Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Decomposed Structured Subsets for Semidefinite Optimization ⋆

Jared Miller * **Yang Zheng** ** **Mario Sznaier** *
Antonis Papachristodoulou ***

* *ECE Department, Northeastern University, Boston, MA 02115. (Emails: miller.jare@husky.neu.edu, msznaier@coe.neu.edu).*
** *SEAS and CGBC at Harvard University, Cambridge, MA, 02138. (Email: zhengy@g.havard.edu).*
*** *Department of Engineering Science, University of Oxford, Oxford, UK OX1 3PJ. (Email: antonis@eng.ox.ac.uk).*

Abstract:
Semidefinite programs (SDPs) are important computational tools in controls, optimization, and operations research. Standard interior-point methods scale poorly for solving large-scale SDPs. With certain compromise of solution quality, one method for scalability is to use the notion of structured subsets (e.g. diagonally-dominant (DD) and scaled-diagonally dominant (SDD) matrices), to derive inner/outer approximations for SDPs. For sparse SDPs, chordal decomposition techniques have been widely used to derive equivalent SDP reformations with smaller PSD constraints. In this paper, we investigate a notion of *decomposed structured subsets* by combining chordal decomposition with DD/SDD approximations. This notion takes advantage of any underlying sparsity via chordal decomposition, while embracing the scalability of DD/SDD approximations. We discuss the applications of decomposed structured subsets to semidefinite optimization. Basis pursuit for refining DD/SDD approximations are also incorporated into the decomposed structured subset framework, and numerical performance is improved as compared to standard DD/SDD approximations. These results are demonstrated on H_∞ norm estimation problems for networked systems.

1. INTRODUCTION

Semidefinite programs (SDPs) are a class of convex optimization problems with a linear objective, affine constraints, and an additional positive semidefinite (PSD) constraint on the variable. SDPs include common optimization problems such as Linear Programs (LPs) and Second-order Cone Programs (SOCPs). A general conic program has the following primal and dual forms:

$$\begin{aligned} p^* = \min_{X} \quad & \langle C, X\rangle \\ & \langle A_i, X\rangle = b_i, i = 1, \dots, m, \\ & X \in K, \end{aligned} \tag{1a}$$

$$\begin{aligned} d^* = \max_{y,Z} \quad & \langle b, y\rangle \\ & Z + \sum_{i=1}^{m} y_i A_i = C, \\ & Z \in K^*, \end{aligned} \tag{1b}$$

where $C, A_1, \dots, A_m \in \mathbb{S}^n$ and $b \in \mathbf{R}^m$ are problem data, K is a proper cone with its dual as K^*, and $\langle \cdot, \cdot\rangle$ denotes the standard inner product. Semidefinite programming occurs over the self-dual cone $K = K^* = \mathbb{S}^n_+$ of positive semidefinite matrices.

Standard interior-point methods (IPMs) can solve an SDP to arbitrary precision in polynomial time, scaling as $O(n^2m^2 + n^3m)$ [Alizadeh, 1995]. When m is fixed, the speed of IPMs can be greatly improved by reducing the size of PSD cone $\mathbb{S}^n_+$. This motivates a variety of decomposition methods, which exploit problem structures to break up a large PSD constraint into a product of smaller PSD constraints. For example, sparsity in problem data (C, A_i) motivates a notion of chordal decomposition [Agler et al., 1988, Grone et al., 1984], and symmetry/common *-algebra structure of (C, A_i) restricts optimization to an invariant subspace [Vallentin, 2009].

A notion of *structured subset* method is to restrict (1a) to inner/outer cones $K_{\text{inner}} \subset \mathbb{S}^n_+ \subset K_{\text{outer}}$ to form optima $p^*_{\text{outer}} \le p^*_{\text{SDP}} \le p^*_{\text{inner}}$. Typical subset sets for K_{inner} are (scaled-) diagonally dominant (DD or SDD) matrices. Optimizing over structured subsets $K_{\text{inner}} \subset \mathbb{S}^n_+$ may lead to computationally simpler problems such as LPs or SOCPs. These approximations can be iteratively refined through a change of basis scheme [Hall, 2018]. Majumdar et al. [2019] provides an overview of decomposition methods and structured subsets for solving SDPs. Note that standard structured subsets (*e.g.* DD/SDD matrices) ignore any sparsity of the original problem. Large semidefinite programs may run into numerical issues, and structured subsets will retain the ill-conditioning. Exploiting chordal sparsity can lead to an equivalent problem with a set of smaller PSD constraints, but some PSD constraints

⋆ J. Miller and M. Sznaier were partially supported by NSF grants CNS–1646121, CMMI–1638234, IIS–1814631 and ECCSâĂŤ1808381 and AFOSR grant FA9550-19-1-0005.

may be still overly large and dominant the computational complexity.

In this paper, we merge structured subsets with decomposition methods to form *decomposed structured subsets*, a cone where each decomposed block is a member of a structured subset. In the framework of decomposed structured subsets, we first apply possible decomposition methods to exploit any underlying sparsity and structure in the problem, leading to an equivalent problem with smaller PSD constraints. Standard structured sets are then used to approximate the large PSD constraints. We show the notion of decomposed structured subsets is a strictly improved approximation to sparse PSD cone as compared to the standard structured sets, and that basis pursuit schemes tend to yield tighter approximations when applied to the decomposed problem. We note that this paper primarily focuses on chordal structures. Symmetry/*-algebra structure are applied in the sequel.

