: Randomized algorithm to find flat solutions for $(\text{DSDP}_{\text{tr}-\min})_s$

In Step 3 we solve the SDP which is obtained from $(\text{DSDP}_{\text{tr}-\min})_s$ by fixing the upper left corner of the Hankel matrix to be equal to left upper corner of the Hankel matrix of $L^{(s)}$ and by taking a full random objective function — like in $(\text{DSDP}_{\text{rand}})_s$. We repeat this step several (e.g. 10) times. In our experiments, this algorithm often returns flat solutions if the module $M_{S,d}$ is archimedean. On the other hand, there is little theoretical evidence supporting this performance.

We repeat Steps 1–3 at most $\delta_{\max}+1$ times, where $\delta_{\max}$ is for computational complexity reasons chosen so that $d+\delta+\delta_{\max}$ is at most 10, if we have 2 nc variables and is at most 8 if we have 3 nc variables. Otherwise the complexity of the underlying SDP exceeds capability of our current hardware. We implemented Steps 1–3 from Algorithm 2 in the `NCSOStools` function `NCtraceOptRand`. Here is a simple demonstration.

Example 6.2.
```
>> NCvars X Y
>> w = 2 - X^2 + X*Y^2*X - Y^2;
>> [X,fX,trace_val]=NCtraceOptRand(w,{4-X^2-Y^2,X*Y+Y*X-2},4);
```

This gives a matrix X of size 2×16; each row represents one symmetric 4×4 matrix,

$$A = \text{reshape}(X(1,:),4,4) = \begin{bmatrix} -0.0000 & 1.4044 & -0.1666 & -0.0000 \\ 1.4044 & 0.0000 & 0.0000 & 1.1329 \\ -0.1666 & 0.0000 & -0.0000 & -0.8465 \\ -0.0000 & 1.1329 & -0.8465 & 0.0000 \end{bmatrix}$$

$$B = \text{reshape}(X(2,:),5,5) = \begin{bmatrix} -0.0000 & 0.8465 & 1.1329 & 0.0000 \\ 0.8465 & 0.0000 & 0.0000 & -0.1666 \\ 1.1329 & 0.0000 & 0.0000 & -1.4044 \\ 0.0000 & -0.1666 & -1.4044 & 0.0000 \end{bmatrix}$$

such that A and B are from $\mathcal{D}_S(4)$ and

$$fX = w(A,B) = \begin{bmatrix} -1.0000 & 0.0000 & 0.0000 & -0.0000 \\ 0.0000 & -1.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & -1.0000 & 0.0000 \\ -0.0000 & 0.0000 & 0.0000 & -1.0000 \end{bmatrix}$$

with (normalized) trace equal to `trace_val` $= -1$.

6.2. **Numerical results on random NC polynomials.** In this section we report numerical results obtained by running Algorithm 2 on random polynomials. We generated random polynomials as in Subsection 3.2.

Like in Subsection 3.2 we check for δ-flatness by computing ranks in three different ways: using Matlab functions `rank, rref` and by SVD decomposition. In all three cases we take the tolerance to be $\min\{30 \cdot \mathtt{err_{flat}}, 10^{-3}\}$.

With this tolerance we again observed that in all tested (random) cases Algorithm 2 returned a flat optimal solution already after the first step, i.e., for $s = d + \delta$. In the following table we report numerical results:

n	d	# of rand inst.	% of flat sol.	average $\mathtt{err_{flat}}$
2	2	100	100	0,000023
2	4	100	100	0,00033
2	6	100	98	$< 10^{-6}$
3	4	100	100	0,00026

TABLE 4. Numerical results obtained by running Algorithm 2 on random nc polynomials in n variables of degree $2d$. For every n and d we generated 100 instances and computed the percentage of 2-flat solutions obtained by Algorithm 2. For almost all randomly generated instances we found a 2-flat solution. The last column contains the average of $\mathtt{err_{flat}}$ over all 100 tested random instances.

7. Concluding Remarks

In this paper we have shown how to effectively compute the smallest (or biggest) eigenvalue or trace a noncommutative (nc) polynomial can attain on a free semialgebraic set. This in turn allows us to prove or produce new matrix inequalities in a dimension-free setting subject to polynomial constraints.

Our algorithm is based on sums of hermitian squares and commutators and implements the noncommutative Lasserre relaxation scheme as a sequence of semidefinite programs (SDPs). To prove exactness, we investigate the solutions of the dual SDPs. If one of these has a rank-preserving property called flatness, we use it to extract eigenvalue or trace optimizers with a procedure based on the solution to a truncated noncommutative moment problem, the Gelfand-Naimark-Segal (GNS) construction and the Artin-Wedderburn theory. To enforce flatness we employ a noncommutative variant of Nie's novel randomization technique.

