

Exploiting Sparsity in the Coefficient Matching Conditions in Sum-of-Squares Programming using ADMM

Yang Zheng, Giovanni Fantuzzi and Antonis Papachristodoulou, *Senior Member, IEEE*

Abstract—This paper introduces an efficient first-order method based on the alternating direction method of multipliers (ADMM) to solve semidefinite programs (SDPs) arising from sum-of-squares (SOS) programming. We exploit the sparsity of the *coefficient matching conditions* when SOS programs are formulated in the usual monomial basis to reduce the computational cost of the ADMM algorithm. Each iteration of our algorithm requires one projection onto the positive semidefinite cone and the solution of multiple quadratic programs with closed-form solutions free of any matrix inversion. Our techniques are implemented in the open-source MATLAB solver SOSADMM. Numerical experiments on SOS problems arising from unconstrained polynomial minimization and from Lyapunov stability analysis for polynomial systems show speed-ups compared to the interior-point solver SeDuMi, and the first-order solver CDCS.

Index Terms—Sum-of-squares, ADMM, large-scale problems.

I. INTRODUCTION

CHECKING whether a given polynomial is nonnegative has applications in many areas (see [1], [2] and the references therein). For example, the unconstrained polynomial optimization problem $\min_{x \in \mathbb{R}^n} p(x)$ is equivalent to

$$\begin{aligned} \max \quad & \gamma \\ \text{subject to} \quad & p(x) - \gamma \geq 0. \end{aligned} \quad (1)$$

Moreover, the stability of an equilibrium x^* of a polynomial dynamical system $\dot{x}(t) = f(x(t))$, $x(t) \in \mathbb{R}^n$ in a neighbourhood \mathcal{D} of x^* (we assume $x^* = 0$ without loss of generality)—a fundamental problem in control theory—can be established by constructing a polynomial $V(x)$ (called *Lyapunov function*) that satisfies the polynomial inequalities [3]

$$\begin{cases} V(x) > 0, & \forall x \in \mathcal{D} \setminus \{0\}, \\ -\dot{V}(x) = -\langle \nabla V(x), f(x) \rangle \geq 0, & \forall x \in \mathcal{D}. \end{cases} \quad (2)$$

Throughout this work, $\langle \cdot, \cdot \rangle$ denotes the inner product in the appropriate Hilbert space.

A powerful way to test polynomial inequalities is to employ a sum-of-squares (SOS) relaxation (we refer the reader to [4], [5] for details on SOS relaxations in polynomial optimization,

and to [3] for a tutorial on SOS techniques for systems analysis). In fact, while testing the non-negativity of a polynomial is NP-hard in general, the existence of an SOS decomposition can be checked in polynomial time by solving a semidefinite program (SDP) [4]. Unfortunately, however, the size of the SDP for the SOS relaxation of a degree- d polynomial in n variables is $\binom{n+d}{d}$. Consequently, SOS relaxations are limited to small problem instances; with the current technology, for example, Lyapunov-based analysis is impractical for general systems with ten or more states.

In order to mitigate scalability issues, one can act at the *modeling* level, *i.e.* one can try to replace the SDP obtained from an SOS relaxation with an optimization problem that is cheaper to solve still using second-order interior-point methods (IPMs), implemented in efficient solvers such as SeDuMi [6]. One approach is to exploit structural properties of the polynomial whose positivity is being tested [7]–[11]. For example, computing the Newton polytope [7] or checking for diagonal inconsistency [8] can restrict the monomial basis required in the SOS decomposition by eliminating redundant monomials. Further improvements are possible by group-theoretic symmetry reduction techniques [9] and graph-theoretic correlative sparsity [10]. Facial reduction has also been applied to select a reduced monomial basis for SOS programs in [11]. A second approach is to approximate the positive semidefinite (PSD) cone using diagonally dominant or scaled diagonally dominant matrices [12], [13]. These relaxations can be solved with linear programs (LPs) or second-order-cone programs (SOCPs), rather than SDPs, and the conservativeness introduced by approximating the PSD cone can be reduced with a recently proposed basis pursuit algorithm [14].

Further improvements are available on the *computational* level if IPMs are replaced by more scalable first-order methods (FOMs) at the cost of reduced accuracy. The design of efficient first-order algorithms for large-scale SDPs has received particular attention in recent years. For instance, Wen *et al.* proposed an alternating direction augmented Lagrangian method for large-scale dual SDPs [15]. O’Donoghue *et al.* developed an operator-splitting method to solve the homogeneous self-dual embedding of conic programs [16], which has recently been extended by the authors to exploit aggregate sparsity via chordal decomposition [17], [18]. In the context of SOS programming, Bertsimas *et al.* proposed an accelerated FOM for unconstrained polynomial optimization [19], while Henrion & Mallick introduced a projection-based method for SOS relaxations [20]. However, both approaches are only applicable

Y. Zheng is supported by the Clarendon Scholarship and the Jason Hu Scholarship. A. Papachristodoulou is supported by project EP/M002454/1.