The rest of this paper is organized as follows. Section 2 introduces preliminaries regarding chordal decomposition and structured subsets. Section 3 unites these concepts with decomposed structured subsets and performs a containment analysis. Section 4 discusses how to apply decomposed structured subsets to semidefinite programs and the change of basis algorithm. This approach is demonstrated through H_∞ norm estimation of networked systems in Section 5. We conclude this paper in Section 6.

2. PRELIMINARIES

2.1 Structured Subsets

A basic structured subset of the PSD cone $\mathbb{S}^n_+$ is diagonal PSD matrices $\mathcal{D}$. Two additional subsets are the cones of diagonally dominant (DD) [Barker and Carlson, 1975] and scaled diagonally dominant (SDD) matrices [Boman et al., 2005]:

$$\begin{aligned} \mathcal{D}^n &= \{A \in \mathbb{S}^n : A = \mathrm{diag}(a_1, \ldots, a_n), a_i \geq 0\}, \\ \mathcal{DD}^n &= \{A \in \mathbb{S}^n : \ a_{ii} \geq \sum_{j \neq i} |a_{ij}|, i = 1, 2, \ldots, n\}, \\ \mathcal{SDD}^n &= \{A \in \mathbb{S}^n : \ \exists D \in \mathcal{D}^n \mid DAD \in \mathcal{DD}^n\}. \end{aligned} \tag{2}$$

It is known that $\mathcal{D}^n, \mathcal{DD}^n$, and $\mathcal{SDD}^n$ are inner approximations to $\mathbb{S}^n_+$, *i.e.*

$$\mathcal{D}^n \subset \mathcal{DD}^n \subset \mathcal{SDD}^n \subset \mathbb{S}^n_+. \tag{3}$$

Linear optimization over $\mathcal{D}^n$ and $\mathcal{DD}^n$ (*i.e.*, setting $K = \mathcal{D}^n$ or $K = \mathcal{DD}^n$ in (1a)) is an LP, and over $\mathcal{SDD}^n$ (*i.e.*, setting $K = \mathcal{SDD}^n$ in (1a)) is an SOCP [Ahmadi and Majumdar, 2017]. As there exist very efficient solvers for LPs and SOCPs, these inner approximations to SDPs can scale to very large-dimension problems. Figure 1 shows a PSD-representable feasible set (black), along with inner approximation found by optimizing over the cone $\mathcal{DD}$ (red). An outer (gray) LP approximation is also shown, corresponding to the dual cone $\mathcal{DD}^*$. This example is discussed further in Section 3.2.

Factor width matrices also form a structured subset of $\mathbb{S}^n_+$. A matrix $M \in \mathcal{FW}^n_k$ if $\exists U$ such that $M = UU^\mathsf{T}$ where each column of U has cardinality at most k [Boman et al., 2005]. An intuitive interpretation is that factor width-k matrices are the sum of $k \times k$ PSD matrices that are embedded in $n \times n$ larger matrices. Factor width matrices can be extended to partitions of indices; see [Zheng et al., 2019b] for details about block factor-width matrices. In Section 5, B_k is defined as the set of block factor-width 2 matrices where each block is of size k.

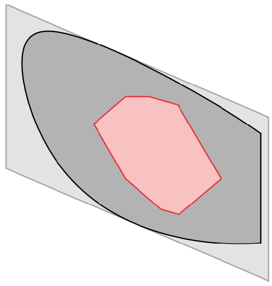

Fig. 1. Cone slices of Eq. (7) showing $\mathcal{DD} \subset \mathbb{S}_+ \subset \mathcal{DD}^*$

Change of Basis Change of Basis is an iterative method that refines an existing structured subset approximation for SDPs [Ahmadi and Hall, 2017]. The underlying idea is that a matrix X may not be a member of a structured subset K in one basis, but may have the correct form in another basis. Given a matrix L and a cone K, a basis-changed cone is $K(L) = \{X \mid LXL^\mathsf{T} \in K\}$. After finding an optimum X_0 of the conic optimization (1a), form a matrix factorization $X_0 = L_0 L_0^\mathsf{T}$, and then solve the modified problem:

$$\begin{aligned} X_1 = \operatorname*{argmin}_X \quad & \langle C, X\rangle \\ \text{subject to} \quad & \langle A_i, X\rangle = b_i, i = 1, \ldots, m, \\ & X \in K(L_0). \end{aligned} \tag{4}$$

This modified problem (4) has a new feasible region, where $X = I$ is an initial feasible point. Change of Basis proceeds as long as desired, forming the accumulated basis matrix $\hat{L} = L_0 L_1 \ldots L_{t-1} L_t$. An analogous process can occur on the dual side; see Hall [2018] for details. This method may reduce objective values between iterations, but is not guaranteed to converge to the true SDP optimum. If $K = \mathcal{DD}^n$ (or $K = \mathcal{SDD}^n$), then each step requires solving an LP (or SOCP) rather than an SDP.

2.2 Chordal Decomposition

In sparse SDPs, only a small number of entries of X are used in the cost C and constraints A_i. All other entries of X can be set arbitrarily to ensure X is PSD. The aggregate sparsity pattern of (C, A_i) can be encoded by a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where there is an edge between vertices i and j if any of $C, A_1, \ldots, A_m$ is nonzero at indices (i, j). A chord in a graph is an edge between two non-consecutive vertices in a cycle, and a graph is chordal if every cycle of length 4 or more has a chord [Vandenberghe et al., 2015]. Graphs that are not chordal can be chordal-extended by adding edges. A clique $\mathcal{C}$ is a set of vertices that forms a complete graph: $\forall v_i, v_j \in \mathcal{C}, (v_i, v_j) \in \mathcal{E}$. Maximal cliques are cliques that are not contained in any other cliques. The cardinality of a maximal clique is denoted as $|\mathcal{C}|$. Figure 2 shows a chordal graph and its maximal cliques $\mathcal{C}_k, k = 1, \ldots, 4$.

