

Vision article

Chordal and factor-width decompositions for scalable semidefinite and polynomial optimization

Yang Zheng^{a,*}, Giovanni Fantuzzi^b, Antonis Papachristodoulou^c^a Department of Electrical and Computer Engineering, University of California, San Diego, CA 92093, United States of America^b Department of Aeronautics, Imperial College London, London, SW7 2AZ, UK^c Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK

ARTICLE INFO

Keywords:

Chordal sparsity
Semidefinite optimization
Polynomial optimization
Sum-of-squares
Matrix decomposition
Factor-width decomposition
Large-scale systems
Scalability

ABSTRACT

Chordal and factor-width decomposition methods for semidefinite programming and polynomial optimization have recently enabled the analysis and control of large-scale linear systems and medium-scale nonlinear systems. Chordal decomposition exploits the sparsity of semidefinite matrices in a semidefinite program (SDP), in order to formulate an equivalent SDP with smaller semidefinite constraints that can be solved more efficiently. Factor-width decompositions, instead, relax or strengthen SDPs with dense semidefinite matrices into more tractable problems, trading feasibility or optimality for lower computational complexity. This article reviews recent advances in large-scale semidefinite and polynomial optimization enabled by these two types of decomposition, highlighting connections and differences between them. We also demonstrate that chordal and factor-width decompositions allow for significant computational savings on a range of classical problems from control theory, and on more recent problems from machine learning. Finally, we outline possible directions for future research that have the potential to facilitate the efficient optimization-based study of increasingly complex large-scale dynamical systems.

1. Introduction

The design of innovative technology capable of addressing the challenges of the 21st century relies on the ability to analyze, predict, and control large-scale complex systems, which are typically nonlinear and may interact over networks (Astrom & Kumar, 2014; Murray, Astrom, Boyd, Brockett, & Stein, 2003). Convex optimization is one of the key tools for achieving these goals, because many questions related to the stability and operational safety of dynamical systems, the synthesis of optimal control policies, and the certification of robust performance can be posed as (or relaxed into) convex optimization problems. Very often, these take the form of semidefinite programs (SDPs)—a generalization of linear optimization problems with positive semidefinite matrix variables.

For linear systems, well-known methods based on linear matrix inequalities (LMIs) enable one to tackle a wide range of problems, including the study of stability, reachability, input-to-state and input-to-output properties, and the design of optimal and robust control strategies (Boyd, El Ghaoui, Feron, & Balakrishnan, 1994; Kailath, 1980; Zhou, Doyle, Glover, et al., 1996). Methods based on LMIs have been successfully applied across a broad spectrum of applications, including automotive applications (Rajamani, 2011), flight control (Giulietti, Pollini, & Innocenti, 2000), power grids (Riverso, Sarzo,

& Ferrari-Trecate, 2014; Sadabadi, Shafiee, & Karimi, 2016), and traffic systems (Li, Zheng, et al., 2017; Ploeg, Shukla, van de Wouw, & Nijmeijer, 2013; Zheng, Wang, & Li, 2020). For nonlinear systems with polynomial dynamics, SDP relaxations based on sum-of-squares polynomials (or, equivalently, moment sequences) enable stability analysis (Anderson & Papachristodoulou, 2015; Henrion & Garulli, 2005; Parrilo, 2000; Peet & Papachristodoulou, 2012), the estimation of regions of attractions (Chesi, 2011; Henrion & Korda, 2014; Korda, Henrion, & Jones, 2013; Topcu, Packard, Seiler, & Balas, 2009; Valmorbida & Anderson, 2017) and reachable sets (Jones & Peet, 2019; Magron, Garoche, Henrion, & Thirioux, 2019), safety verification (Miller, Henrion, & Szaier, 2021; Prajna, Jadbabaie, & Pappas, 2007), analysis of extreme or average behavior (Fantuzzi & Goluskin, 2020; Fantuzzi, Goluskin, Huang, & Chernyshenko, 2016; Goluskin, 2020; Korda, Henrion, & Mezić, 2021; Kuntz, Ottobre, Stan, & Barahona, 2016), and optimal control (Han & Tedrake, 2018; Henrion & Lasserre, 2006; Lasagna, Huang, Tutty, & Chernyshenko, 2016; Lasserre, Henrion, Prieur, & Trélat, 2008; Majumdar, Vasudevan, Tobenkin, & Tedrake, 2014; Prajna, Papachristodoulou, & Wu, 2004).

A widespread view since the 1990s (Boyd et al., 1994; Parrilo & Lall, 2003) is that, once a control problem is reformulated as an

* Corresponding author.

E-mail addresses: zhengy@eng.ucsd.edu (Y. Zheng), giovanni.fantuzzi10@imperial.ac.uk (G. Fantuzzi), antonis@eng.ox.ac.uk (A. Papachristodoulou).




Fig. 1.1. Illustration of chordal decomposition, where \succeq denotes positive semidefiniteness and $X \in \mathbb{S}_+^n(\mathcal{E}, ?)$ is a positive semidefinite completion constraint (see Section 2.2 for a precise definition). (a) A chordal graph with five vertices, six edges, and two maximal cliques (complete connected subgraphs), $C_1 = \{1, 2, 3\}$ and $C_2 = \{3, 4, 5\}$; (b) Chordal decomposition of a semidefinite constraint on a sparse matrix X into smaller positive semidefinite constraints on matrices Y_1, Y_2 with nonzero entries indexed by the cliques C_1 and C_2 ; (c) Chordal decomposition of a positive semidefinite completion constraint on a sparse matrix X with smaller positive semidefinite constraints on its principal submatrices X_1 and X_2 indexed by the cliques C_1 and C_2 .