Y. Zheng and A. Papachristodoulou are with the Department of Engineering Science, University of Oxford, Parks Road, Oxford, OX1 3PJ, U.K. (e-mail: {yang.zheng, antonis}@eng.ox.ac.uk)

G. Fantuzzi is with the Department of Aeronautics, Imperial College London, South Kensington Campus, London, SW7 2AZ, United Kingdom (e-mail: gf910@ic.ac.uk)

to a small subset of SOS programs as they rely on the constraint of the corresponding SDP being orthogonal, which is not the case for SOS problems with free variables.

In this paper, we propose a first-order algorithm based on the alternating direction method of multipliers (ADMM) to solve the SDPs arising from SOS optimization. In contrast to [19] and [20], we exploit the sparsity in the coefficient matching conditions, making our approach suitable for a larger class of SOS programs. While the aggregate sparsity pattern of these SDPs is dense (so that the methods of [17], [18] are not very advantageous), each equality constraint in the SDP only involves a small subset of decision variables when an SOS program is formulated in the usual monomial basis. This sparsity can be exploited to formulate an efficient ADMM algorithm, the iterations of which consist of conic projections and optimization problems with closed-form solutions that—crucially—are free of any matrix inversion. We implement our techniques in SOSADMM, an open-source MATLAB solver. The efficiency of our methods compared to the IPM solver SeDuMi [6] and the first-order solver CDCS [21] is demonstrated on SOS problems arising from unconstrained polynomial optimization and from Lyapunov stability analysis of polynomial systems.

The rest of this paper is organized as follows. Section II reviews SOS polynomials and the ADMM algorithm. Sparsity for SDPs arising in SOS programs is discussed in Section III, and we show how to exploit it to build an efficient ADMM algorithm in Section IV. Numerical experiments are reported in Section V. Section VI concludes the paper.

II. PRELIMINARIES

A. SOS polynomials and SDPs

The sets of real and natural numbers (including zero) are denoted by \mathbb{R} and \mathbb{N} , respectively. Let $x \in \mathbb{R}^n$, $\alpha \in \mathbb{N}^n$, and let $x^\alpha = x_1^{\alpha_1} x_2^{\alpha_2} \cdots x_n^{\alpha_n}$ denote a monomial in x of degree $|\alpha| = \sum_{i=1}^n \alpha_i$. Given an integer $d \in \mathbb{N}$, we denote $\mathbb{N}_d^n = \{\alpha \in \mathbb{N}^n : |\alpha| \leq d\}$, and the vector of all monomials of degree no greater than d by

$$v_d(x) = \{x^\alpha \mid \alpha \in \mathbb{N}_d^n\} \\ = [1, x_1, x_2, \dots, x_n, x_1^2, x_1 x_2, \dots, x_n^d]^T. \quad (3)$$

The length of $v_d(x)$ is $|\mathbb{N}_d^n| = \binom{n+d}{d}$. A real polynomial $p(x)$ is a finite, real linear combination of monomials of x

$$p(x) = \sum_{\alpha \in \mathbb{N}^n} p_\alpha x^\alpha, \quad p_\alpha \in \mathbb{R}.$$

The degree of $p(x)$ is the maximum of the degrees of all monomials with nonzero coefficients. We denote the set of real polynomials in x by $\mathbb{R}[x]$.

Definition 1: A polynomial $p(x) \in \mathbb{R}[x]$ of degree $2d$ is a sum-of-squares (SOS) if there exist polynomials $f_i(x) \in \mathbb{R}[x]$, $i = 1, \dots, m$ of degree no greater than d such that

$$p(x) = \sum_{i=1}^m [f_i(x)]^2.$$

Clearly, the existence of an SOS representation guarantees that $p(x) \geq 0$. The following theorem gives an equivalent characterization of SOS polynomials.

Proposition 1 ([4]): A polynomial $p(x) \in \mathbb{R}[x]$ of degree $2d$ is an SOS polynomial if and only if there exists a $\binom{n+d}{d} \times \binom{n+d}{d}$ symmetric PSD matrix $X \succeq 0$ such that

$$p(x) = v_d(x)^T X v_d(x). \quad (4)$$

The equality in (4) gives a set of affine equalities on the elements of X to match the coefficients of $p(x)$. Together with $X \succeq 0$, this makes the problem of finding an SOS representation for $p(x)$ an SDP. The formulation of such SDPs can be assisted by software packages, such as SOSTOOLS [22] and GloptiPoly [23].

Remark 1: The size of the PSD matrix X in (4) is $\binom{n+d}{d} \times \binom{n+d}{d}$ because we have used the full set of monomials of degree no greater than d in our representation. This number might be reduced by inspecting the structural properties of $p(x)$ to identify and eliminate redundant monomials in $v_d(x)$; well-known techniques include Newton polytope [7], diagonal inconsistency [8], symmetry property [9], and facial reduction [11].