Given a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, let $\mathcal{E}^*$ be an edge set $\mathcal{E}$ augmented with self-loops. The cone of sparse symmetric matrices

$$\mathbb{S}^n(\mathcal{E}, 0) = \{X \in \mathbb{S}^n \mid \ X_{ij} = 0, \ \forall (i,j) \notin \mathcal{E}^*\},$$

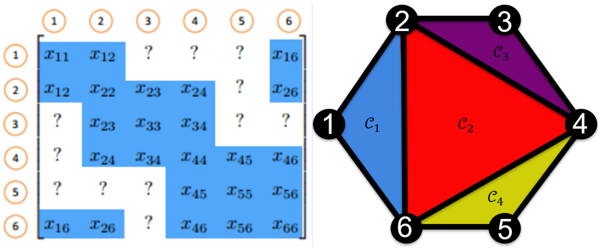

Fig. 2. Left: A 6×6 PSD completable cone $\mathbb{S}^n_+(\mathcal{E},?)$. Right: corresponding chordal graph $\mathcal{G}(\mathcal{V},\mathcal{E})$ with maximal cliques $\{\mathcal{C}_k\}_{k=1}^4$.

and its subcone of sparse PSD symmetric matrices is

$$\mathbb{S}^n_+(\mathcal{E},0) = \mathbb{S}(\mathcal{E},0)\cap\mathbb{S}^n_+.$$

The dual space $\mathbb{S}^n_+(\mathcal{E},?) = [\mathbb{S}^n_+(\mathcal{E},0)]^*$ are matrices with entries on $\mathcal{E}^*$ that can be completed into PSD matrices. Let $E_{\mathcal{C}_k} \in \mathbb{R}^{|\mathcal{C}_k|\times n}$ be 0/1 entry selector matrices that index out entries in clique $\mathcal{C}_k$. For sparse PSD matrices, we have the following two decomposition results.

Theorem 1 ([Grone et al., 1984]). *Let $\mathcal{G}(\mathcal{V},\mathcal{E})$ be a chordal graph with a set of maximal cliques $\{\mathcal{C}_1,\mathcal{C}_2,\dots,\mathcal{C}_p\}$. Then, $X\in\mathbb{S}^n_+(\mathcal{E},?)$ if and only if*

$$X_k = E_{\mathcal{C}_k} X E_{\mathcal{C}_k}^{\mathsf{T}} \in \mathbb{S}_+^{|\mathcal{C}_k|}, \qquad k=1,\,\dots,\,p.$$

Theorem 2 ([Agler et al., 1988]). *Let $\mathcal{G}(\mathcal{V},\mathcal{E})$ be a chordal graph with a set of maximal cliques $\{\mathcal{C}_1,\mathcal{C}_2,\dots,\mathcal{C}_p\}$. Then, $Z\in\mathbb{S}^n_+(\mathcal{E},0)$ if and only if there exist $Z_k\in\mathbb{S}_+^{|\mathcal{C}_k|}, k = 1,\,\dots,\,p$, such that*

$$Z = \sum_{k=1}^{p} E_{\mathcal{C}_k}^{\mathsf{T}} Z_k E_{\mathcal{C}_k}.$$

Theorem 1 breaks up a large sparse PSD constraint $X\in\mathbb{S}^n_+(\mathcal{E},?)$ into a series of smaller coupled PSD constraints $X_k\succeq 0, k=1,\dots,p$. This result can be applied to primal SDPs with a chordal sparsity pattern $\mathcal{E}$, *i.e.*, problem (1a) with $K=\mathbb{S}^n_+(\mathcal{E},?)$ can be decomposed as

$$\begin{aligned}\min_{X}\quad & \langle C,X\rangle\\ \text{subject to}\quad & \langle A_i,X\rangle = b_i, i=1,\dots,m,\\ & E_{\mathcal{C}_k} X E_{\mathcal{C}_k}^{\mathsf{T}}\in\mathbb{S}_+^{|\mathcal{C}_k|}, k=1,\dots,p.\end{aligned}\tag{5}$$

Analogous results can be obtained for sparse dual SDPs, with a characterization of the dual variable $Z\in\mathbb{S}^n_+(\mathcal{E},0)$ using Theorem 2. These decomposed SDPs can be solved using first order methods via variable splits $E_{\mathcal{C}_k}XE_{\mathcal{C}_k}^{\mathsf{T}} = X_k$ (see [Zheng et al., 2019a] for details), but interior point methods may suffer from the increase of the equality constraints introduced by the decomposition. Conversion utilities such as SparseCoLO [Fujisawa et al., 2009] internally perform domain and range space decompositions to take advantage of the chordal sparse structure.

3. DECOMPOSED STRUCTURED SUBSETS

This section introduces a natural idea to combine structured subsets with existing decomposition methods. Consider problem (1a) with $K=\mathbb{S}^6_+(\mathcal{E},?)$ where $\mathcal{E}$ is the sparsity pattern shown in Figure 2. Theorem 1 poses an optimization problem over the cliques $\{X_k\in\mathbb{S}^3_+\}_{k=1}^4$. Now consider a structured subset restriction. If we require $X=[x_{ij}]\in\mathcal{DD}^6$, this constraint requires $x_{11}\geq \sum_{i=2}^{6}|x_{1i}|$. Instead, if we consider a decomposition and impose structured subset restriction on the cliques, *e.g.*. $X_1\in\mathcal{DD}^3$, then it requires $x_{11}\geq|x_{12}|+|x_{16}|$, which is less restrictive than competing against all variables in the same row/column. *Decomposed Structured Subsets* arise from performing decompositions before applying structured subsets, and are presented in detail in this section.