SDP or relaxed into one, then it is effectively solved because SDPs can be solved computationally to global optimality using algorithms with polynomial-time complexity (Nemirovski, 2006; Nesterov, 2003; Nesterov & Nemirovski, 1994; Vandenberghe & Boyd, 1996; Ye, 2011). While this is true in theory, in practice solving large-scale SDPs accurately and without recourse to extremely large computational resources is a widely recognized challenge that remains largely unresolved. The need for scalable algorithms that can handle the very-large-scale SDPs encountered in real-life applications is even more pressing in today's world of complex, interconnected systems. Barring significant breakthrough, therefore, the relaxation of control problems into SDPs that can be solved with currently available algorithms requires further thought.

One particular bottleneck is the complexity of handling large semidefinite constraints; for instance, each iteration of classical interior-point algorithms requires $\mathcal{O}(n^3 m + n^2 m^2 + m^3)$ time and $\mathcal{O}(n^2 + m^2)$ memory (Nesterov, 2003, Section 4.3.3), where n is the size of semidefinite constraint and m is the number of equality constraints (see Section 3 for precise definitions). The majority of established general-purpose SDP solvers currently available, therefore, cannot handle large problems (e.g., with n larger than a few hundreds and m larger than a few thousands) on a regular computer. Consequently, the application of SDP-based frameworks for analysis and control is currently limited to medium-scale linear systems and small-scale nonlinear ones.

Overcoming these scalability issues is a problem that has received much attention in recent years (Ahmadi, Hall, Papachristodoulou, Saunderson, & Zheng, 2017; De Klerk, 2010; Majumdar, Hall, & Ahmadi, 2020; Vandenberghe & Andersen, 2015), and significant progress has been made through a number of different approaches. Most of them are related by a simple, yet powerful, underlying idea: *decompose a large positive semidefinite matrix X as a sum of structured ones, for which it is easier to impose positivity*.

One type of structured decomposition considers sums of low-rank matrices (Burer & Choi, 2006; Burer & Monteiro, 2003, 2005; Burer, Monteiro, & Zhang, 2002). Specifically, one writes $X = \sum_{i=1}^t v_i v_i^T$ for some vectors $v_1, \dots, v_t \in \mathbb{R}^n$, where $t \leq n$ is a parameter to be chosen, and optimizes over the choice of such vectors. Such a decomposition is guaranteed to exist for a properly chosen t , and there are explicit lower bounds on this parameter ensuring that the global minimum of the decomposed problem coincides with that of the original SDP (Barvinok, 1995; Pataki, 1998). However, while low-rank decomposition can bring considerable performance gains on large SDPs, it transforms a convex problem into a nonconvex one. Solution algorithms for the latter cannot be guaranteed to converge to the global minimum unless the original SDP is sufficiently “smooth” and t is large enough (Boumal, Voroninski, & Bandeira, 2020; Waldspurger & Waters, 2020).

A second type of structured decomposition, which we focus on in this paper, considers sums of sparse matrices. In this case, one writes $X = \sum_{i=1}^t Y_i$ for positive semidefinite matrices Y_1, \dots, Y_t that are nonzero only on a certain (and, ideally, small) principal submatrix. The choice of these principal submatrices is crucial in determining the particular type of matrix decomposition, as well as its properties. Two common selection strategies distinguish whether the original matrix X is dense or sparse.

If X is sparse, the principal submatrices are usually indexed by the maximal cliques of the sparsity graph of X ; these notions will be defined precisely in Section 2, but are illustrated in Fig. 1.1. When the sparsity graph is *chordal*, meaning that all cycles of length larger than three have an edge between nonconsecutive vertices, the existence of a clique-based decomposition is guaranteed (Agler, Helton, McCullough, & Rodman, 1988; Griewank & Toint, 1984; Kakimura, 2010). One can therefore replace the optimization of the large matrix X with the optimization of the matrices Y_1, \dots, Y_t without any loss of generality. Together with a dual result on the existence of positive semidefinite matrix completions (Grone, Johnson, Sá, & Wolkowicz, 1984), this *chordal decomposition* strategy enables one to significantly reduce the computational complexity of SDPs involving sparse positive semidefinite matrices (Fukuda, Kojima, Murota, & Nakata, 2001; Kim, Kojima, Mevissen, & Yamashita, 2011; Nakata, Fujisawa, Fukuda, Kojima, & Murota, 2003; Vandenberghe & Andersen, 2015).

When X is dense, instead, each matrix Y_i in the decomposition $X = \sum_{i=1}^t Y_i$ is chosen to be nonzero only on one of the $t = \binom{n}{k}$ possible $k \times k$ principal submatrices of X , where the parameter $k \geq 2$ is specified *a priori*. This type of decomposition leads to *factor-width- k* inner approximations of the positive semidefinite cone (Boman, Chen, Parekh, & Toledo, 2005), which are conservative but improve as k is increased. When $k \ll n$, optimizing over the matrices Y_1, \dots, Y_t , rather than over the original dense matrix X , leads to SDPs with small positive semidefinite cones, which can often be handled efficiently. In the extreme case $k = 2$, one obtains a second-order cone program, which can often be solved more efficiently in practice (Alizadeh & Goldfarb, 2003).

