

Fast ADMM for Semidefinite Programs with Chordal Sparsity

Yang Zheng^{†,1}, Giovanni Fantuzzi^{†,2}, Antonis Papachristodoulou¹, Paul Goulart¹, Andrew Wynn²

Abstract—Many problems in control theory can be formulated as semidefinite programs (SDPs). For large-scale SDPs, it is important to exploit the inherent sparsity to improve scalability. This paper develops efficient first-order methods to solve SDPs with chordal sparsity based on the alternating direction method of multipliers (ADMM). We show that chordal decomposition can be applied to either the primal or the dual standard form of a sparse SDP, resulting in scaled versions of ADMM algorithms with the same computational cost. Each iteration of our algorithms consists of a projection on the product of small positive semidefinite cones, followed by a projection on an affine set, both of which can be carried out efficiently. Our techniques are implemented in CDCS, an open source add-on to MATLAB. Numerical experiments on large-scale sparse problems in SDPLIB and random SDPs with block-arrow sparse patterns show speedups compared to some common state-of-the-art software packages.

I. INTRODUCTION

Semidefinite programs (SDPs) are convex optimization problems over the cone of positive semidefinite (PSD) matrices. Given $b \in \mathbb{R}^m$, $C \in \mathbb{S}^n$, and matrices $A_1, \dots, A_m \in \mathbb{S}^n$ that define the operators

$$\mathcal{A}(X) = \begin{bmatrix} \langle A_1, X \rangle \\ \vdots \\ \langle A_m, X \rangle \end{bmatrix}, \quad \mathcal{A}^*(y) = \sum_{i=1}^m A_i y_i,$$

SDPs are typically written in the *standard primal form*

$$\begin{aligned} \min_X \quad & \langle C, X \rangle \\ \text{subject to} \quad & \mathcal{A}(X) = b, \\ & X \in \mathbb{S}_+^n, \end{aligned} \quad (1)$$

or in the *standard dual form*

$$\begin{aligned} \max_{y, Z} \quad & \langle b, y \rangle \\ \text{subject to} \quad & \mathcal{A}^*(y) + Z = C, \\ & Z \in \mathbb{S}_+^n. \end{aligned} \quad (2)$$

In the above and throughout this work, \mathbb{R}^m is the m -dimensional Euclidean space, \mathbb{S}^n is the space of $n \times n$ symmetric matrices, \mathbb{S}_+^n is the subspace of PSD matrices, and $\langle \cdot, \cdot \rangle$ denotes the inner product in the appropriate space.

[†]Y. Zheng and G. Fantuzzi contributed equally to this work. Y. Zheng is supported by the Clarendon Scholarship and the Jason Hu Scholarship. G. Fantuzzi was partially supported by the EPSRC grant EP/J010537/1.

¹Department of Engineering Science, University of Oxford, Parks Road, Oxford, OX1 3PJ, United Kingdom (e-mail: yang.zheng@eng.ox.ac.uk; paul.goulart@eng.ox.ac.uk; antonis@eng.ox.ac.uk).

²Department of Aeronautics, Imperial College London, South Kensington Campus, London, SW7 2AZ, United Kingdom (e-mail: gf910@ic.ac.uk; a.wynn@imperial.ac.uk).

SDPs have applications in control theory, machine learning, combinatorics, and operations research [1]. Moreover, linear, quadratic, and second-order-cone programs, are particular instances of SDPs [2]. Small to medium-sized SDPs can be solved in polynomial time [3] using efficient second-order interior-point methods (IPMs) [4], [5]. However, many real-life problems are too large for the state-of-the-art interior-point algorithms, due to memory and CPU time constraints.

One approach is to abandon IPMs in favour of faster first-order methods (FOMs) with modest accuracy. For instance, Wen *et al.* proposed an alternating-direction augmented Lagrangian method for large-scale SDPs in the dual standard form [6]. More recently, O'Donoghue *et al.* developed a first-order operator-splitting method to solve the homogeneous self-dual embedding (HSDE) of a primal-dual pair of conic programs, which has the advantage of being able to provide primal or dual certificates of infeasibility [7]. A second approach relies on the fact that large-scale SDPs are often structured and/or sparse [1]. Exploiting sparsity in SDPs is an active and challenging area of research [8], one main difficulty being that the optimal solution is typically dense despite the sparsity of the data. If, however, the sparsity pattern of the data is *chordal* or has sparse *chordal extensions*, Grone's and Agler's theorems [9], [10] allow us to replace the PSD constraint with a set of smaller semidefinite constraints, plus an additional set of equality constraints. In some cases, the modified SDP can then be solved more efficiently than the original problem. These ideas underly the *domain-* and *range-space* conversion techniques [11], [12], implemented in SparseCoLO [13].