3.1 Definition of decomposed structured subsets

Let $\mathcal{G}(\mathcal{V},\mathcal{E})$ be a graph with maximal cliques $\mathcal{C}_1,\dots,\mathcal{C}_p$. We define sparse DD and SDD matrices

$$\begin{aligned}\mathcal{DD}^n(\mathcal{E},0) &= \mathcal{DD}^n\cap\mathbb{S}^n(\mathcal{E},0),\\ \mathcal{SDD}^n(\mathcal{E},0) &= \mathcal{SDD}^n\cap\mathbb{S}^n(\mathcal{E},0).\end{aligned}$$

It follows that

$$\mathcal{DD}^n(\mathcal{E},0)\subset\mathcal{SDD}^n(\mathcal{E},0)\subset\mathbb{S}^n_+(\mathcal{E},0).$$

We therefore have the following decomposition result:

Proposition 1. *Let $\mathcal{G}(\mathcal{V},\mathcal{E})$ be a (not necessarily chordal) graph with a set of maximal cliques $\{\mathcal{C}_1,\mathcal{C}_2,\dots,\mathcal{C}_p\}$. Then,*

(1) $Z\in\mathcal{DD}^n(\mathcal{E},0)$ if and only if

$$Z=\sum_{k=1}^{p}E_{\mathcal{C}_k}^{\mathsf{T}}Z_kE_{\mathcal{C}_k},\quad Z_k\in\mathcal{DD}^{|\mathcal{C}_k|},\ k=1,\,\dots,\,p.$$

(2) $Z\in\mathcal{SDD}^n(\mathcal{E},0)$ if and only if

$$Z=\sum_{k=1}^{p}E_{\mathcal{C}_k}^{\mathsf{T}}Z_kE_{\mathcal{C}_k},\quad Z_k\in\mathcal{SDD}^{|\mathcal{C}_k|},\ k=1,\,\dots,\,p.$$

The proof is omitted for space reasons, but can be obtained by contacting the authors or from `arxiv.org/abs/1911.12859`. These results hold for an arbitrary clique edge cover that covers all maximal cliques, given that finding all maximal cliques is an NP-hard problem for generic graphs.

Motivated by Theorems 1 and 2, and Proposition 1, we let $\mathcal{E}$ be a sparsity pattern, $\mathcal{K}=\{K_k\}_{k=1}^p$ be a set of cones corresponding to a clique edge cover $\mathcal{C}_1,\dots,\mathcal{C}_p$, where each individual cone K_k is some structured subset in $\mathbb{S}^{|\mathcal{C}_k|}$. We define two *decomposed structured subsets*:

$$\begin{aligned}\mathcal{K}(\mathcal{E},0) &:= \Big\{Z\in\mathbb{S}^n \mid Z=\sum_{k=1}^{p}E_{\mathcal{C}_k}^{\mathsf{T}}Z_kE_{\mathcal{C}_k},\\ &\qquad\qquad Z_k\in K_k,\ k=1,\,\dots,\,p.\Big\},\\ \mathcal{K}(\mathcal{E},?) &:= \Big\{X\in\mathbb{S}^n\mid E_{\mathcal{C}_k}XE_{\mathcal{C}_k}^{\mathsf{T}}\in K_k,\ k=1,\,\dots,\,p.\Big\}.\end{aligned}\tag{6}$$

The decomposed structured subset $\mathcal{K}(\mathcal{E},0)$ is a generalization of $\mathcal{DD}^n(\mathcal{E},0),\mathcal{SDD}^n(\mathcal{E},0)$ and $\mathbb{S}^n_+(\mathcal{E},0)$:

$$\begin{aligned}\mathcal{K}(\mathcal{E},0) &= \mathcal{DD}^n(\mathcal{E},0), &&\text{if } K_k=\mathcal{DD}^{|\mathcal{C}_k|}, k=1,\dots,p.\\ \mathcal{K}(\mathcal{E},0) &= \mathcal{SDD}^n(\mathcal{E},0), &&\text{if } K_k=\mathcal{SDD}^{|\mathcal{C}_k|}, k=1,\dots,p.\end{aligned}$$

If $\mathcal{E}$ is chordal, the additional results hold for PSD cones:

$$\begin{aligned}\mathcal{K}(\mathcal{E},0) &= \mathbb{S}^n_+(\mathcal{E},0), &&\text{if } K_k=\mathbb{S}_+^{|\mathcal{C}_k|}, k=1,\dots,p\\ \mathcal{K}(\mathcal{E},?) &= \mathbb{S}^n_+(\mathcal{E},?), &&\text{if } K_k=\mathbb{S}_+^{|\mathcal{C}_k|}, k=1,\dots,p.\end{aligned}$$

The notion of decomposition structured subsets (6) gives more freedom to choose the individual cones K_k. We give a detailed analysis below.

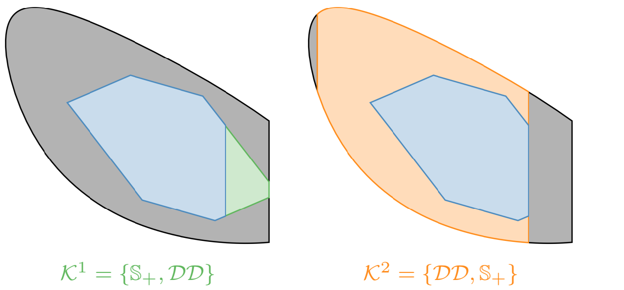

Fig. 3. Mixing cones broadens feasibility regions for $M(a,b) \in \mathcal{K}(\mathcal{E}, ?)$ in (7).