B. ADMM algorithm

The ADMM algorithm solves the optimization problem

$$\begin{aligned} \min_{y,z} \quad & f(y) + g(z) \\ \text{subject to} \quad & Ay + Bz = c, \end{aligned} \quad (5)$$

where $y \in \mathbb{R}^n$ and $z \in \mathbb{R}^m$ are the decision variables, $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $g : \mathbb{R}^m \rightarrow \mathbb{R}$ are convex functions, and $A \in \mathbb{R}^{l \times n}$, $B \in \mathbb{R}^{l \times m}$ and $c \in \mathbb{R}^l$ are the constraint data. Given a penalty parameter $\rho > 0$ and a multiplier $\lambda \in \mathbb{R}^l$ (known as the *dual variable*), the ADMM algorithm solves (5) by finding a saddle point [24, Chapter 5.4] of the augmented Lagrangian

$$L_\rho(y, z, \lambda) = f(y) + g(z) + \lambda^T (Ay + Bz - c) \\ + \frac{\rho}{2} \|Ay + Bz - c\|^2 \quad (6)$$

with the following steps:

$$y^{k+1} = \arg \min_y L_\rho(y, z^k, \lambda^k), \quad (7a)$$

$$z^{k+1} = \arg \min_z L_\rho(y^{k+1}, z, \lambda^k), \quad (7b)$$

$$\lambda^{k+1} = \lambda^k + \rho(Ay^{k+1} + Bz^{k+1} - c). \quad (7c)$$

In these equations, the superscript k denotes the value of a variable at the k -th iteration of the algorithm, and $\|\cdot\|$ denotes the standard Euclidean norm, i.e., $\|x\| = \sqrt{x^T x}$ for $x \in \mathbb{R}^n$. Then, from a computational perspective, steps (7a) and (7b) are equivalent to the minimizations of

$$\tilde{L}_\rho(y, z, \lambda) = f(y) + g(z) + \frac{\rho}{2} \left\| Ay + Bz - c + \frac{1}{\rho} \lambda \right\|^2 \quad (8)$$

over y and z , respectively, with λ fixed. Under very mild conditions, ADMM converges to a solution with a rate $\mathcal{O}(\frac{1}{k})$, which is independent of ρ , although its value can affect convergence in practice; see [25, Section 3.2] for details.

TABLE I
DENSITY OF NONZERO ELEMENTS IN THE EQUALITY CONSTRAINTS OF SDP (12)

n	4	6	8	10	12	14	16
$2d = 4$	1.42×10^{-2}	4.76×10^{-3}	2.02×10^{-3}	9.99×10^{-4}	5.49×10^{-4}	3.27×10^{-4}	2.06×10^{-4}
$2d = 6$	4.76×10^{-3}	1.08×10^{-3}	3.33×10^{-4}	1.25×10^{-4}	5.39×10^{-5}	2.58×10^{-5}	1.34×10^{-5}
$2d = 8$	2.02×10^{-3}	3.33×10^{-4}	7.77×10^{-5}	2.29×10^{-5}	7.94×10^{-6}	3.13×10^{-6}	1.36×10^{-6}

III. ROW SPARSITY IN SDPS FROM SOS PROGRAMS

A. SDP formulations of SOS relaxations

Let A_α be the indicator matrix for the monomials x^α in the rank-one matrix $v_d(x)v_d(x)^T$; in other words, the entry of A_α with row index β and column index γ (where the natural ordering for multi-indices $\beta, \gamma \in \mathbb{N}_d^n$ is used) satisfies

$$(A_\alpha)_{\beta, \gamma} = \begin{cases} 1 & \text{if } \beta + \gamma = \alpha \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

The SOS constraint (4) can then be reformulated as

$$p(x) = \langle v_d(x)v_d(x)^T, X \rangle = \sum_{\alpha \in \mathbb{N}_{2d}^n} \langle A_\alpha, X \rangle x^\alpha. \quad (10)$$

Matching the coefficients of the left- and right-hand sides gives the equality constraints

$$\langle A_\alpha, X \rangle = p_\alpha \quad \forall \alpha \in \mathbb{N}_{2d}^n. \quad (11)$$

We refer to these equalities as *coefficient matching conditions*. The existence of an SOS decomposition for $p(x)$ (or lack thereof) can then be checked with the feasibility SDP

$$\begin{aligned} & \text{find } X \\ & \text{subject to } \langle A_\alpha, X \rangle = p_\alpha, \quad \alpha \in \mathbb{N}_{2d}^n, \\ & \quad X \succeq 0. \end{aligned} \quad (12)$$

When the full monomial basis is used, as in this case, the dimension of X and the number of constraints in (12) are, respectively,

$$N = |\mathbb{N}_d^n| = \binom{n+d}{d}, \quad m = |\mathbb{N}_{2d}^n| = \binom{n+2d}{2d}. \quad (13)$$

B. Properties of the coefficient matching conditions

In this section, for simplicity, we re-index the constraint matching conditions (11) using integer indices $i = 1, \dots, m$ instead of the multi-indices α .