We have implemented these procedures in our open source computer algebra system `NCSOStools`, freely available at `http://ncsostools.fis.unm.si/` and this is demonstrated throughout the paper with several illustrative examples.

References

[AL12] M. F. Anjos and J. B. Lasserre. *Handbook of Semidefinite, Conic and Polynomial Optimization: Theory, Algorithms, Software and Applications*, volume 166 of *International Series in Operational Research and Management Science*. Springer, 2012. 4

[Bar02] A. Barvinok. *A course in convexity*, volume 54 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2002. 10, 12

[BCKP13] S. Burgdorf, K. Cafuta, I. Klep, and J. Povh. The tracial moment problem and trace-optimization of polynomials. *Math. Program.*, 137(1-2):557–578, 2013. 14

[BK09] M. Brešar and I. Klep. Noncommutative polynomials, Lie skew-ideals and tracial Nullstellensätze. *Math. Res. Lett.*, 16(4):605–626, 2009. 10, 11

[BK12] S. Burgdorf and I. Klep. The truncated tracial moment problem. *J. Operator Theory*, 68:141–163, 2012. 13, 19, 20

[Brä11] P. Brändén. Obstructions to determinantal representability. *Adv. Math.*, 226(2):1202–1212, 2011. 3

[CF96] R. E. Curto and L. A. Fialkow. Solution of the truncated complex moment problem for flat data. *Mem. Amer. Math. Soc.*, 119(568):x+52, 1996. 3

[CF98] R. E. Curto and L. A. Fialkow. Flat extensions of positive moment matrices: recursively generated relations. *Mem. Amer. Math. Soc.*, 136(648):x+56, 1998. 3

[Cim10] J. Cimprič. A method for computing lowest eigenvalues of symmetric polynomial differential operators by semidefinite programming. *J. Math. Anal. Appl.*, 369(2):443–452, 2010. 1

[CKP11] K. Cafuta, I. Klep, and J. Povh. NCSOStools: a computer algebra system for symbolic and numerical computation with noncommutative polynomials. *Optim. Methods. Softw.*, 26(3):363–380, 2011. Available from http://ncsostools.fis.unm.si/. 1, 7

[CKP12] K. Cafuta, I. Klep, and J. Povh. Constrained polynomial optimization problems with noncommuting variables. *SIAM J. Optim.*, 22(2):363–383, 2012. 2, 3, 4, 5, 6, 10, 19

[Con76] A. Connes. Classification of injective factors. Cases II_1, II_∞, III_λ, $\lambda \neq 1$. *Ann. of Math. (2)*, 104(1):73–115, 1976. 9

[dK02] E. de Klerk. *Aspects of semidefinite programming*, volume 65 of *Applied Optimization*. Kluwer Academic Publishers, Dordrecht, 2002. 17

[HKM12] J. W. Helton, I. Klep, and S. McCullough. The convex Positivstellensatz in a free algebra. *Adv. Math.*, 231(1):516–534, 2012. 3, 4

[HLL09] D. Henrion, J. B. Lasserre, and J. Löfberg. GloptiPoly 3: moments, optimization and semidefinite programming. *Optim. Methods Softw.*, 24(4-5):761–779, 2009. Available from http://www.laas.fr/~henrion/software/gloptipoly3/. 1

[HM04] J. W. Helton and S. McCullough. A Positivstellensatz for non-commutative polynomials. *Trans. Amer. Math. Soc.*, 356(9):3721–3737, 2004. 3, 5

[HM12] J. W. Helton and S. McCullough. Every convex free basic semi-algebraic set has an LMI representation. *Ann. of Math. (2)*, 176(2):979–1013, 2012. 3

[HMdOP08] J. W. Helton, S. McCullough, M. C. de Oliveira, and M. Putinar. Engineering systems and free semi-algebraic geometry. In *Emerging Applications of Algebraic Geometry*, volume 149 of *IMA Vol. Math. Appl.*, pages 17–62. Springer, 2008. 1

[KS08] I. Klep and M. Schweighofer. Connes' embedding conjecture and sums of Hermitian squares. *Adv. Math.*, 217(4):1816–1837, 2008. 9, 10, 11, 15, 17

[KVV14] D. S. Kaliuzhnyi-Verbovetskyi and V. Vinnikov. *Foundations of free noncommutative function theory*, volume 199. American Mathematical Society, 2014. 11