However, adding equality constraints often offsets the benefit of working with smaller PSD cones. One possible solution is to exploit chordal sparsity directly in the IPMs: Fukuda *et al.* used Grone's theorem [9] to develop a primal-dual path-following method for SDPs [11]; Burer proposed a nonsymmetric primal-dual IPM using Cholesky factors of the dual variable and maximum determinant completion of the primal variable [14]; and Andersen *et al.* developed fast recursive algorithms for SDPs with chordal sparsity [15]. Alternatively, one can solve the decomposed SDP with FOMs: Sun *et al.* proposed a first-order splitting method for decomposable conic programs [16]; Kalbat & Lavaei applied the alternating-direction method of multipliers (ADMM) to SDPs with fully decomposable constraints [17]; Madani *et al.* developed a highly-parallelizable ADMM algorithm for sparse SDPs with inequality constraints with optimal power flow applications [18].

In this work we adopt the strategy of exploiting sparsity using first-order algorithms in the spirit of [16]–[18], and

develop efficient ADMM algorithms to solve large-scale sparse SDPs. Our contributions are:

- 1) We combine ADMM and chordal decomposition to solve sparse SDPs in primal or dual standard form. The resulting algorithms are scaled versions of each other. This gives a conversion framework for the application of FOMs, analogous to that of [11], [12] for IPMs.
- 2) In each iteration, the PSD constraint is enforced via parallel projections onto small PSD cones. The affine constraints are imposed by a quadratic program with equality constraints, and its KKT system matrix can be factorized before iterating the ADMM algorithm since it only depends on the problem data.
- 3) We implement our methods in the open-source MATLAB solver CDCS (Cone Decomposition Conic Solver) [19]. Numerical simulations on random SDPs with block-arrow sparsity patterns and on four large-scale sparse problems in SDPLIB [20] demonstrate the efficiency of our algorithms compared to other solvers.

The rest of this paper is organized as follows. Section II reviews chordal sparsity and decomposition techniques. We show how to apply the ADMM to primal and dual standard-form SDPs in Sections III–IV, respectively, and report our numerical experiments in Section V. Finally, Section VI concludes the paper.

II. PRELIMINARIES: CHORDAL DECOMPOSITION AND THE ADMM ALGORITHM

A. Chordal graphs

Let $\mathcal{G}(\mathcal{V}, \mathcal{E})$ be an undirected graph with vertices $\mathcal{V} = \{1, 2, \dots, n\}$ and edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. A *clique* $\mathcal{C} \subseteq \mathcal{V}$ is a subset of vertices such that $(i, j) \in \mathcal{E}$ for any distinct vertices $i, j \in \mathcal{C}$, and the number of vertices in \mathcal{C} is denoted by $|\mathcal{C}|$. If \mathcal{C} is not a subset of any other clique, then it is referred to as a *maximal clique*. A cycle of length k in \mathcal{G} is a set of pairwise distinct vertices $\{v_1, v_2, \dots, v_k\} \subset \mathcal{V}$ such that $(v_k, v_1) \in \mathcal{E}$ and $(v_i, v_{i+1}) \in \mathcal{E}$ for $i = 1, \dots, k-1$. A chord is an edge joining two non-adjacent vertices in a cycle.

Definition 1 (Chordal graph): An undirected graph is *chordal* if all cycles of length four or higher have a chord.

If $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is not chordal, it can be *chordal extended*, i.e., we can construct a chordal graph $\mathcal{G}'(\mathcal{V}, \mathcal{E}')$ by adding edges to \mathcal{E} such that \mathcal{G}' is chordal. Finding the chordal extension with the minimum number of additional edges is an NP-complete problem [21], but good chordal extensions can be computed efficiently using several heuristics [22].

B. Sparse matrices defined by graphs

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be an undirected graph such that $(i, i) \in \mathcal{E}$, i.e., each node has a self-loop. We say that X is a sparse symmetric matrix defined by \mathcal{G} if $X_{ij} = X_{ji} = 0$ whenever $(i, j) \notin \mathcal{E}$. The spaces of sparse and PSD sparse symmetric matrices defined by \mathcal{G} are

$$\begin{aligned}\mathbb{S}^n(\mathcal{E}, 0) &= \{X \in \mathbb{S}^n \mid X_{ij} = X_{ji} = 0 \text{ if } (i, j) \notin \mathcal{E}\}, \\ \mathbb{S}_+^n(\mathcal{E}, 0) &= \{X \in \mathbb{S}^n(\mathcal{E}, 0) \mid X \succeq 0\}.\end{aligned}$$

Similarly, we say that X is a partial symmetric matrix defined by \mathcal{G} if $X_{ij} = X_{ji}$ are given when $(i, j) \in \mathcal{E}$, and arbitrary otherwise. Moreover, we say that M is a PSD completion of the partial symmetric matrix X if $M \succeq 0$ and $M_{ij} = X_{ij}$ when $(i, j) \in \mathcal{E}$. We can then define the spaces

$$\begin{aligned}\mathbb{S}^n(\mathcal{E}, ?) &= \{X \in \mathbb{S}^n \mid X_{ij} = X_{ji} \text{ given if } (i, j) \in \mathcal{E}\}, \\ \mathbb{S}_+^n(\mathcal{E}, ?) &= \{X \in \mathbb{S}^n(\mathcal{E}, ?) \mid \exists M \succeq 0, M_{ij} = X_{ij} \forall (i, j) \in \mathcal{E}\}.\end{aligned}$$

Finally, given a clique \mathcal{C}_k of \mathcal{G} , $E_k \in \mathbb{R}^{|\mathcal{C}_k| \times n}$ is the matrix with $(E_k)_{ij} = 1$ if $\mathcal{C}_k(i) = j$ and zero otherwise, where $\mathcal{C}_k(i)$ is the i -th vertex in \mathcal{C}_k , sorted in the natural ordering. The submatrix of $X \in \mathbb{S}^n$ defined by \mathcal{C}_k is $E_k X E_k^T \in \mathbb{S}^{|\mathcal{C}_k|}$.