The conditions (11) inherit two important properties from the data matrices A_i , $i = 1, \dots, m$. The first one follows from the fact that the matrices A_i are orthogonal. If n_i denotes the number of nonzero entries in A_i we have

$$\langle A_i, A_j \rangle = \begin{cases} n_i & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases} \quad (14)$$

After letting $\text{vec} : \mathbb{S}^N \rightarrow \mathbb{R}^{N^2}$ be the usual operator mapping a matrix to the stack of its columns, and defining

$$A = [\text{vec}(A_1) \quad \dots \quad \text{vec}(A_m)]^T, \quad (15)$$

the equality constraints in (12) can be rewritten as the matrix-vector product $A \cdot \text{vec}(X) = b$, where $b \in \mathbb{R}^m$ is a vector collecting the coefficients p_i , $i = 1, \dots, m$. Property (14)




Fig. 1. Sparsity pattern of AA^T for the example (16)

directly implies the following lemma, which forms the basis of the FOMs of [19], [20].

Lemma 1 (Orthogonality of constraints): AA^T is an $m \times m$ diagonal matrix with $(AA^T)_{ii} = n_i$.

The second property of the coefficient matching conditions is that they are sparse, in the sense that each equality constraint in (12) only involves a small subset of entries of X , because only a small subset of entries of the product $v_d(x)v_d(x)^T$ are equal to a given monomial x^α . Thus, the vectorized matrix A is *row sparse*, meaning that each row is a sparse vector. In particular, the following result holds.

Lemma 2 (Sparsity of constraints): Let A be the vectorized matrix for (12), and let N and m be as in (13). The number of nonzero elements in A is N^2 , and the density of nonzero elements in A is equal to $m^{-1} = \mathcal{O}(n^{-2d})$.

Proof: Since the matrix $v_d(x)v_d(x)^T$ contains all monomials x^α , $\alpha \in \mathbb{N}_{2d}^n$, all entries of the PSD matrix X enter at least one of the equality constraints in (12). Moreover, (14) implies that each entry of X enters at most one constraint. Therefore, A must contain N^2 nonzero elements. Its density is then given by

$$\frac{N^2}{N^2 \times m} = \frac{1}{m} = \left[\binom{n+2d}{2d} \right]^{-1} = \mathcal{O}(n^{-2d}).$$

Remark 2: While the constraint matrix A is sparse (see typical values in Table I), the aggregate sparsity pattern of the SDP (12) is dense because all entries of the matrix variable X appear in the equality constraints. This implies that X is generally a dense variable, so the first-order algorithms of [17], [18] are not particularly suitable.

Remark 3: The property of orthogonality in Lemma 1 holds for standard SOS feasibility problems. However, this property fails for the following example:

$$\begin{aligned} & \text{find } a, b \\ & \text{subject to } ax^4 + bx^2 + x + 1 \text{ is SOS,} \\ & \quad bx^4 + ax^2 + x + 1 \text{ is SOS.} \end{aligned} \quad (16)$$

Fig. 1 shows the sparsity pattern of AA^T for (16) obtained using SOSTOOLS, demonstrating that the constraints are not orthogonal. The reason is that (16) involves the free parameters a, b as well as the PSD matrices for the SOS representation,

and this destroys the orthogonality. This issue is common in control applications; see, *e.g.*, the condition (2) when finding Lyapunov functions. Consequently, the first-order algorithms in [19], [20] cannot be applied to many problems with SOS constraints because they rely on constraint orthogonality.

IV. EXPLOITING ROW SPARSITY IN SDPs

As we have seen, the algorithms in [17]–[20] are not useful for generic SOS problems because their aggregate sparsity pattern is dense, and the orthogonality property only holds for simple SOS feasibility problems. However, the data matrix A is always row-sparse due to the coefficient matching conditions in SOS programs. This property can be exploited to construct an efficient ADMM algorithm. In the following, we consider a generic SDP in the vectorized form

$$\begin{aligned} \min_x \quad & c^T x \\ \text{subject to} \quad & Ax = b, \quad x \in \mathcal{K}, \end{aligned} \quad (17)$$

where x is the optimization variable, $A \in \mathbb{R}^{m \times \hat{n}}$, $b \in \mathbb{R}^m$ and $c \in \mathbb{R}^{\hat{n}}$ are the problem data, and \mathcal{K} is a product of cones, at least one of which is the PSD cone. Throughout this section, $\delta_{\mathcal{S}}(x)$ denotes the indicator function of a set \mathcal{S} ,

$$\delta_{\mathcal{S}}(x) = \begin{cases} 0, & \text{if } x \in \mathcal{S}, \\ +\infty, & \text{if } x \notin \mathcal{S}. \end{cases}$$