[Las01] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2000/01. 2, 4

[Las09] J. B. Lasserre. *Moments, Positive Polynomials and Their Applications*, volume 1. Imperial College Press, 2009. 4

[Lau09] M. Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging applications of algebraic geometry*, pages 157–270. Springer, 2009. 4

[Löf04] J. Löfberg. YALMIP: A toolbox for modeling and optimization in MATLAB. In *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004. Available from http://control.ee.ethz.ch/~joloef/wiki/pmwiki.php. 1

[Mit03] D. Mittelmann. An independent benchmarking of SDP and SOCP solvers. *Math. Program. B*, 95:407–430, 2003. http://plato.asu.edu/bench.html. 4

[MPRW09] J. Malick, J. Povh, F. Rendl, and A. Wiegele. Regularization methods for semidefinite programming. *SIAM J. Optim.*, 20(1):336–356, 2009. 4

[Nie14] J. Nie. The $\mathcal{A}$-truncated $\mathcal{K}$-moment problem. *Found. Comput. Math.*, 14(6):1243–1276, 2014. 2, 4, 5

[NT14] T. Netzer and A. Thom. Hyperbolic polynomials and generalized Clifford algebras. *Discrete Comput. Geom.*, 51:802–814, 2014. 3

[PNA10] S. Pironio, M. Navascués, and A. Acín. Convergent relaxations of polynomial optimization problems with noncommuting variables. *SIAM J. Optim.*, 20(5):2157–2180, 2010. 1, 2, 4, 5, 19

[PPSP05] S. Prajna, A. Papachristodoulou, P. Seiler, and P.A. Parrilo. SOSTOOLS and its control applications. In *Positive polynomials in control*, volume 312 of *Lecture Notes in Control and Inform. Sci.*, pages 273–292. Springer, Berlin, 2005. 1

[Pro76] C. Procesi. The invariant theory of $n \times n$ matrices. *Adv. Math.*, 19(3):306–381, 1976. 11

[PS76] C. Procesi and M. Schacher. A non-commutative real nullstellensatz and hilbert's 17th problem. *Ann. Math.*, 104(3):395–406, 1976. 11

[PS01] V. Powers and C. Scheiderer. The moment problem for non-compact semialgebraic sets. *Adv. Geom.*, 1(1):71–88, 2001. 12

[Put93] M. Putinar. Positive polynomials on compact semi-algebraic sets. *Indiana Univ. Math. J.*, 42(3):969–984, 1993. 3

[Row80] L. H. Rowen. *Polynomial identities in ring theory*, volume 84 of *Pure and Applied Mathematics*. Academic Press Inc., New York, 1980. 10, 11

[Stu99] J. F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optim. Methods Softw.*, 11/12(1-4):625–653, 1999. Available from http://sedumi.ie.lehigh.edu/. 17

[Tak03] M. Takesaki. *Theory of operator algebras. III.* Berlin: Springer, 2003. 9

[TTT99] K. C. Toh, M. J. Todd, and R. H. Tütüncü. SDPT3—a MATLAB software package for semidefinite programming, version 1.3. *Optim. Methods Softw.*, 11/12(1-4):545–581, 1999. Available from http://www.math.nus.edu.sg/~mattohkc/sdpt3.html. 17

[VB96] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Rev.*, 38(1):49–95, 1996. 17

[WKK+09] H. Waki, S. Kim, M. Kojima, M. Muramatsu, and H. Sugimoto. Algorithm 883: sparsePOP—a sparse semidefinite programming relaxation of polynomial optimization problems. *ACM Trans. Math. Software*, 35(2):Art. 15, 13, 2009. 1

[WSV00] H. Wolkowicz, R. Saigal, and L. Vandenberghe. *Handbook of Semidefinite Programming*. Kluwer, 2000. 4

[XL08] S. Xu and J. Lam. A survey of linear matrix inequality techniques in stability analysis of delay systems. *Internat. J. Systems Sci.*, 39(12):1095–1113, 2008. 4

[YFK03] M. Yamashita, K. Fujisawa, and M. Kojima. Implementation and evaluation of SDPA 6.0 (semidefinite programming algorithm 6.0). *Optim. Methods Softw.*, 18(4):491–505, 2003. Available from http://sdpa.sourceforge.net/. 17

Igor Klep, Department of Mathematics, The University of Auckland, New Zealand

E-mail address: igor.klep@auckland.ac.nz

Janez Povh, Fakulteta za informacijske študije v Novem mestu, Slovenia

E-mail address: janez.povh@fis.unm.si

NOT FOR PUBLICATION

Contents