C. Chordal decomposition of PSD matrices

The spaces $\mathbb{S}_+^n(\mathcal{E}, ?)$ and $\mathbb{S}_+^n(\mathcal{E}, 0)$ are a pair of dual convex cones for any undirected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ [15], [22]. If \mathcal{G} is chordal, $\mathbb{S}_+^n(\mathcal{E}, ?)$ and $\mathbb{S}_+^n(\mathcal{E}, 0)$ can be expressed in terms of several coupled smaller convex cones:

Theorem 1 (Grone's theorem [9]): Let $\{\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_p\}$ be the set of maximal cliques of a chordal graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$. Then, $X \in \mathbb{S}_+^n(\mathcal{E}, ?)$ if and only if $X_k := E_k X E_k^T \in \mathbb{S}_+^{|\mathcal{C}_k|}$ for all $k = 1, \dots, p$.

Theorem 2 (Agler's theorem [10]): Let $\{\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_p\}$ be the set of maximal cliques of a chordal graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$. Then, $Z \in \mathbb{S}_+^n(\mathcal{E}, 0)$ if and only if there exist matrices $Z_k \in \mathbb{S}_+^{|\mathcal{C}_k|}$ for $k = 1, \dots, p$ such that $Z = \sum_{k=1}^p E_k^T Z_k E_k$.

These results can be proven individually, but can also be derived from each other using the duality of the cones $\mathbb{S}_+^n(\mathcal{E}, ?)$ and $\mathbb{S}_+^n(\mathcal{E}, 0)$ [22].

D. ADMM algorithm

The ADMM algorithm solves the optimization problem

$$\begin{aligned}\min \quad & f(x) + g(y) \\ \text{subject to} \quad & Ax + By = c,\end{aligned}$$

where f and g are convex functions, $x \in \mathbb{R}^{n_x}$, $y \in \mathbb{R}^{n_y}$, $A \in \mathbb{R}^{n_c \times n_x}$, $B \in \mathbb{R}^{n_c \times n_y}$ and $c \in \mathbb{R}^{n_c}$. Given a penalty parameter $\rho > 0$ and a dual multiplier $z \in \mathbb{R}^{n_c}$, the ADMM algorithm minimizes the augmented Lagrangian

$$L_\rho(x, y, z) = f(x) + g(y) + \frac{\rho}{2} \left\| Ax + By - c + \frac{1}{\rho} z \right\|^2$$

with respect to the variables x and y separately, followed by a dual variable update:

$$x^{(n+1)} = \arg \min_x L_\rho(x, y^{(n)}, z^{(n)}), \quad (3a)$$

$$y^{(n+1)} = \arg \min_y L_\rho(x^{(n+1)}, y, z^{(n)}), \quad (3b)$$

$$z^{(n+1)} = z^{(n)} + \rho(Ax^{(n+1)} + By^{(n+1)} - c). \quad (3c)$$

The superscript (n) indicates that a variable is fixed to its value at the n -th iteration. ADMM is particularly suitable when the minimizations with respect to each of the variables x and y in (3a) and (3b) can be carried out efficiently through closed-form expressions.

III. ADMM FOR SPARSE PRIMAL-FORM SDPs

A. Reformulation and decomposition of the PSD constraint

Let the primal-standard-form SDP (1) be sparse with an *aggregate sparsity pattern* described by the graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, meaning that $(i, j) \in \mathcal{E}$ if and only if the entry ij of at least one of the data matrices C, A_0, \dots, A_m , is nonzero. We assume that \mathcal{G} is chordal (otherwise it can be chordal extended) and that its maximal cliques $\mathcal{C}_1, \dots, \mathcal{C}_p$ are small. Then, only the entries of the matrix variable X corresponding to the graph edges \mathcal{E} appear in the cost and constraint functions, so the constraint $X \in \mathbb{S}_+^n$ can be replaced by $X \in \mathbb{S}_+^n(\mathcal{E}, ?)$. Using Theorem 1, we can then reformulate (1) as

$$\begin{aligned} \min_{X, X_1, \dots, X_p} \quad & \langle C, X \rangle \\ \text{subject to} \quad & \mathcal{A}(X) = b, \\ & X_k - E_k X E_k^T = 0, \quad k = 1, \dots, p, \\ & X_k \in \mathbb{S}_+^{|\mathcal{C}_k|}, \quad k = 1, \dots, p. \end{aligned} \quad (4)$$

In other words, we can decompose the original large semidefinite cone into multiple smaller cones, at the expense of introducing a set of consensus constraints between the variables.