3.2 Containment Analysis

In Theorems 1 and 2 and Proposition 1, the cones corresponding to each clique are of the same type, *i.e.*, K_k are all either $\mathbb{S}_+^{|\mathcal{C}_k|}$ or $\mathcal{DD}^{|\mathcal{C}_k|}$. Additional freedom can be gained by allowing cliques to reside in different cones.

Given a graph with clique edge cover $\mathcal{C}_1, \ldots, \mathcal{C}_p$, we consider two sets of cones $\mathcal{K} = \{K_k\}_{k=1}^p$ and $\tilde{\mathcal{K}} = \{\tilde{K}_k\}_{k=1}^p$, where K_k or $\tilde{K}_k$ is a cone in $\mathbb{S}^{|\mathcal{C}_k|}$. We define the partial ordering $\subseteq$ on decomposed structured subsets:

$$\mathcal{K} \subseteq \tilde{\mathcal{K}} \text{ iff } K_k \subseteq \tilde{K}_k \quad \forall k = 1 \ldots p.$$

Then, we have the following proposition.

Proposition 2. *Given two sets of cones $\mathcal{K} = \{K_k\}_{k=1}^p$, and $\tilde{\mathcal{K}} = \{\tilde{K}_k\}_{k=1}^p$, if $\mathcal{K} \subseteq \tilde{\mathcal{K}}$, then we have*

$$\mathcal{K}(\mathcal{E}, 0) \subseteq \tilde{\mathcal{K}}(\mathcal{E}, 0) \quad \textit{and} \quad \mathcal{K}(\mathcal{E}, ?) \subseteq \tilde{\mathcal{K}}(\mathcal{E}, ?)$$

The result is true by definition. In the context of optimization, DD/SDD constraints offer scalable computation while PSD constraints are close (or exactly meet when the underlying graph is chordal) the true feasible region. Proposition 2 suggests some flexibility of choosing the individual cones K_k. As an example, consider the following matrix parameterized by (a, b):

$$M(a,b) = \begin{bmatrix} 1 & 1/2+a & ? & ? \\ 1/2+a & 2 & -2a & a+b \\ ? & -2a & 5 & b/2 \\ ? & a+b & b/2 & 2 \end{bmatrix}, \tag{7}$$

where ? denotes unspecified entries. The sparsity pattern of $M(a,b)$ has two maximal cliques: $\{2,3,4\}$ and $\{1,2\}$. By Theorem 1, $M(a,b) \in \mathbb{S}_+^4(\mathcal{E}, ?)$ if its two cliques are PSD:

$$M_1(a,b) = \underbrace{\begin{bmatrix} 2 & -2a & a+b \\ -2a & 5 & b/2 \\ a+b & b/2 & 2 \end{bmatrix}}_{\succeq 0}, M_2(a,b) = \underbrace{\begin{bmatrix} 1 & 1/2+a \\ 1/2+a & 2 \end{bmatrix}}_{\succeq 0}.$$

We can define feasibility sets for a cone-set $\mathcal{K}$ as $\{(a,b) \mid M_1(a,b) \in \mathcal{K}_1 \text{ and } M_2(a,b) \in \mathcal{K}_2\}$. Figure 3 compares feasibility sets when both cliques are $\mathcal{DD}$ (blue) and when both are $\mathbb{S}_+$ (black). As expected, the blue set is contained within the black set since $\mathcal{DD}^4(\mathcal{E}, ?) \subset \mathbb{S}_+^4(\mathcal{E}, ?)$. The orange set in the right panel has $M_1(a,b) \in \mathbb{S}_+^3$ and $M_2(a,b) \in \mathcal{DD}^2$. Note how the orange set includes the blue set (all $\mathcal{DD}$) and expands to nearly fill the left side of the black set (all $\mathbb{S}_+$). The green set in the left panel has $M_2(a,b) \in \mathbb{S}_+^2$ instead, which expands the all $\mathcal{DD}$ blue set with a small rightward bump.

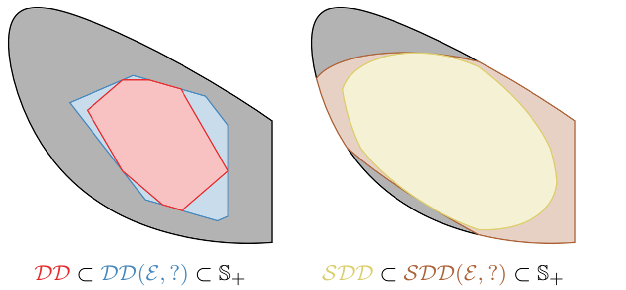

Fig. 4. $M(a,b)$ feasibility sets and containments.

Given $X \in \mathcal{DD}^n(\mathcal{E}, ?)$, we say X is DD-completable if $\exists X^c \in \mathcal{DD}^n$ such that X and X^c agree on entries in $\mathcal{E}$. A similar definition applies to $X \in \mathcal{SDD}^n(\mathcal{E}, ?)$. These generalize the concept of a PSD completion. For such an X with entries in $\mathcal{E}$, we write $X \in K^n$ if X is K-completable, and $X \in K^n(\mathcal{E}, ?)$ if each maximal clique $\mathcal{C}_k$ has $E_{\mathcal{C}_k}^{\mathsf{T}} X E_{\mathcal{C}_k} \in K^{|\mathcal{C}_k|}$. Under this definition:

Proposition 3. *Let $K^n \subset \tilde{K}^n$ be cones in $\mathbb{S}^n$. Given a sparsity pattern $\mathcal{E}$ outside of which are entries '?', the following containment holds:*

$$K^n \subseteq K^n(\mathcal{E}, ?), \quad \tilde{K}^n \subseteq \tilde{K}^n(\mathcal{E}, ?), \tag{8}$$

and $K^n(\mathcal{E}, ?) \subseteq \tilde{K}^n(\mathcal{E}, ?)$.