This paper offers a comprehensive review of chordal and factor-width- k decomposition methods, as well as of their application to large-scale semidefinite programming and polynomial optimization. Our goal is to introduce practitioners in control theory to the latest advances in these fields, which over the last decade or so have increased the scale of systems for which optimization-based frameworks for analysis and control can be implemented at a reasonable cost. Examples of problems that can now be handled efficiently include the analysis and synthesis of large-scale linear networked systems (Andersen, Pakazad, Hansson & Rantzer, 2014; Mason & Papachristodoulou, 2014; Zheng, Kamgarpour, Sootla & Papachristodoulou, 2018; Zheng, Mason & Papachristodoulou, 2018), the stability analysis and the approximation of regions of attraction for sparse nonlinear systems (Ahmadi & Majumdar, 2019; Schlosser & Korda, 2020; Tacchi, Cardozo, Henrion & Lasserre, 2019; Zheng, Fantuzzi & Papachristodoulou, 2019), optimal power flow in power grids (Andersen, Hansson, & Vandenberghe, 2014; Jabr, 2011; Molzahn, Holzer, Lesieutre, & DeMarco, 2013), and numerous problems in machine learning (Batten, Kouvaros, Lomuscio, & Zheng, 2021; Chen, Lasserre, Magron & Pauwels, 2020; Dahl, Vandenberghe, & Roychowdhury, 2008; Kim, Kojima, & Waki, 2009; Latorre, Rolland, & Cevher, 2020; Newton & Papachristodoulou, 2021; Zhang, Fattahi, & Sojoudi, 2018). We hope that knowledge of the advanced optimization techniques discussed here can assist control theorists in developing efficient modeling frameworks that can be applied much more widely and, crucially, to increasingly complex large-scale systems.

1.1. Outline

After introducing relevant graph-theoretic notions in Section 2, we discuss chordal decomposition methods for general SDPs in Section 3. Section 4 looks at decomposition methods for sparse polynomial optimization problems, which arise when relaxing analysis and control problems for nonlinear systems. Factor-width- k decompositions for dense matrices are discussed in Section 5. Section 6 presents examples of how matrix decomposition can be applied to some classical control problems and to some recent problems in machine learning. Section 7 draws conclusions and outlines possible directions for future research.

1.2. Basic notation

Mathematical symbols are defined as necessary in each of the following sections, but we summarize common notation here. The m -dimensional Euclidean space, the vector space of $n \times n$ real symmetric matrices, and the cone of $n \times n$ positive semidefinite symmetric matrices are denoted, respectively, by \mathbb{R}^m , \mathbb{S}^n , and \mathbb{S}_+^n . Angled brackets are used to denote the inner product in any of these spaces; in particular, $\langle x, y \rangle = x^T y$ when $x, y \in \mathbb{R}^m$ and $\langle X, Y \rangle = \text{trace}(XY)$ when $X, Y \in \mathbb{S}^n$. We often write $X \geq 0$ instead of $X \in \mathbb{S}_+^n$ when the matrix size is clear from the context or is unimportant, and write $X > 0$ if X is strictly positive definite.

2. Chordal graphs and matrix decomposition

This section reviews chordal graphs and their applications to sparse matrix decomposition. Matrix decomposition is central to many sparsity-exploiting techniques for semidefinite and polynomial optimization. Detailed introductions to chordal graphs can be found in the surveys by Blair and Peyton (1993) and Rose (1970), and in the monographs by Vandenberghe and Andersen (2015) and Golumbic (2004). We first introduce some graph-theoretic notions in Section 2.1, and then given an overview of classical matrix decomposition and completion results in Section 2.2. Extensions to sparse block-partitioned matrices are discussed in Section 2.3.

2.1. Chordal graphs

A graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is defined by a set of vertices $\mathcal{V} = \{1, 2, \dots, n\}$ and a set of edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. A graph \mathcal{G} is undirected if $(v_i, v_j) \in \mathcal{E}$ implies that $(v_j, v_i) \in \mathcal{E}$. A path in $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is a sequence of edges that connect a sequence of distinct vertices. A graph is *connected* if there is a path between any two vertices, and *complete* if any two vertices are connected by an edge, i.e., $\mathcal{E} = \mathcal{V} \times \mathcal{V}$. The subgraph induced by a subset of vertices $\mathcal{W} \subset \mathcal{V}$ is the undirected graph with vertices \mathcal{W} and edges $\mathcal{E} \cap (\mathcal{W} \times \mathcal{W})$. A subset of vertices $C \subseteq \mathcal{V}$ is called a *clique* if the subgraph induced by C is complete. If C is not contained in any other clique, it is a *maximal clique*. The number of vertices in C is denoted by $|C|$.

A *cycle* of length $k \geq 3$ in a graph \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* in a cycle is an edge connecting two nonconsecutive vertices.

Definition 2.1. An undirected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is *chordal* if every cycle of length $k \geq 4$ has at least one chord.