To simplify the exposition, we rewrite (4) in a vectorized form. Letting $\text{vec} : \mathbb{S}^n \rightarrow \mathbb{R}^{n^2}$ be the usual operator mapping a matrix to the stack of its column, define the vectorized data $c := \text{vec}(C)$, $A := [\text{vec}(A_0) \dots \text{vec}(A_m)]^T$, the vectorized variables $x := \text{vec}(X)$, $x_k := \text{vec}(X_k)$, $k = 1, \dots, p$, and the matrices $H_k := E_k \otimes E_k$ such that $x_k = \text{vec}(X_k) = \text{vec}(E_k X E_k^T) = H_k x$. In other words, the matrices H_1, \dots, H_p are “entry-selector” matrices of 1’s and 0’s, whose rows are orthonormal, that project x onto the subvectors x_1, \dots, x_p , respectively. Denoting the constraints $X_k \in \mathbb{S}_+^{|\mathcal{C}_k|}$ by $x_k \in \mathcal{S}_k$, we can rewrite (4) as

$$\begin{aligned} \min_{x, x_1, \dots, x_p} \quad & \langle c, x \rangle \\ \text{subject to} \quad & Ax = b, \\ & x_k = H_k x, \quad k = 1, \dots, p, \\ & x_k \in \mathcal{S}_k, \quad k = 1, \dots, p. \end{aligned} \quad (5)$$

B. The ADMM algorithm for primal SDPs

Moving the constraints $Ax = b$ and $x_k \in \mathcal{S}_k$ in (5) to the objective using the indicator functions $\delta_0(\cdot)$ and $\delta_{\mathcal{S}_k}(\cdot)$ gives

$$\begin{aligned} \min_{x, x_1, \dots, x_p} \quad & \langle c, x \rangle + \delta_0(Ax - b) + \sum_{k=1}^p \delta_{\mathcal{S}_k}(x_k) \\ \text{subject to} \quad & x_k = H_k x, \quad k = 1, \dots, p. \end{aligned} \quad (6)$$

This problem is in the standard form for the application of ADMM. Given a penalty parameter $\rho > 0$ and a Lagrange multiplier λ_k for each constraint $x_k = H_k x$, we define

$$\begin{aligned} \mathcal{L} := & \langle c, x \rangle + \delta_0(Ax - b) \\ & + \sum_{k=1}^p \left[\delta_{\mathcal{S}_k}(x_k) + \frac{\rho}{2} \left\| x_k - H_k x + \frac{1}{\rho} \lambda_k \right\|^2 \right], \end{aligned} \quad (7)$$

and group the variables as $\mathcal{X} := \{x\}$, $\mathcal{Y} := \{x_1, \dots, x_p\}$, and $\mathcal{Z} := \{\lambda_1, \dots, \lambda_p\}$. As in (3), in each ADMM iteration, we minimize (7) with respect to \mathcal{X} and \mathcal{Y} , then update \mathcal{Z} .

Algorithm 1 ADMM for decomposed primal form SDPs

```

1: Input:  $\rho > 0$ ,  $\epsilon_{\text{tol}} > 0$ , initial guesses  $x^{(0)}$ ,
    $x_1^{(0)}, \dots, x_p^{(0)}$ ,  $\lambda_1^{(0)}, \dots, \lambda_p^{(0)}$ 
2: Setup: Chordal decomposition, KKT factorization.
3: while  $\max(\epsilon_p, \epsilon_d) \geq \epsilon_{\text{tol}}$  do
4:   Compute  $x^{(n)}$  using (9).
5:   for  $k = 1, \dots, p$ : Compute  $x_k^{(n)}$  using (11).
6:   for  $k = 1, \dots, p$ : Compute  $\lambda_k^{(n)}$  using (12).
7:   Update the residuals  $\epsilon_p, \epsilon_d$ .
8: end while

```

1) **Minimization over \mathcal{X} :** Minimizing (7) over \mathcal{X} is equivalent to the equality-constrained quadratic program

$$\min_x \quad \langle c, x \rangle + \frac{\rho}{2} \sum_{k=1}^p \left\| x_k^{(n)} - H_k x + \frac{1}{\rho} \lambda_k^{(n)} \right\|^2 \quad (8)$$

subject to $Ax = b$.

Define $D := \sum_{k=1}^p H_k^T H_k$ and let ρy be the multiplier for the equality constraint. We can write the optimality conditions for (8) as the KKT system

$$\begin{bmatrix} D & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \sum_{k=1}^p H_k^T \left(x_k^{(n)} + \rho^{-1} \lambda_k^{(n)} \right) - \rho^{-1} c \\ b \end{bmatrix}. \quad (9)$$

Note that D is a diagonal matrix, because the rows of each matrix H_k are orthonormal, so (9) can be solved efficiently, e.g., by block elimination. Moreover, the coefficient matrix is the same at every iteration, so its factorization can be pre-computed and cached before starting the ADMM iterations.