Note that by Theorem 1, a matrix X can be PSD-completable if and only if $X \in \mathbb{S}_+^n(\mathcal{E}, ?)$.

Figure 4 illustrates and compares feasibility sets for $M(a,b) \in K(\mathcal{E}, ?)$ (cliques of $M(a,b)$ in K) and $M(a,b) \in K$ ($M(a,b)$ has a K-completion). The blue $\mathcal{DD}^4(\mathcal{E}, ?)$ and black $\mathbb{S}_+(\mathcal{E}, ?)$ feasibility set are the same in Figure 4 as in 3. The left panel additionally shows feasible regions for the set $\mathcal{DD}^4$ (red) and $\mathcal{DD}^4(\mathcal{E}, ?)$. Constraining that $M(a,b)$ has a $\mathcal{DD}$-completion is stricter than restricting cliques to be $\mathcal{DD}$, so the feasibility sets will have $\mathcal{DD}^4 \subset \mathcal{DD}^4(\mathcal{E}, ?)$. The right plot echoes the left plot, where imposing that $M_1(a,b) \in \mathcal{SDD}, M_2(a,b) \in \mathcal{SDD}$ yields a broader feasibility set than $M(a,b) \in \mathcal{SDD}$.

4. APPLICATIONS TO SEMIDEFINITE OPTIMIZATION

Inner and outer approximations of semidefinite programs can be developed through decomposed structured subsets. A semidefinite program in primal form (1a) with ($X \in \mathbb{S}_+^n$) and dual form (1b) ($Z \in \mathbb{S}_+^n$) will have matching optima $p^* = d^*$ when strong duality holds. By complementary slackness, $\langle X, Z \rangle = 0$. Assume this semidefinite program has an aggregate sparsity pattern $\mathcal{E}$. With an optimization problem (1a) over $\mathbb{S}_+^n(\mathcal{E}, ?)$ and a cone set $\mathcal{K} = \{K_k\}_{k=1}^p$, an upper bound is attained by imposing $X \in \mathcal{K}(\mathcal{E}, ?)$, and a lower bound is found by restricting $Z \in \mathcal{K}(\mathcal{E}, 0)$ in the dual form. Lower bounds may also be found by setting $X \in \mathcal{K}^*(\mathcal{E}, ?)$, where $\mathcal{K}^* = \{K_k^*\}_{k=1}^p$.

The conic optimization problem for a decomposed structured subset $\mathcal{K}(\mathcal{E}, ?)$ is:

$$\begin{aligned} \min_X \quad & \langle C, X \rangle \\ \text{subject to} \quad & \langle A_i, X \rangle = b_i, i = 1, \ldots, m, \\ & E_{\mathcal{C}_k} X E_{\mathcal{C}_k}^{\mathsf{T}} \in K_k, \ k = 1, \ldots, p. \end{aligned} \tag{9}$$

4.1 Certifying Optimality

An outer structured subset approximation over $Z \in K^* \supset \mathbb{S}_+$ is tight (lower bound has the same optimum as the original SDP) if $Z \in \mathbb{S}_+$. Given an optimal $(X, y) \in (K, \mathbb{R}^m)$ for an inner approximation, this upper bound is likewise tight for if the optimal dual variable $Z \in \mathbb{S}_+$ for $Z = C - \sum_{i=1}^m y_i A_i$ [Ahmadi et al., 2017].

SDP-optimality of decomposed structured subsets can be certified in the same framework given a set of clique cones $\mathcal{K}$. When finding lower bounds to semidefinite programs, the clique cones $\mathcal{K}$ have $K_k \supseteq \mathbb{S}_+$. Tightness is certified if each $X_k \in \mathbb{S}_+$. Upper bounds have $K_k \subseteq \mathbb{S}_+$. For each clique $\mathcal{C}_k$ in the clique cone $\mathcal{K}$, check if the corresponding dual block $Z_k \in \mathbb{S}_+$. The dual clique blocks Z_k can be obtained by computing $Z_k = C_k - \sum_{i=1}^m y_i (A_i)_k$.

4.2 Decomposed Change of Basis

The change of basis algorithm in 2.1.1 can be extended to decomposed structured subsets. Let X_0 be the solution to Problem (9), and let $\mathcal{L} = \{L_k\}_{k=1}^p$ be Cholesky factorizations such that $L_k L_k^\top = E_{\mathcal{C}_k} X E_{\mathcal{C}_k}^\top$. The first iteration of change of basis will solve:

$$\begin{aligned} \min_X \quad & \langle C, X\rangle \\ \text{subject to} \quad & \langle A_i, X\rangle = b_i,\ i = 1, \ldots, m, \\ & X_k = E_k X E_k^T \\ L_k X_k L_k^T \in K_k & \end{aligned} \tag{10}$$

As before, an accumulated change of basis $\hat{L}_k$ can be formed for each clique, forming the agglomerated $\hat{\mathcal{L}}$. Different basis matrices $\hat{L}_k$ may cover the same entries due to clique overlap constraints, and this freedom may lead to better quality optima.