Examples of chordal graphs are given in Fig. 2.1. Observe also that many common types of graphs are chordal, including chains, acyclic undirected graphs (i.e., graphs with no cycles, such as trees), undirected graphs with cycles of length no greater than three, and complete graphs.

Chordal graphs have a number of properties that make them easy to handle computationally. For example, a connected chordal graph has at most $n-1$ maximal cliques, and they can be identified in linear




Fig. 2.1. Examples of chordal graphs: (a) A generic chordal graph. (b) A “banded” chordal graph; (c) A “block-arrow” graph. The names “banded” and “block-arrow” are motivated by the fact that, as explained in Section 2.2.1, these graphs describe the sparsity patterns of the matrices in Fig. 2.3.




Fig. 2.2. (a) A nonchordal graph: the cycle (1-2-3-4) is of length four but has no chords. (b) Minimal chordal extension obtained by adding edge (2,4). The maximal cliques are $C_1 = \{1, 2, 4\}$ and $C_2 = \{2, 3, 4\}$. (c) Minimal chordal extension obtained by adding edge (1,3). The maximal cliques are $C_1 = \{1, 2, 3\}$ and $C_2 = \{1, 3, 4\}$. (d) Trivial chordal extension by completion.

time with respect to the number of vertices and edges (Vandenberghe & Andersen, 2015) using, for instance, Algorithm 2 in Appendix C. In addition, any induced subgraph of a chordal graph is chordal because cycles in the subgraph are also cycles in the original graph. This is a useful fact in several induction proofs using chordality (Blair & Peyton, 1993, Section 2). Finally, chordal graphs admit a so-called *perfect elimination ordering* of the vertices, which is central to the zero fill-in property of Cholesky factorizations for sparse matrices. These two properties are reviewed in Appendix A.

Given the rich structure implied by chordality, it is very often convenient to extend a nonchordal graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ into a chordal graph $\hat{\mathcal{G}}(\mathcal{V}, \hat{\mathcal{E}})$ with larger edge set $\hat{\mathcal{E}} \supset \mathcal{E}$, which is called a *chordal extension* of \mathcal{G} . Usually, a graph admits many different chordal extensions, including the trivial one with edge set $\hat{\mathcal{E}} = \mathcal{V} \times \mathcal{V}$ obtained by completion, and the one obtained by completing only the graph’s connected components. Finding a *minimal* chordal extension, meaning that the smallest possible number of additional edges has been added, is an NP-complete problem (Yannakakis, 1981). However, approximately minimal chordal extensions can often be constructed in practice using heuristic strategies such as the maximum cardinality search (Berry, Blair, Heggernes, & Peyton, 2004) and the symbolic Cholesky factorization with an approximate minimum degree ordering (Fukuda et al., 2001; Vandenberghe & Andersen, 2015).

Fig. 2.2 illustrates these concepts. The graph in Fig. 2.2(a) is not chordal, but can be extended to the chordal graph in Fig. 2.2(b) by adding edge (2,4), edge (1,3), or both. The first two extensions are minimal, while the latter is the trivial extension by completion. The minimal chordal extension obtained by adding edge (2,4) has two maximal cliques, $C_1 = \{1, 2, 4\}$ and $C_2 = \{2, 3, 4\}$.

2.2. Sparse matrix decomposition

This subsection reviews two fundamental results on the decomposition of sparse positive semidefinite matrices whose sparsity can be described using chordal graphs.

2.2.1. Sparse symmetric matrices

Fix any positive integer n and set $\mathcal{V} = \{1, \dots, n\}$. Given an undirected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, we say that a symmetric matrix $X \in \mathbb{S}^n$ has a sparsity




Fig. 2.3. Sparsity patterns of 8×8 matrices corresponding to the chordal graphs in Figs. 2.1(a)–2.1(c), respectively (throughout this paper, \times denotes a real number).

graph \mathcal{G} (alternatively, sparsity pattern \mathcal{E}) if $X_{ij} = X_{ji} = 0$ when $(i, j) \notin \mathcal{E}$. We denote the space of sparse symmetric matrices by

$$\mathbb{S}^n(\mathcal{E}, 0) := \{X \in \mathbb{S}^n \mid X_{ij} = X_{ji} = 0, \text{ if } (i, j) \notin \mathcal{E}\}.$$

For example, the graph¹ in Fig. 2.2(b) describes the sparsity pattern of the matrix

$$X = \begin{bmatrix} X_{11} & X_{12} & 0 & X_{14} \\ X_{21} & X_{22} & X_{23} & X_{24} \\ 0 & X_{32} & X_{33} & X_{34} \\ X_{41} & X_{42} & X_{43} & X_{44} \end{bmatrix} \in \mathbb{S}^4, \quad (2.1)$$

where each entry X_{ij} may be nonzero or zero. Similarly, the symbolic matrices in Fig. 2.3 have sparsity patterns described by the graphs in Fig. 2.1.

Given $X \in \mathbb{S}^n(\mathcal{E}, 0)$, the diagonal elements X_{ii} and the off-diagonal elements X_{ij} with $(i, j) \in \mathcal{E}$ may be nonzero or zero. Thus, if $X \in \mathbb{S}^n(\mathcal{E}, 0)$ and $\hat{\mathcal{E}} \supset \mathcal{E}$ is an extension of the edge set, then we also have $X \in \mathbb{S}^n(\hat{\mathcal{E}}, 0)$. In this paper, we are especially interested in chordal extensions of sparsity pattern. For simplicity, we will say that a matrix X has a chordal sparsity pattern if its corresponding sparsity graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is chordal. Of course, this can always be achieved via chordal extension.