2) **Minimization over \mathcal{Y} :** Minimizing (7) over \mathcal{Y} is equivalent to the p independent problems

$$\min_{x_k} \left\| x_k - H_k x^{(n+1)} + \rho^{-1} \lambda_k^{(n)} \right\|^2 \quad \text{subject to } x_k \in \mathcal{S}_k. \quad (10)$$

In terms of the original matrix variables X_1, \dots, X_p , this amounts to a projection on the PSD cone. More precisely, if \mathbb{P}_k denotes the projection onto $\mathbb{S}_+^{|\mathcal{C}_k|}$ we have

$$x_k^{(n+1)} = \text{vec} \left\{ \mathbb{P}_k \left[\text{vec}^{-1} \left(H_k x^{(n+1)} - \rho^{-1} \lambda_k^{(n)} \right) \right] \right\}. \quad (11)$$

Since computing \mathbb{P}_k amounts to an eigenvalue decomposition and each cone $\mathbb{S}_+^{|\mathcal{C}_k|}$ is small by assumption, we can compute $x_1^{(n+1)}, \dots, x_p^{(n+1)}$ efficiently and in parallel.

3) **Updating the multipliers \mathcal{Z} :** Each multiplier λ_k , $k = 1, \dots, p$, is updated with the usual gradient ascent rule,

$$\lambda_k^{(n+1)} = \lambda_k^{(n)} + \rho \left(x_k^{(n+1)} - H_k x^{(n+1)} \right). \quad (12)$$

This computation is cheap, and can be parallelized.

The ADMM algorithm is stopped after the n -th iteration if the relative primal/dual error measures ϵ_p and ϵ_d are smaller than a specified tolerance, ϵ_{tol} ; see [23] for more details on stopping conditions for a generic ADMM algorithm. Algorithm 1 summarizes the steps to solve a decomposable SDP in standard primal form (5).

IV. ADMM FOR SPARSE DUAL-FORM SDPs

A. Reformulation of decomposition of the PSD constraint

Similar to Section III, suppose the aggregate sparsity pattern of an SDP in standard dual form (2) is described by the chordal graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$. The equality constraint in (2) implies that the PSD variable Z has the same sparsity pattern as the aggregate sparsity pattern of the problem data, *i.e.*, $Z \in \mathbb{S}_+^n(\mathcal{E}, 0)$, so using Theorem 2 we can rewrite (2) as

$$\begin{aligned} \min_{y, Z_1, \dots, Z_p} \quad & -\langle b, y \rangle \\ \text{subject to} \quad & \mathcal{A}^*(y) + \sum_{k=1}^p E_k^T Z_k E_k = C, \\ & Z_k \in \mathbb{S}_+^{|\mathcal{C}_k|}, \quad k = 1, \dots, p. \end{aligned} \quad (13)$$

While the original PSD constraint has been replaced by multiple smaller PSD constraints, it is not convenient to apply ADMM to this problem form because the PSD variables Z_1, \dots, Z_k in the equality constraint are weighted by the matrices E_k . Instead, we replace Z_1, \dots, Z_k in the equality constraint with slack variables V_1, \dots, V_p such that $Z_k = V_k$, $k = 1, \dots, p$. Defining $z_k := \text{vec}(Z_k)$ and $v_k := \text{vec}(V_k)$ for all $k = 1, \dots, p$, and using the same vectorized notation as in Section III we then reformulate (13) in the vectorized form

$$\begin{aligned} \min_{y, z_1, \dots, z_p, v_1, \dots, v_p} \quad & -\langle b, y \rangle \\ \text{subject to} \quad & A^T y + \sum_{k=1}^p H_k^T v_k = c, \\ & z_k - v_k = 0, \quad k = 1, \dots, p, \\ & z_k \in \mathcal{S}_k, \quad k = 1, \dots, p. \end{aligned} \quad (14)$$

Remark 1: Although we have derived (14) by applying Theorem 2, (14) is exactly the dual of the decomposed primal SDP (5). Consequently, our analysis provides a decomposition framework for the application of FOMs analogous to that of [11], [12] for IPMs. This elegant picture, in which the decomposed SDPs inherit the duality between the original ones by virtue of the duality between Grone's and Agler's theorems, is shown in Fig. 1.

B. The ADMM algorithm for dual SDPs

Using indicator functions to move all but the equality constraints $z_k = v_k$, $k = 1, \dots, p$, to the objective gives

$$\begin{aligned} \min \quad & -\langle b, y \rangle + \delta_0 \left(c - A^T y - \sum_{k=1}^p H_k^T v_k \right) + \sum_{k=1}^p \delta_{\mathcal{S}_k}(z_k) \\ \text{subject to} \quad & z_k = v_k, \quad k = 1, \dots, p. \end{aligned} \quad (15)$$

Given a penalty parameter $\rho > 0$ and a Lagrange multiplier λ_k for each constraint $z_k = v_k$, $k = 1, \dots, p$, we define

$$\begin{aligned} \mathcal{L} := & -\langle b, y \rangle + \delta_0 \left(c - A^T y - \sum_{k=1}^p H_k^T v_k \right) \\ & + \sum_{k=1}^p \left[\delta_{\mathcal{S}_k}(z_k) + \frac{\rho}{2} \left\| z_k - v_k + \frac{1}{\rho} \lambda_k \right\|^2 \right], \end{aligned} \quad (16)$$

and group the variables as $\mathcal{X} := \{y, v_1, \dots, v_p\}$, $\mathcal{Y} := \{z_1, \dots, z_p\}$, and $\mathcal{Z} := \{\lambda_1, \dots, \lambda_p\}$.