A. Reformulation considering individual row sparsity

Let us represent $A = [a_1, a_2, \dots, a_m]^T$, so each vector a_i is a row of A , and let $H_i \in \mathbb{R}^{|a_i| \times \hat{n}}$, $i = 1, \dots, m$ be “entry-selector” matrices of 1’s and 0’s selecting the nonzero elements of a_i , where $|a_i|$ denotes the number of nonzero elements of a_i . Note that the rows of H_i are orthonormal, since each selects a different entry of a_i . Then,

$$Ax = b \Leftrightarrow \begin{cases} (H_i a_i)^T z_i = b_i, & i = 1, \dots, m, \\ z_i = H_i x, & i = 1, \dots, m. \end{cases} \quad (18)$$

In (18), z_i is a copy of the elements of x which enter the i -th affine constraint. It is also convenient to introduce an additional slack variable $u = x$, so the affine constraints in (18) and conic constraint in (17) are decoupled when applying the ADMM algorithm. We can then reformulate (17) as

$$\begin{aligned} \min_{z_i, u, x} \quad & c^T x \\ \text{subject to} \quad & (H_i a_i)^T z_i = b_i \quad i = 1, \dots, m, \\ & z_i = H_i x, \quad i = 1, \dots, m, \\ & u = x, u \in \mathcal{K}. \end{aligned} \quad (19)$$

B. ADMM steps

To apply ADMM, we move the affine constraints $(H_i a_i)^T z_i = b_i$ and the conic constraint $u \in \mathcal{K}$ in (19) to the objective using the indicator functions $\delta_0(\cdot)$ and $\delta_{\mathcal{K}}(\cdot)$, respectively:

$$\begin{aligned} \min_{z_i, u, x} \quad & c^T x + \delta_{\mathcal{K}}(u) + \sum_{i=1}^m \delta_0((H_i a_i)^T z_i - b_i) \\ \text{subject to} \quad & z_i = H_i x, \quad i = 1, \dots, m, \\ & u = x. \end{aligned} \quad (20)$$

The augmented Lagrangian of (20) is

$$\begin{aligned} L = c^T x + \delta_{\mathcal{K}}(u) + \sum_{i=1}^m \delta_0[(H_i a_i)^T z_i - b_i] \\ + \sum_{i=1}^m \mu_i^T (z_i - H_i x) + \frac{\rho}{2} \|z_i - H_i x\|^2 \\ + \xi^T (u - x) + \frac{\rho}{2} \|u - x\|^2, \end{aligned} \quad (21)$$

and we group the variables as

$$\mathcal{Y} = \{x\}, \quad \mathcal{Z} = \{u, z_1, \dots, z_m\}, \quad \mathcal{D} = \{\mu_1, \dots, \mu_m, \xi\}.$$

Then, the ADMM steps (7a)–(7c) become the following:

1) *Minimization over \mathcal{Y}* : The minimization of (21) over the variables in \mathcal{Y} is equivalent to an unconstrained quadratic program,

$$\min_x c^T x + \frac{\rho}{2} \sum_{i=1}^m \left\| z_i^k - H_i x + \frac{\mu_i^k}{\rho} \right\|^2 + \frac{\rho}{2} \left\| u^k - x + \frac{\xi^k}{\rho} \right\|^2. \quad (22)$$

The updated variable x^{k+1} is then simply given by

$$x^{k+1} = D^{-1} \left[\sum_{i=1}^m H_i^T \left(z_i^k + \frac{\mu_i^k}{\rho} \right) + \left(u^k + \frac{\xi^k}{\rho} \right) - \frac{1}{\rho} c \right], \quad (23)$$

where the matrix $D = I + \sum_{i=1}^m H_i^T H_i$ is diagonal because the rows of each matrix H_i are orthonormal. This means that (23) is cheap to calculate.

2) *Minimization over \mathcal{Z}* : Minimizing (21) over the variables in \mathcal{Z} amounts to a conic projection,

$$\begin{aligned} \min_u \quad & \|u - x^{k+1} + \rho^{-1} \xi^k\|^2 \\ \text{subject to} \quad & u \in \mathcal{K}, \end{aligned} \quad (24)$$

plus m independent quadratic programs

$$\begin{aligned} \min_{z_i} \quad & \|z_i - H_i x^{k+1} + \rho^{-1} \mu_i^k\|^2 \\ \text{subject to} \quad & (H_i a_i)^T z_i = b_i. \end{aligned} \quad (25)$$

The projection (24) is easy to compute when \mathcal{K} is a product of \mathbb{R}^n , the non-negative orthant, second-order cones, and PSD cones; for example, a projection onto the PSD cone only requires one eigen-decomposition. As for problem (25), its KKT conditions are

$$z_i - H_i x^{k+1} + \rho^{-1} \mu_i^k + (H_i a_i) \omega_i = 0, \quad (26a)$$

$$(H_i a_i)^T z_i = b_i, \quad (26b)$$

where ω_i is the Lagrangian multiplier for the equality constraint in (25). Simple algebra shows that

$$\omega_i = \frac{1}{\|H_i a_i\|^2} \left(-b_i + (H_i a_i)^T H_i x^{k+1} - \frac{1}{\rho} (H_i a_i)^T \mu_i^k \right),$$

so the solution z_i^{k+1} to (25) can be calculated easily with (26a). Note that this step is free of any matrix inversion.