In what follows, it will be convenient to refer to particular principal submatrices of a sparse matrix, indexed by the maximal cliques of its sparsity graph. Given a clique C_k of $\mathcal{G}(\mathcal{V}, \mathcal{E})$, we define a matrix $E_{C_k} \in \mathbb{R}^{|C_k| \times n}$ with entries

$$(E_{C_k})_{ij} = \begin{cases} 1, & \text{if } C_k(i) = j, \\ 0, & \text{otherwise,} \end{cases} \quad (2.2)$$

where $C_k(i)$ is the i th vertex.² Given $X \in \mathbb{S}^n$, the definition of E_{C_k} implies that the operation $E_{C_k} X E_{C_k}^\top \in \mathbb{S}^{|C_k|}$ extracts the principal submatrix of X indexed by the clique C_k . Conversely, the operation $E_{C_k}^\top Y E_{C_k}$ “inflates” a $|C_k| \times |C_k|$ matrix Y into a sparse $n \times n$ symmetric matrix that has Y as its principal submatrix indexed by C_k , and is zero otherwise. For example, the chordal graph in Fig. 2.2(b) has a maximal clique $C_1 = \{1, 2, 4\}$, and the corresponding matrix E_{C_1} is

$$E_{C_1} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

¹ Throughout, we assume that each vertex has a self-loop, unless otherwise noted. We omit the self-loops when plotting a graph.

² The elements of C_k can be sorted in any convenient order. We implicitly use the natural ordering in this work, but using a different one simply amounts to a permutation of the rows of E_{C_k} .

For the sparse matrix $X \in \mathbb{S}^4$ in (2.1) and any matrix $Y \in \mathbb{S}^3$, we have

$$E_{C_1} X E_{C_1}^\top = \begin{bmatrix} X_{11} & X_{12} & X_{14} \\ X_{21} & X_{22} & X_{24} \\ X_{41} & X_{42} & X_{44} \end{bmatrix},$$

$$E_{C_1}^\top Y E_{C_1} = \begin{bmatrix} Y_{11} & Y_{12} & 0 & Y_{13} \\ Y_{21} & Y_{22} & 0 & Y_{23} \\ 0 & 0 & 0 & 0 \\ Y_{31} & Y_{32} & 0 & Y_{33} \end{bmatrix}.$$

2.2.2. Cone of sparse positive semidefinite matrices

Denote the set of positive semidefinite matrices with sparsity pattern \mathcal{E} by

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

This set is a convex cone because it is the intersection of a subspace and a convex cone. If $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is a chordal graph, $\mathbb{S}_+^n(\mathcal{E}, 0)$ can be represented using smaller but coupled convex cones, as stated in the following result (Agler et al. 1988, Theorem 2.3; Griewank and Toint 1984, Theorem 4; Kakimura 2010, Theorem 1).

Theorem 2.1. Let $\mathcal{G}(\mathcal{V}, \mathcal{E})$ be a chordal graph with maximal cliques C_1, \dots, C_t . Then, $Z \in \mathbb{S}_+^n(\mathcal{E}, 0)$ if and only if there exist matrices $Z_k \in \mathbb{S}_+^{|C_k|}$ for $k = 1, \dots, t$ such that

$$Z = \sum_{k=1}^t E_{C_k}^\top Z_k E_{C_k}. \quad (2.3)$$

The “if” part of Theorem 2.1 is immediate, since a sum of positive semidefinite matrices is positive semidefinite. The “only if” part, instead, can be proven using the zero fill-in property of sparse Cholesky factorization for $Z \in \mathbb{S}_+^n(\mathcal{E}, 0)$ (Vandenberghe & Andersen, 2015, Section 9.2); see Appendices A and B for details. A similar elementary proof given by Kakimura (2010), based on simple linear algebra and perfect elimination orderings for chordal graphs, reveals that one can impose a rank constraint in the decomposition (2.3): there exist Z_k with $\text{rank}(Z) = \sum_{k=1}^t \text{rank}(Z_k)$ such that (2.3) holds.

Remark 2.1. The chordality assumption in Theorem 2.1 is necessary. For every nonchordal pattern \mathcal{E} , while particular matrices in $\mathbb{S}_+^n(\mathcal{E}, 0)$ admit the decomposition (2.3), there always exist matrices in $\mathbb{S}_+^n(\mathcal{E}, 0)$ that do not; see Vandenberghe and Andersen (2015, p. 342) for an explicit example. In addition, the decomposition (2.3) generally requires all maximal cliques C_1, \dots, C_t , even when a subset of maximal cliques has already covered the sparsity pattern \mathcal{E} (that is $\mathcal{E} = \bigcup_{k \in I} C_k \times C_k$ with $I \subset \{1, \dots, t\}$). An example of this is given in Appendix C. ■