Fig. 1. Duality relationships between primal and dual SDPs, and the decomposed primal and dual SDPs.

1) **Minimization over \mathcal{X} :** Minimizing (16) over block \mathcal{X} is equivalent to the equality-constrained quadratic program

$$\begin{aligned} \min_{y, v_1, \dots, v_p} \quad & -\langle b, y \rangle + \frac{\rho}{2} \sum_{k=1}^p \left\| z_k^{(n)} - v_k + \frac{1}{\rho} \lambda_k^{(n)} \right\|^2 \\ \text{subject to} \quad & c - A^T y - \sum_{k=1}^p H_k^T v_k = 0. \end{aligned} \quad (17)$$

Let ρx be the multiplier for the equality constraint. After some algebra, the optimality conditions for (17) can be written as the KKT system

$$\begin{bmatrix} D & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} c - \sum_{k=1}^p H_k^T \left(z_k^{(n)} + \rho^{-1} \lambda_k^{(n)} \right) \\ -\rho^{-1} b \end{bmatrix}, \quad (18)$$

plus a set of p uncoupled equations for the variables v_k ,

$$v_k = z_k^{(n)} + \frac{1}{\rho} \lambda_k^{(n)} + H_k x, \quad k = 1, \dots, p. \quad (19)$$

The KKT system (18) is the same as (9) after rescaling $x \mapsto -x$, $y \mapsto -y$, $c \mapsto \rho^{-1}c$ and $b \mapsto \rho b$, and as in Section III-B.1 the factors of the coefficient matrix required to solve (18) can be pre-computed and cached. Consequently, updating \mathcal{X} has the same the cost as in Section III-B.1 plus the cost of (19), which is cheap and can also be parallelized.

2) **Minimization over \mathcal{Y} :** As in Section III-B.2, z_1, \dots, z_p are updated with p independent and efficient projections

$$z_k^{(n+1)} = \text{vec} \left\{ \mathbb{P}_k \left[\text{vec}^{-1} \left(v_k^{(n+1)} - \rho^{-1} \lambda_k^{(n)} \right) \right] \right\}. \quad (20)$$

3) **Updating the multipliers \mathcal{Z} :** The multipliers λ_k , $k = 1, \dots, p$, are updated with the usual gradient ascent rule

$$\lambda_k^{(n+1)} = \lambda_k^{(n)} + \rho \left(z_k^{(n+1)} - v_k^{(n+1)} \right). \quad (21)$$

As in Section III-B, we stop the ADMM algorithm when the relative primal/dual error measures ϵ_p and ϵ_d are smaller than a specified tolerance, ϵ_{tol} . Algorithm 2 summarizes the full ADMM algorithm for sparse dual-standard-form SDPs.

Remark 2: The computational cost of (19) is the same as (12), so the ADMM iterations for the decomposed dual-standard-form SDP (14) have the same cost as those for the decomposed primal-standard-form SDP (5), plus the cost of (21). However, if one minimizes (16) over \mathcal{Y} *before* minimizing it over \mathcal{X} , substituting (19) into (21) gives

$$\lambda_k^{(n+1)} = \rho H_k x^{(n+1)}, \quad k = 1, \dots, p. \quad (22)$$

Algorithm 2 ADMM for decomposed dual form SDPs

- 1: **Input:** $\rho > 0$, $\epsilon_{\text{tol}} > 0$, initial guesses $y^{(0)}$, $z_1^{(0)}, \dots, z_p^{(0)}$, $\lambda_1^{(0)}, \dots, \lambda_p^{(0)}$
- 2: **Setup:** Chordal decomposition, KKT factorization.
- 3: **while** $\max(\epsilon_p, \epsilon_d) \geq \epsilon_{\text{tol}}$ **do**
- 4: **for** $k = 1, \dots, p$: Compute $z_k^{(n)}$ using (20).
- 5: Compute $y^{(n)}, x$ using (17).
- 6: **for** $k = 1, \dots, p$: Compute $v_k^{(n)}$ using (19)
- 7: Compute $\lambda_k^{(n)}$ using (22) (no cost).
- 8: Update the residuals ϵ_p and ϵ_d .
- 9: **end while**

Since H_1x, \dots, H_px have already been computed to update v_1, \dots, v_p , updating $\lambda_1, \dots, \lambda_p$ requires only a scaling operation. Consequently, the ADMM algorithms for the primal- and dual-standard-form SDPs can be considered as scaled versions of each other, with the same leading-order computational cost at each iteration.

V. NUMERICAL SIMULATIONS

We have implemented our techniques in CDCS (Cone Decomposition Conic Solver) [19], an open-source MATLAB solver. CDCS supports cartesian products of the following cones: \mathbb{R}^n , non-negative orthant, second-order cone, and the PSD cone. Currently, only chordal decomposition techniques for semidefinite cones are implemented, while the other cone types are not decomposed. Our code can be downloaded from <https://github.com/OxfordControl/CDCS>.