3) *Update multipliers \mathcal{D}* : According to (7c), the multipliers in \mathcal{D} are updated with inexpensive and parallelizable gradient ascent steps:

$$\begin{aligned} \mu_i^{k+1} &= \mu_i^k + \rho(z_i^{k+1} - H_i x^{k+1}), \quad i = 1, \dots, m, \\ \xi^{k+1} &= \xi^k + \rho(u^{k+1} - x^{k+1}). \end{aligned} \quad (27)$$

C. Summary of the computations in the ADMM algorithm

In the proposed ADMM algorithm, subproblems (7a) and (7b) have explicit closed-form solutions. Each iteration requires solving

- 1) one unconstrained quadratic program, given by (22);
- 2) one conic projection, given by (24);
- 3) m independent quadratic programs, given by (25).

Note that only the nonzero elements of a_i appear in (25). Since we have assumed that a_i is sparse, only operations on vectors of small size are required. Besides, our algorithm is free of matrix inversion (with the exception of the $m \times m$ diagonal matrix D , requiring $O(m)$ flops), which results from introducing the local variables z_i so each affine constraint can be considered individually. The cost is that our algorithm needs to maintain multiple local variables z_i , which may have adverse effects on the convergence speed of ADMM.

In contrast, the FOMs in [19], [20] fail to deal with general SOS programs since they rely on orthogonality of constraints, and those in [17], [18] require factorizing the $m \times m$ matrix AA^T (in general, $O(m^3)$ flops). In SDPs arising from generic SOS relaxations, this step is computationally demanding because AA^T is not diagonal (as seen in Remark 3, this is often the case) and the number m is usually large ($m = 18564$ if $n = 12$ and $2d = 6$ in (12)). Note that all these algorithms converge at rate $O(\frac{1}{k})$ because they are based on ADMM or its variants.

V. NUMERICAL EXPERIMENTS

We implemented our techniques in SOSADMM, an open-source first-order MATLAB solver for conic programs with row sparsity. Currently, SOSADMM supports cartesian products of the following cones: \mathbb{R}^n , non-negative orthant, second-order cone, and the positive definite cone. SOSADMM, the numerical examples presented in this section, and additional examples are available from

<https://github.com/oxfordcontrol/SOSADMM>

We tested SOSADMM on random unconstrained polynomial optimization problems and Lyapunov stability analysis of polynomial systems. To assess the suboptimality of the solution returned by SOSADMM, we compared it to the accurate one computed with the interior-point solver SeDuMi [6]. CPU times were compared to the first-order solver CDCS [21], which exploits aggregate sparsity in SDPs; in particular, the primal method in CDCS was used [17]. In our experiments, the termination tolerance for SOSADMM and CDCS was set to 10^{-4} , and the maximum number of iterations was 2000. To improve convergence, SOSADMM employs an adaptive penalty parameter update rule [25], with an initial value $\rho = 1$. All tests were run on a PC with a 2.8 GHz Intel® Core™ i7 CPU and 8GB of RAM.

A. Unconstrained polynomial optimization

Consider the global polynomial minimization problem

$$\min_{x \in \mathbb{R}^n} p(x), \quad (28)$$

where $p(x)$ is a given polynomial. This problem is equivalent to (1), and we can obtain an SDP relaxation by replacing the

TABLE II
CPU TIME (S) TO SOLVE THE SDP RELAXATIONS OF (28). N IS THE SIZE OF THE PSD CONE, m IS THE NUMBER OF CONSTRAINTS.

Dimensions			CPU time (s)		
n	N	m	SeDuMi	CDCS (primal)	SOS-ADMM
2	6	14	0.23	0.08	0.05
4	15	69	0.13	0.11	0.06
6	28	209	0.24	0.16	0.14
8	45	494	1.16	0.18	0.18
10	66	1000	3.17	0.25	0.39
12	91	1819	13.89	0.46	0.55
14	120	3059	54.63	0.79	0.84
16	153	4844	187.0	0.92	0.82
18	190	7314	610.2	2.91	1.92
20	231	10625	1739	4.93	2.32

TABLE III
LYAPUNOV FUNCTIONS FOR THE SYSTEM (29)

Solver	Time (s)	Lyapunov function $V(x)$
SeDuMi	0.054	$6.659x_1^2 + 4.628x_2^2 + 2.073x_3^2$
CDCS-primal	0.21	$7.008x_1^2 + 1.477x_2^2 + 2.172x_3^2$
SOSADMM	0.58	$6.699x_1^2 + 1.803x_2^2 + 2.172x_3^2$

non-negativity constraint with an SOS condition on $p(x) - \gamma$. Motivated by [20], we generated $p(x)$ according to

$$p(x) = p_0(x) + \sum_{i=1}^n x_i^{2d},$$

where $p_0(x)$ is a random polynomial of degree strictly less than $2d$. We used GloptiPoly [23] to generate the examples.