Example 2.1. Consider the positive semidefinite matrix

$$Z = \begin{bmatrix} 2 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 2 \end{bmatrix}, \quad (2.4)$$

whose sparsity graph is a chordal chain graph with three vertices, edge set $\mathcal{E} = \{(1, 1), (2, 2), (3, 3), (1, 2), (2, 3)\}$, and maximal cliques $C_1 = \{1, 2\}$ and $C_2 = \{2, 3\}$. Theorem 2.1 guarantees that the decomposition (2.3) exists. Indeed, we have

$$E_{C_1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}, \quad E_{C_2} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

and

$$Z = \underbrace{E_{C_1}^\top \begin{bmatrix} 2 & 1 \\ 1 & 0.5 \end{bmatrix} E_{C_1}}_{Z_1 \geq 0} + \underbrace{E_{C_2}^\top \begin{bmatrix} 0.5 & 1 \\ 1 & 2 \end{bmatrix} E_{C_2}}_{Z_2 \geq 0}.$$

This decomposition satisfies the rank constraint mentioned above since $\text{rank}(Z) = 2$ and $\text{rank}(Z_1) = \text{rank}(Z_2) = 1$. ■




Fig. 2.4. Joint feasible set of the decomposed LMIs in (2.6): (a) projection onto the (x_1, x_2) plane; (b) projection onto the (x_1, d) plane; (c) projection onto the (x_2, d) plane. Panel (a) also shows the boundary (thick black line) of the feasible set of the original 3×3 LMI (2.5).

Example 2.2. Given a variable $x \in \mathbb{R}^2$, consider the 3×3 linear matrix inequality (LMI)

$$Z(x) := \begin{bmatrix} 2x_1 & x_1 + x_2 & 0 \\ x_1 + x_2 & 5 - x_1 - x_2 & x_1 \\ 0 & x_1 & x_2 + 1 \end{bmatrix} \geq 0. \quad (2.5)$$

This LMI has the same chordal sparsity pattern as the matrix in (2.4). Consequently, Theorem 2.1 implies that (2.5) holds if and only if there exist matrices

$$Z_1 := \begin{bmatrix} a & b \\ b & c \end{bmatrix} \geq 0 \quad \text{and} \quad Z_2 := \begin{bmatrix} d & e \\ e & f \end{bmatrix} \geq 0$$

such that

$$\begin{bmatrix} a & b & 0 \\ b & c+d & e \\ 0 & e & f \end{bmatrix} = Z(x).$$

After eliminating the variables a, b, c, e and f using this matching condition, we conclude that (2.5) holds if and only if there exists d such that

$$\begin{bmatrix} 2x_1 & x_1 + x_2 \\ x_1 + x_2 & 5 - x_1 - x_2 - d \end{bmatrix} \geq 0, \quad (2.6)$$

$$\begin{bmatrix} d & x_1 \\ x_1 & x_2 + 1 \end{bmatrix} \geq 0.$$

Fig. 2.4 shows two-dimensional projections of the three-dimensional feasible set of the two LMIs in (2.6). As expected, the projection on the (x_1, x_2) plane coincides with the feasible set of LMI (2.5), which is contained inside the thick black line in Fig. 2.4(a). This confirms that the LMIs in (2.6) are equivalent to the LMI (2.5). Therefore, we have decomposed a 3×3 LMI into two coupled LMIs of size 2×2 . ■

2.2.3. Cone of positive-semidefinite-completable matrices

A concept related to the matrix decomposition above is that of positive semidefinite matrix completion. Given a matrix $X \in \mathbb{S}^n$, let

$$\mathbb{P}_{\mathbb{S}^n(\mathcal{E},0)}(X) = \begin{cases} X_{ij} & \text{if } (i, j) \in \mathcal{E}, \\ 0 & \text{otherwise} \end{cases} \quad (2.7)$$

be its projection onto the space of sparse matrices $\mathbb{S}^n(\mathcal{E}, 0)$ with respect to the Frobenius matrix norm. We define the cone

$$\mathbb{S}_+^n(\mathcal{E}, ?) := \mathbb{P}_{\mathbb{S}^n(\mathcal{E},0)}(\mathbb{S}_+^n).$$

Using (2.7), it is not hard to see that a sparse matrix X is in $\mathbb{S}_+^n(\mathcal{E}, ?)$ if and only if it has a positive semidefinite completion, meaning that some (or all) of the zero entries X_{ij} with $(i, j) \notin \mathcal{E}$ can be replaced with nonzeros to obtain a positive semidefinite matrix \bar{X} . We call \bar{X} the *completion* of X and refer to $\mathbb{S}_+^n(\mathcal{E}, ?)$ as the cone of positive-semidefinite-completable matrices.