We tested CDCS on four sparse large-scale ($n \geq 1000, m \geq 1000$) problems in SDPLIB [20], as well as on randomly generated SDPs with block-arrow sparse pattern, used as a benchmark in [16]. The performance is compared to that of the IPM solver SeDuMi [24] and of the first-order solver SCS [25] on both the full SDPs (without decomposition) and the SDPs decomposed by SparseCoLO [13].

The comparison has two purposes: 1) SeDuMi computes accurate optimal points, which can be used to assess the quality of the solution computed by CDCS; 2) SCS is a high-performance first-order solver for general conic programs, so we can assess the advantages of chordal decomposition. SeDuMi should not be compared to the other solvers on CPU time, because the latter only aim to achieve moderate accuracy. In the experiments reported below, the termination tolerance for CDCS and SCS was set to $\epsilon_{\text{tol}} = 10^{-3}$, with a maximum of 2000 iterations. All experiments were carried out on a PC with an Intel(R) Core(TM) i7 CPU, 2.8 GHz processor and 8GB of RAM.

A. SDPs with block-arrow pattern

SDPs with the block-arrow sparsity pattern shown in Fig. 2—which is chordal—are used as a benchmark in [16]. The SDP parameters are: the number of blocks, l ; the block size, d ; the size of the arrow head, h ; the number of constraints, m . Here, we consider the following cases: 1) Fix $l = 40$, $d = 10$, $h = 20$, vary m ; 2) Fix $m = 1000$, $d = 10$, $h = 20$, vary l ; 3) Fix $l = 40$, $h = 10$, $m = 1000$, vary d .




Fig. 2. Block-arrow sparsity pattern: the number of blocks, l ; block size, d ; the size of the arrow head, h .

TABLE I
PROBLEM STATISTICS FOR SDPLIB PROBLEMS

	maxG32	maxG51	thetaG51	qpG51
Affine constraints, m	2000	1000	6910	1000
Original cone size, n	2000	1000	1001	2000
Number of cliques, p	1499	674	674	1675
Maximum clique size	60	322	323	304
Minimum clique size	5	6	7	1

The CPU times for different solvers, averaged over five random problem instances, are shown in Fig. 3. CDCS is approximately 10 times faster than SeDuMi and the combination SparseCoLO+SeDuMi, our Algorithm 2 being the fastest. Besides, the optimal value from CDCS was always within 0.02% of the accurate value from SeDuMi.

B. Sparse SDPs from SDPLIB

We consider two max-cut problems (maxG32 and maxG51), a Lovász theta problem (thetaG51), and a box-constrained quadratic problem (qpG51) from SDPLIB [20]. Table I reports the dimensions and chordal decomposition details of these large, sparse SDPs. Table II summarizes our numerical results; maxG51, thetaG51 and qpG51 could not be solved by the combination SparseCoLO+SeDuMi due to memory overflow. For all four problems, CDCS (primal and dual) is faster than SeDuMi and can give speedups compared to SCS and SparseCoLO+SCS. We remark that the stopping objective value from CDCS is within 2% of the optimal value returned by SeDuMi (which is highly accurate, and can be considered exact) in all four cases, and within 0.08% for the max-cut problems maxG32 and maxG51 — a negligible difference in applications.

VI. CONCLUSION

We proposed a conversion framework for SDPs characterized by chordal sparsity suitable for the application of FOMs, analogous to the conversion techniques for IPMs of [11], [12]. We also developed efficient ADMM algorithms for sparse SDPs in primal or dual standard form, which are implemented in the solver CDCS. Numerical experiments on SDPs with block-arrow sparsity patterns and on large sparse problems in SDPLIB show that our methods can provide speedups compared to both IPM solvers such as SeDuMi [24]—even when the chordal sparsity is exploited using SparseCoLO [13]—and the state-of-the-art first-order solver SCS [25]. Exploiting chordal sparsity in a first-order HSDE formulation similar to that of [7] would be desirable in the future to be able to detect infeasibility.




Fig. 3. CPU time for SDPs with block-arrow patterns. Left to right: varying number of constraints; varying number of blocks; varying block size.

TABLE II
RESULTS FOR THE PROBLEM INSTANCES IN SDPLIB

		SeDuMi	SparseCoLo+ SeDuMi	SCS	SparseCoLo+ SCS	CDCS (primal)	CDCS (dual)
maxG32	Total time (s)	974.6	355.2	2.553×10^3	65.1	88.6	53.1
	Pre-processing time (s)	0	3.18	0.43	3.24	21.2	21.4
	Objective value	1.568×10^3	1.568×10^3	1.568×10^3	1.566×10^3	1.569×10^3	1.568×10^3
	Iterations	14	15	2000	960	238	127
maxG51	Total time (s)	134.5	—	87.9	1.201×10^3	110.9	75.9
	Pre-processing time (s)	0	—	0.11	2.87	3.30	3.20
	Objective value	4.006×10^3	—	4.006×10^3	3.977×10^3	4.005×10^3	4.006×10^3
	Iterations	16	—	540	2000	235	157
thetaG51	Total time (s)	2.218×10^3	—	424.2	1.346×10^3	471.2	735.1
	Pre-processing time (s)	0	—	0.30	5.30	25.1	25.0
	Objective value	349	—	350.6	341.3	354.5	355.9
	Iterations	20	—	2000	2000	394	646
qpG51	Total time (s)	1.407×10^3	—	2.330×10^3	985.8	727.1	606.2
	Pre-processing time (s)	0	—	0.47	190.2	12.3	12.3
	Objective value	1.182×10^3	—	1.288×10^3	1.174×10^3	1.195×10^3	1.194×10^3
	Iterations	22	—	2000	2000	1287	1048