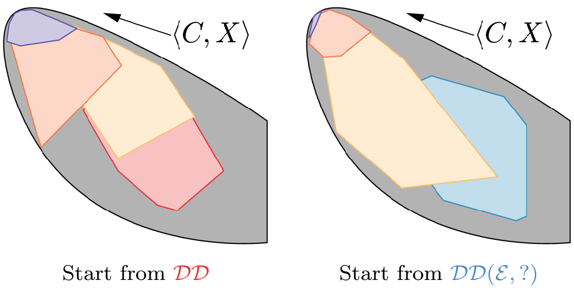

Fig. 5. Decomposed vs. Standard Change of Basis on (7)

Figure 5 illustrates the change of basis technique on a standard (left) and decomposed (right) structured subsets in the direction of $\langle C, X\rangle$. The intermediate costs are recorded in the below table for the first three iterations:

	Change of Basis Iteration			
	0	1	2	3
$\mathcal{DD}$	-1.41	-2.50	-3.08	-3.15
$\mathcal{DD}(\mathcal{E}, ?)$	-1.41	-3.02	-3.13	-3.17

A similar process can be done over the sparse cone $Z \in \mathcal{K}(\mathcal{E}, 0)$ (dual SDP), where bases $\mathcal{L} = \{L_k\}_{k=1}^p$ are tracked for each clique component Z_k forming the clique-sum $Z = \sum_k E_{\mathcal{C}_k}^\top Z_k E_{\mathcal{C}_k}$ for $Z \in \mathcal{K}(\mathcal{L})(\mathcal{E}, 0)$.

5. H-INFINITY NORM ESTIMATION FOR NETWORKED SYSTEMS

Consider a state-space stable dynamical system $G(s)$:

$$\begin{aligned} \dot{x} &= Ax + Bu, \\ y &= Cx + Du. \end{aligned}$$

The Bounded Real Lemma can establish an upper bound on H_∞ norm of $G(s)$: [Boyd et al., 1994]:

Theorem 3 (Bounded Real Lemma). *The following two statements are equivalent:*

(1) $\|G\|_\infty < \gamma$
(2) There exists a $P \succ 0$ *such that*

$$\begin{bmatrix} PA + A^\top P + C^\top C & P^\top B + C^\top D \\ B^\top P + D^\top C & -\gamma^2 I \end{bmatrix} \prec 0.$$

If the dynamical system is sparse (has a network structure), a dense $P \succ 0$ will give the tightest H_∞ approximation but will destroy the sparsity pattern. Choosing a P structure to be compatible with the LMI sparsity pattern will form a computationally tractable upper bound of $\|G(s)\|_\infty$. One such structure is a block-diagonal P where the size of each block is the corresponding agent's number of states [Zheng et al., 2018].

As an example of applying decomposed structured subsets to H_∞ estimation, we present the 'sea star' networked system. The sea star system is composed of a set of agents clustered into a head and a set of arms. Each agent has internal linear dynamics (n_i states, m_i inputs, d_i outputs), and they communicate and respond to a sparse selection of other agents. Figure 6 shows a sea star network with 70 densely connected agents in the head and other agents distributed into 12 arms. Each arm is composed of 2 densely connected 'knuckles'. Each knuckle has 10 agents, and every knuckle in the arm communicates with 4 agents in the next and previous knuckle (or the head as appropriate). The individual agent dynamics combine to form global dynamics $[A, B, C, D]$, where A is Hurwitz.

Estimating $\|G(s)\|_\infty = \|C(sI - A)^{-1}B + D\|_\infty$ can be accomplished by using the bounded real lemma to minimize γ^2. The resultant LMI has two semidefinite variables, as displayed in Figure 6. The top left corner of the Bounded Real LMI shows a structure induced by the network interconnections. On their own, the two semidefinite blocks are of size 1760 and 2691. This LMI system strongly exhibits chordal sparsity with edges $\mathcal{E}$, and can be posed as an optimization problem over the cone $\mathbb{S}_+(\mathcal{E}, 0)$.

Results of H_∞ norm estimation of the sea star system are presented in Figures 7 and 8. Columns are cones K where $K = \mathcal{DD}$ or $K = B_q$ if q is an integer 2.1. Rows are

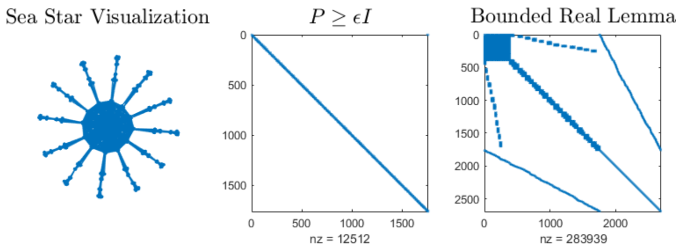

Fig. 6. Sea Star network topology and LMI sparsity

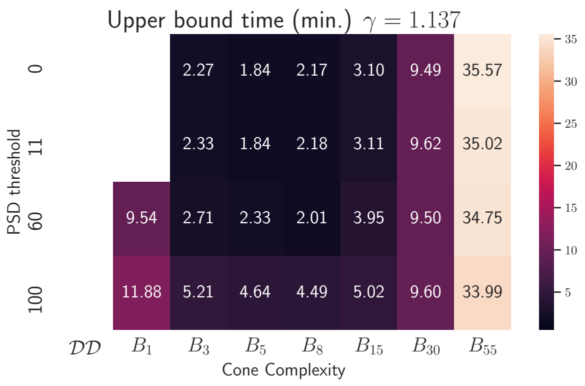

Fig. 7. Time to find γ by upper bound K (minutes)