Remark 2.2 (Nonuniqueness of the Positive Semidefinite Completion). The positive semidefinite completion of a matrix $X \in \mathbb{S}_+^n(\mathcal{E}, ?)$ with sparsity pattern \mathcal{E} is generally not unique. For a chordal sparsity pattern \mathcal{E} , two widely used and efficient strategies to compute a completion \bar{X} are the maximum determinant completion (Vandenberghe & Andersen, 2015, Chapter 10.2), which maximizes $\det(\bar{X})$, and the minimum rank completion (see Dancis 1992; Jiang 2017; Sun 2015, Chapter 3.3), which minimizes $\text{rank}(\bar{X})$. In particular, there exists a positive semidefinite completion \bar{X} whose rank agrees with the maximum rank of the principal submatrices $E_{C_i} X E_{C_i}^T$ (Dancis, 1992, Theorem 1.5), i.e.,

$$\text{rank}(\bar{X}) = \max_{k=1,2,\dots,t} \text{rank}(E_{C_k} X E_{C_k}^T). \quad (2.8)$$

For any undirected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, the cones $\mathbb{S}_+^n(\mathcal{E}, ?)$ and $\mathbb{S}_+^n(\mathcal{E}, 0)$ are dual to each other with respect to the trace inner product $\langle X, Z \rangle = \text{trace}(XZ)$ in the space $\mathbb{S}^n(\mathcal{E}, 0)$ (Vandenberghe & Andersen, 2015, Chapter 10). To see this, observe that

$$\begin{aligned} (\mathbb{S}_+^n(\mathcal{E}, ?))^* &= \{Z \in \mathbb{S}^n(\mathcal{E}, 0) \mid \langle X, Z \rangle \geq 0 \forall X \in \mathbb{S}_+^n(\mathcal{E}, ?)\} \\ &= \{Z \in \mathbb{S}^n(\mathcal{E}, 0) \mid \langle \mathbb{P}_{\mathbb{S}^n(\mathcal{E},0)}(X), Z \rangle \geq 0 \forall X \geq 0\} \\ &= \{Z \in \mathbb{S}^n(\mathcal{E}, 0) \mid \langle X, Z \rangle \geq 0 \forall X \geq 0\} \\ &= \{Z \in \mathbb{S}^n(\mathcal{E}, 0) \mid Z \geq 0\} \\ &= \mathbb{S}_+^n(\mathcal{E}, 0). \end{aligned}$$

For a chordal matrix sparsity pattern, Theorem 2.1 on the decomposition of the cone $\mathbb{S}_+^n(\mathcal{E}, 0)$ can be dualized to obtain the following characterization of $\mathbb{S}_+^n(\mathcal{E}, ?)$, first proved by Grone et al. (1984, Theorem 7).

Theorem 2.2. Let $\mathcal{G}(\mathcal{V}, \mathcal{E})$ be a chordal graph with maximal cliques C_1, \dots, C_t . Then, $X \in \mathbb{S}_+^n(\mathcal{E}, ?)$ if and only if

$$E_{C_k} X E_{C_k}^T \in \mathbb{S}_+^{|C_k|} \quad \forall k = 1, \dots, t. \quad (2.9)$$

The “only if” part of Theorem 2.2 is immediate, since any principal submatrix of a positive semidefinite matrix is positive semidefinite. The “if” part, instead, relies on the properties of chordal graphs and, as mentioned above, can be proven by combining the duality between $\mathbb{S}_+^n(\mathcal{E}, 0)$ and $\mathbb{S}_+^n(\mathcal{E}, ?)$ with Theorem 2.1 (Vandenberghe & Andersen, 2015, p. 357). Precisely,

$$\begin{aligned} X \in \mathbb{S}_+^n(\mathcal{E}, ?) &\Leftrightarrow \langle X, Z \rangle \geq 0 \quad \forall Z \in \mathbb{S}_+^n(\mathcal{E}, 0), \\ &\Leftrightarrow \left\langle X, \sum_{k=1}^t E_{C_k}^T Z_k E_{C_k} \right\rangle \geq 0 \quad \forall Z_k \in \mathbb{S}_+^{|C_k|}, \\ &\Leftrightarrow \sum_{k=1}^t \langle E_{C_k} X E_{C_k}^T, Z_k \rangle \geq 0 \quad \forall Z_k \in \mathbb{S}_+^{|C_k|}, \\ &\Leftrightarrow E_{C_k} X E_{C_k}^T \in \mathbb{S}_+^{|C_k|} \quad \forall k = 1, \dots, t. \end{aligned}$$

The first equivalence expresses the duality between $\mathbb{S}_+^n(\mathcal{E}, 0)$ and $\mathbb{S}_+^n(\mathcal{E}, ?)$, the second one follows from Theorem 2.1, and the third one follows from the cyclic property of the trace operator: $\text{trace}(MN) = \text{trace}(NM)$ for any matrices M, N of compatible dimensions. Fig. 2.5 illustrates how the duality between $\mathbb{S}_+^n(\mathcal{E}, 0)$ and $\mathbb{S}_+^n(\mathcal{E}, ?)$ is mirrored in the duality between Theorem 2.1 and Theorem 2.2 for chordal graphs.

Example 2.3. Consider the symmetric matrix

$$X = \begin{bmatrix} 2 & 1 & 0 \\ 1 & 0.5 & 1 \\ 0 & 1 & 2 \end{bmatrix},$$

whose sparsity pattern is the (by now usual) 3-node chordal chain graph with maximal cliques $C_1 = \{1, 2\}$ and $C_2 = \{2, 3\}$. It is easy to check that, while X is not positive semidefinite, the principal submatrices indexed by the cliques C_1 and C_2 are. Then, Theorem 2.2




Fig. 2.5. Summary of duality between $\mathbb{S}_+^n(\mathcal{E}, 0)$ and $\mathbb{S}_+^n(\mathcal{E}, ?)$ and duality between Theorems 2.1 and 2.2 for a chordal graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with maximal cliques C_1, \dots, C_t .