REFERENCES

- [1] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*. SIAM, 1994.
- [2] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
- [3] L. Vandenberghe and S. Boyd, “Semidefinite programming,” *SIAM review*, vol. 38, no. 1, pp. 49–95, 1996.
- [4] F. Alizadeh, J.-P. A. Haeberly, and M. L. Overton, “Primal-dual interior-point methods for semidefinite programming: convergence rates, stability and numerical results,” *SIAM J. Optim.*, vol. 8, no. 3, pp. 746–768, 1998.
- [5] C. Helmberg, F. Rendl, R. J. Vanderbei, and H. Wolkowicz, “An interior-point method for semidefinite programming,” *SIAM J. Optim.*, vol. 6, no. 2, pp. 342–361, 1996.
- [6] 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.
- [7] B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd, “Conic optimization via operator splitting and homogeneous self-dual embedding,” *J. Optim. Theory Appl.*, vol. 169, no. 3, pp. 1042–1068, 2016.
- [8] M. Andersen, J. Dahl, Z. Liu, and L. Vandenberghe, “Interior-point methods for large-scale cone programming,” in *Optimization for machine learning*. MIT Press, 2011, pp. 55–83.
- [9] R. Grone, C. R. Johnson, E. M. Sá, and H. Wolkowicz, “Positive definite completions of partial hermitian matrices,” *Linear Algebra Appl.*, vol. 58, pp. 109–124, 1984.
- [10] J. Agler, W. Helton, S. McCullough, and L. Rodman, “Positive semidefinite matrices with a given sparsity pattern,” *Linear algebra Appl.*, vol. 107, pp. 101–149, 1988.
- [11] M. Fukuda, M. Kojima, K. Murota, and K. Nakata, “Exploiting sparsity in semidefinite programming via matrix completion I: General framework,” *SIAM J. Optim.*, vol. 11, no. 3, pp. 647–674, 2001.
- [12] S. Kim, M. Kojima, M. Mevissen, and M. Yamashita, “Exploiting sparsity in linear and nonlinear matrix inequalities via positive semidefinite matrix completion,” *Math. Program.*, vol. 129, no. 1, pp. 33–68, 2011.
- [13] K. Fujisawa, S. Kim, M. Kojima, Y. Okamoto, and M. Yamashita, “Users manual for SparseCoLo: Conversion methods for sparse conic form linear optimization problems,” Research Report B-453, Tokyo Institute of Technology, Tokyo 152-8552, Japan, Tech. Rep., 2009.
- [14] S. Burer, “Semidefinite programming in the space of partial positive semidefinite matrices,” *SIAM J. Optim.*, vol. 14, no. 1, pp. 139–172, 2003.
- [15] M. S. Andersen, J. Dahl, and L. Vandenberghe, “Implementation of nonsymmetric interior-point methods for linear optimization over sparse matrix cones,” *Math. Program. Comput.*, vol. 2, no. 3-4, pp. 167–201, 2010.
- [16] Y. Sun, M. S. Andersen, and L. Vandenberghe, “Decomposition in conic optimization with partially separable structure,” *SIAM J. Optim.*, vol. 24, no. 2, pp. 873–897, 2014.
- [17] A. Kalbat and J. Lavaei, “A fast distributed algorithm for decomposable semidefinite programs,” in *Proc. 54th IEEE Conf. Decis. Control*, Dec 2015, pp. 1742–1749.
- [18] R. Madani, A. Kalbat, and J. Lavaei, “ADMM for sparse semidefinite programming with applications to optimal power flow problem,” in *Proc. 54th IEEE Conf. Decis. Control*, Dec 2015, pp. 5932–5939.
- [19] Y. Zheng, G. Fantuzzi, A. Papachristodoulou, P. Goulart, and A. Wynn, “CDCS: Cone decomposition conic solver, version 1.0,” <https://github.com/OxfordControl/CDCS>, Sep. 2016.
- [20] B. Borchers, “SDPLIB 1.2, a library of semidefinite programming test problems,” *Optim. Methods Softw.*, vol. 11, no. 1-4, pp. 683–690, 1999.
- [21] M. Yannakakis, “Computing the minimum fill-in is NP-complete,” *SIAM Journal on Algebraic Discrete Methods*, vol. 2, pp. 77–79, 1981.
- [22] L. Vandenberghe and M. S. Andersen, “Chordal graphs and semidefinite optimization,” *Found. Trends® Optim.*, vol. 1, no. 4, pp. 241–433, 2014.
- [23] 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.
- [24] 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.
- [25] B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd, “SCS: Splitting conic solver, version 1.2.6,” <https://github.com/cvxgrp/scs>, Apr. 2016.