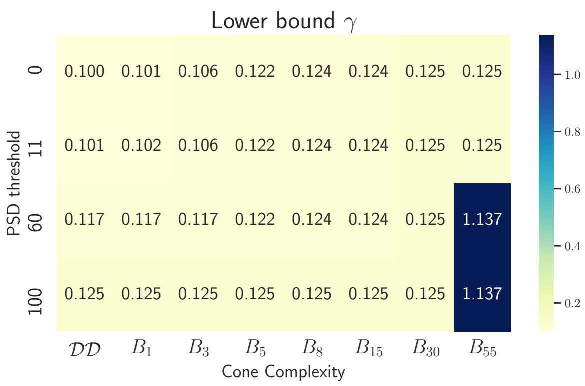

Fig. 8. γ found by lower bound K^*

size thresholds: the cone $K(\mathcal{E}, 0)$ has all cliques in K, and $\mathcal{K}_{60}(\mathcal{E}, 0)$ is a mixed cone where cliques with $|\mathcal{C}| \leq 60$ are PSD and $|\mathcal{C}| > 60$ are in K. All experiments were written in MatlabR2018a and performed on Mosek [Andersen and Andersen, 2000] on a Intel i7 CPU.

Times in Figure 7 were measured solving the primal program over K. All displayed values achieved the SDP optimal solution, as certified in Section 4.1. The cones $\mathcal{DD}$ with size thresholds 0 and 11 were primal infeasible, other non-displayed values did not attain the optimal $\gamma = 1.137$. The cone B_5 was fastest at 1.84 minutes. Figure 8 displays lower bounds for γ by over the dual cone $\mathcal{K}^*(\mathcal{E}, 0)$. Lower bounds tighten as cone complexity and size thresholds increase. The true γ is obtained with a block-size of 55 and size-thresholds of 60 and 100 taking 34.8 and 34.0 minutes. The computer running experiments ran out of memory attempting to solve the LMI over $\mathbb{S}_+(\mathcal{E}, 0)$.

6. CONCLUSIONS

Decomposition methods can break down large structured SDPs into simpler problems. Structured subsets allow for inner and outer approximations of dense SDPs to be quickly estimated. This paper combines the two approaches into decomposed structured subsets, which allow flexibility in choosing cones and form tighter objective approximations. Applications to semidefinite and network optimization are highlighted, specifically with an H_∞ norm estimation problem.

REFERENCES

Agler, J., Helton, W., and et al. (1988). Positive semidefinite matrices with a given sparsity pattern. *Linear algebra and its applications*, 107, 101–149.

Ahmadi, A.A., Dash, S., and Hall, G. (2017). Optimization over structured subsets of positive semidefinite matrices via column generation. *Discrete Optim.*, 24, 129–151.

Ahmadi, A.A. and Hall, G. (2017). Sum of squares basis pursuit with linear and second order cone programming. In *Algebraic and Geometric Methods in Discrete Mathematics*, chapter 2, 27–54. American Mathematical Society, Oxford.

Ahmadi, A.A. and Majumdar, A. (2017). Dsos and sdsos optimization: more tractable alternatives to sum of squares and semidefinite optimization. *arXiv preprint arXiv:1706.02586*.

Alizadeh, F. (1995). Interior point methods in semidefinite programming with applications to combinatorial optimization. *SIAM J OPTIMIZ*, 5(1), 13–51.

Andersen, E.D. and Andersen, K.D. (2000). The mosek interior point optimizer for linear programming: an implementation of the homogeneous algorithm. In *High performance optimization*, 197–232. Springer.

Barker, G. and Carlson, D. (1975). Cones of diagonally dominant matrices. *PAC J. MATH*, 57(1), 15–32.

Boman, E.G., Chen, D., Parekh, O., and Toledo, S. (2005). On factor width and symmetric h-matrices. *Linear algebra and its applications*, 405, 239–248.

Boyd, S., El Ghaoui, L., Feron, E., and Balakrishnan, V. (1994). *Linear matrix inequalities in system and control theory*, volume 15. Siam.

Fujisawa, K., Kim, S., Kojima, M., Okamoto, Y., and Yamashita, M. (2009). User's manual for sparsecolo: Conversion methods for sparse conic-form linear optimization problems. *Research Report B-453, Dept. of Math. and Comp. Sci. Japan, Tech. Rep.*, 152–8552.

Grone, R., Johnson, C.R., Sá, E.M., and Wolkowicz, H. (1984). Positive definite completions of partial hermitian matrices. *Linear Algebra Appl.*, 58, 109–124.

Hall, G. (2018). Optimization over nonnegative and convex polynomials with and without semidefinite programming. *arXiv preprint arXiv:1806.06996*.

Majumdar, A., Hall, G., and Ahmadi, A.A. (2019). A survey of recent scalability improvements for semidefinite programming with applications in machine learning, control, and robotics. *ArXiv*, abs/1908.05209.

Vallentin, F. (2009). Symmetry in semidefinite programs. *Linear Algebra and Its Applications*, 1(430), 360–369.

Vandenberghe, L., Andersen, M.S., et al. (2015). Chordal graphs and semidefinite optimization. *Foundations and Trends in Optimization*, 1(4), 241–433.

Zheng, Y., Fantuzzi, G., Papachristodoulou, A., Goulart, P., and Wynn, A. (2019a). Chordal decomposition in operator-splitting methods for sparse semidefinite programs. *Mathematical Programming*, 1–44.

Zheng, Y., Kamgarpour, M., Sootla, A., and Papachristodoulou, A. (2018). Scalable analysis of linear networked systems via chordal decomposition. In *2018 European Control Conference (ECC)*, 2260–2265. IEEE.

Zheng, Y., Sootla, A., and Papachristodoulou, A. (2019b). Block factor-width-two matrices and their applications to semidefinite and sum-of-squares optimization. *arXiv preprint arXiv:1909.11076*.