Table II compares the CPU time (in seconds) required to solve the SOS relaxation as the number n of variables was increased with $d = 2$. Both SOSADMM and CDCS-primal were faster than SeDuMi on these examples (note that SeDuMi's runtime reduces if a weaker termination tolerance is set, but not significantly). Also, the optimal value returned by SOSADMM was within 0.05% of the high-accuracy value returned by SeDuMi. For all examples in Table II, the cone size N is moderate (less than 300), while the number of constraints m is large. SeDuMi assembles and solves an $m \times m$ linear system at each iteration, which is computationally expensive.

B. Finding Lyapunov functions

Next, we consider the problem of constructing Lyapunov functions to check local stability of polynomial/rational systems when (2) is replaced by SOS conditions. We used SOSTOOLS [22] to generate the corresponding SDPs.

The first system we study is

$$\begin{aligned} \dot{x}_1 &= -x_1^3 - x_1x_3^2, \\ \dot{x}_2 &= -x_2 - x_1^2x_2, \\ \dot{x}_3 &= -x_3 - \frac{3x_3}{x_3^2 + 1} + 3x_1^2x_3, \end{aligned} \quad (29)$$

which is demo 2 in SOSTOOLS. The system has an equilibrium at the origin, and we search for a homogeneous quadratic polynomial Lyapunov function $V(x) = ax_1^2 + bx_2^2 + cx_3^2$ to prove its global stability. The results given by SeDuMi, CDCS-primal and SOSADMM are listed in Table III. For such a small

TABLE IV
CPU TIME (S) TO CONSTRUCT A QUADRATIC LYAPUNOV FUNCTION FOR
RANDOMLY GENERATED POLYNOMIAL SYSTEMS.

n	Statistics		CPU time (s)		
	Size of A	nonzero density	SeDuMi	CDCS (primal)	SOS-ADMM
10	1100×2365	1.50×10^{-3}	3.3	7.5	5.3
12	1963×4407	8.76×10^{-4}	11.0	11.8	7.7
14	3255×7560	5.25×10^{-4}	49.9	21.0	11.2
16	5100×12172	3.13×10^{-4}	181.9	31.7	16.2
18	7638×18639	2.13×10^{-4}	574.8	55.0	24.6
20	11025×27405	1.48×10^{-4}	1617.2	100.3	37.7
22	15433×38962	1.11×10^{-4}	7442.7	265.9	65.6
25	24375×62725	6.87×10^{-5}	*	729.1	104.7
30	47275×124620	3.64×10^{-5}	*	3509.2	259.0

* SeDuMi fails due to memory requirements.

system, SeDuMi was slightly faster than CDCS-primal and SOSADMM, which is expected since IPMs are well-suited for small-scale SDPs. Note that since the problem of constructing a Lyapunov functions is a feasibility problem, the solutions returned by SeDuMi, CDCS-primal and SOSADMM need not be the same (see Table III).

As the last example, we consider randomly generated polynomial dynamical systems $\dot{x} = f(x)$ of degree three with a locally asymptotically stable equilibrium at the origin, and checked for local nonlinear stability in the ball $\mathcal{D} = \{x \in \mathbb{R}^n | 0.1 - \|x\|^2 \geq 0\}$ using a complete quadratic polynomial as the candidate Lyapunov function. Table IV summarizes the average CPU times required to search for such a Lyapunov function, when successful (note that we cannot detect infeasible problems because we only solve the primal form (17)). The results clearly show that SOSADMM is faster than both SeDuMi and CDCS-primal for the largest problem instances ($n \geq 18$). Also, FOMs have much lower memory requirements, and SOSADMM can solve problems that are not accessible with IPM: SeDuMi failed due to memory issues when $n > 22$. Finally, note that for the problem of finding Lyapunov functions the $m \times m$ linear system solved in SeDuMi and CDCS is not diagonal, and solving it is expensive: when $n = 30$ it took over 150 s for CDCS just to factorize AA^T , which is over 50% of the total time taken by SOSADMM to return a solution.

VI. CONCLUSION

In this paper, we proposed an efficient ADMM algorithm to exploit the row-sparsity of SDPs that arise from SOS programming, which are implemented in SOSADMM. The subproblems of our algorithm consist of one conic projection and multiple quadratic programs with closed-form solutions, which can be computed efficiently and—most importantly—do not require any matrix inversion.

Our numerical experiments on random unconstrained polynomial optimization and on Lyapunov stability analysis of polynomial/rational systems demonstrate that our method can provide speed-ups compared to the interior-point solver SeDuMi and the first-order solver CDCS. One major drawback of our method is the inability to detect infeasibility; future work will try to exploit the sparsity of SDPs from SOS relaxations

in a homogeneous self-dual embedding formulation similar to that of [16], [18].

REFERENCES

- [1] J. B. Lasserre, *Moments, positive polynomials and their applications*. World Scientific, 2009, vol. 1.
- [2] G. Chesi and D. Henrion, “Guest editorial: Special issue on positive polynomials in control,” *IEEE Trans. Autom. Control*, vol. 54, no. 5, pp. 935–936, 2009.
- [3] A. Papachristodoulou and S. Prajna, “A tutorial on sum of squares techniques for systems analysis,” in *Am. Control Conf. (ACC)*. IEEE, 2005, pp. 2686–2700.
- [4] P. A. Parrilo, “Semidefinite programming relaxations for semialgebraic problems,” *Math. Program.*, vol. 96, no. 2, pp. 293–320, 2003.
- [5] J. B. Lasserre, “Global optimization with polynomials and the problem of moments,” *SIAM J. Optim.*, vol. 11, no. 3, pp. 796–817, 2001.
- [6] J. F. Sturm, “Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones,” *Optim. Methods Softw.*, vol. 11, no. 1-4, pp. 625–653, 1999.
- [7] B. Reznick *et al.*, “Extremal PSD forms with few terms,” *Duke Math. J.*, vol. 45, no. 2, pp. 363–374, 1978.
- [8] J. Löfberg, “Pre-and post-processing sum-of-squares programs in practice,” *IEEE Trans. Autom. Control*, vol. 54, no. 5, pp. 1007–1011, 2009.
- [9] K. Gatermann and P. A. Parrilo, “Symmetry groups, semidefinite programs, and sums of squares,” *J. Pure Appl. Algebra*, vol. 192, no. 1, pp. 95–128, 2004.
- [10] H. Waki, S. Kim, M. Kojima, and M. Muramatsu, “Sums of squares and semidefinite program relaxations for polynomial optimization problems with structured sparsity,” *SIAM J. Control Optim.*, vol. 17, no. 1, pp. 218–242, 2006.
- [11] F. Permenter and P. A. Parrilo, “Basis selection for SOS programs via facial reduction and polyhedral approximations,” in *53rd IEEE Conf. Decis. Control (CDC)*, 2014, pp. 6615–6620.
- [12] A. Majumdar, A. A. Ahmadi, and R. Tedrake, “Control and verification of high-dimensional systems with DSOS and SDSOS programming,” in *IEEE 53rd Conf. Decis. Control (CDC)*, 2014, pp. 394–401.
- [13] A. A. Ahmadi and A. Majumdar, “DSOS and SDSOS optimization: LP and SOCP-based alternatives to sum of squares optimization,” in *48th Annu. Conf. Inf. Sci. Syst. (CISS)*. IEEE, 2014, pp. 1–5.
- [14] A. A. Ahmadi and G. Hall, “Sum of squares basis pursuit with linear and second order cone programming,” *arXiv preprint arXiv:1510.01597*, 2015.
- [15] Z. Wen, D. Goldfarb, and W. Yin, “Alternating direction augmented Lagrangian methods for semidefinite programming,” *Math. Program. Comput.*, vol. 2, no. 3-4, pp. 203–230, 2010.
- [16] B. O’Donoghue, E. Chu, and S. Parrilo, Nealand Boyd, “Conic optimization via operator splitting and homogeneous self-dual embedding,” *J. Optim. Theory Appl.*, vol. 169, no. 3, pp. 1042–1068, 2016.
- [17] Y. Zheng, G. Fantuzzi, A. Papachristodoulou, P. Goulart, and A. Wynn, “Fast ADMM for semidefinite programs with chordal sparsity,” *arXiv preprint arXiv:1609.06068*, 2016.
- [18] —, “Fast ADMM for homogeneous self-dual embedding of sparse SDPs,” *arXiv preprint arXiv:1611.01828*, 2016.
- [19] D. Bertsimas, R. M. Freund, and X. A. Sun, “An accelerated first-order method for solving SOS relaxations of unconstrained polynomial optimization problems,” *Optim. Methods Softw.*, vol. 28, no. 3, pp. 424–441, 2013.
- [20] D. Henrion and J. Malick, “Projection methods in conic optimization,” in *Handbook on Semidefinite, Conic and Polynomial Optimization*. Springer, 2012, pp. 565–600.
- [21] Y. Zheng, G. Fantuzzi, A. Papachristodoulou, P. Goulart, and A. Wynn, “CDCS: Cone decomposition conic solver, version 1.1,” <https://github.com/oxfordcontrol/CDCS>, Sep. 2016.
- [22] A. Papachristodoulou, J. Anderson, G. Valmorbida, S. Prajna, P. Seiler, and P. Parrilo, “SOSTOOLS version 3.00 sum of squares optimization toolbox for MATLAB,” *arXiv preprint arXiv:1310.4716*, 2013.
- [23] D. Henrion and J.-B. Lasserre, “GloptiPoly: Global optimization over polynomials with MATLAB and SeDuMi,” *ACM Tans. Math. Softw. (TOMS)*, vol. 29, no. 2, pp. 165–194, 2003.
- [24] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
- [25] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Found. Trends® Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.