guarantees that $X \in \mathbb{S}_+^n(\mathcal{E}, ?)$, meaning that the zero entries may be replaced by nonzeros to obtain a positive semidefinite matrix \bar{X} . One possible positive semidefinite completion is

$$\bar{X} = \begin{bmatrix} 2 & 1 & 2 \\ 1 & 0.5 & 1 \\ 2 & 1 & 2 \end{bmatrix}.$$

In fact, this is the minimum-rank completion whose rank, $\text{rank}(X) = 1$, coincides with the maximum rank of individual principal submatrices of X (cf. Remark 2.2). ■

Example 2.4. Consider the problem of finding a variable $x \in \mathbb{R}^3$ such that the matrix

$$X(x) := \begin{bmatrix} 1 - x_1 & x_1 + x_2 & 0 \\ x_1 + x_2 & x_2 & x_2 + x_3 \\ 0 & x_2 + x_3 & 2x_3 + 1 \end{bmatrix} \quad (2.10)$$

admits a positive semidefinite completion. This is equivalent to finding $x \in \mathbb{R}^3$ as well as a corresponding scalar $y \in \mathbb{R}$ such that

$$\begin{bmatrix} 1 - x_1 & x_1 + x_2 & y \\ x_1 + x_2 & x_2 & x_2 + x_3 \\ y & x_2 + x_3 & 2x_3 + 1 \end{bmatrix} \geq 0. \quad (2.11)$$

Since the sparsity graph of $X(x)$ is chordal, Theorem 2.2 implies that (2.10) is equivalent to the two LMIs

$$\begin{bmatrix} 1 - x_1 & x_1 + x_2 \\ x_1 + x_2 & x_2 \end{bmatrix} \geq 0, \quad \begin{bmatrix} x_2 & x_2 + x_3 \\ x_2 + x_3 & 2x_3 + 1 \end{bmatrix} \geq 0. \quad (2.12)$$

Feasible vectors x for the first of these two LMIs can be found by imposing $1 - x_1 + x_2 \geq 0$ and $(1 - x_1)x_2 - (x_1 + x_2)^2 \geq 0$, while feasible x for the second LMI are found by requiring $x_2 + 2x_3 + 1 \geq 0$ and $x_2(2x_3 + 1) - (x_2 + x_3)^2 \geq 0$. The feasible sets obtained in each case are illustrated by the red and green regions in Fig. 2.6, respectively. The blue region in the figure, instead, represents the three-dimensional set of feasible x for (2.11). As expected from Theorem 2.2, this is exactly the intersection of the feasible regions for the two LMIs in (2.12). Similar to Example 2.2, one can therefore replace the original 3×3 completion constraint—which is equivalent to LMI (2.11)—with the two 2×2 LMIs in (2.12) without any loss of generality. ■

2.3. Block-partitioned matrices

Theorems 2.1 and 2.2 can be extended to block-partitioned matrices characterized by block-sparsity. Such matrices arise, for example, when modeling network systems (cf. Section 6.1), where each block in the partition corresponds to an individual subsystem and sparsity in the network connectivity translates into block-sparsity. Block-partitioned matrices are also useful in extending factor-width decomposition that will be discussed in Section 5.




Fig. 2.6. Region of \mathbb{R}^3 where the matrix $X(x)$ in (2.10) admits a positive semidefinite completion (blue shading). This region coincides with the intersection of the region of \mathbb{R}^3 where the first LMI in (2.12) is feasible (red shading; the region extends to infinity in the x_3 direction) and the cylindrical region of \mathbb{R}^3 where the second LMI in (2.12) is feasible (green shading; the region extends to infinity in the x_1 direction). Thick red and green lines highlight the cross section of these two regions.

2.3.1. Sparse block matrices

Given a positive integer n , any finite set of positive integers $\alpha = \{\alpha_1, \alpha_2, \dots, \alpha_p\}$ is called a *partition* of n if $\sum_{i=1}^p \alpha_i = n$. The set of all possible partitions of n can be equipped with the following (partial) order relation.

Definition 2.2. Let $\alpha = \{\alpha_1, \dots, \alpha_p\}$ and $\beta = \{\beta_1, \dots, \beta_q\}$ be two partitions of an integer n with $p < q$. We say that β is *finer* than α (and α is *coarser* than β), denoted by $\beta \sqsubset \alpha$, if there exist integers $\{m_1, m_2, \dots, m_{p+1}\}$ with $m_1 = 1$, $m_{p+1} = q + 1$ and $m_i < m_{i+1}$ for $i = 1, \dots, p$ such that $\alpha_i = \sum_{j=m_i}^{m_{i+1}-1} \beta_j$ for all $i = 1, \dots, p$.

Essentially, a finer partition β breaks some entries of α into smaller ones (conversely, a coarser partition α is obtained by summing some consecutive entries of β). For example, the partitions $\alpha = \{4, 2\}$, $\beta = \{2, 2, 2\}$ and $\gamma = \{1, 1, 1, 1, 1, 1\}$ of $n = 6$ satisfy $\gamma \sqsubset \beta \sqsubset \alpha$.

Given any integer n and any partition $\alpha = \{\alpha_1, \dots, \alpha_p\}$ of n , a matrix $M \in \mathbb{R}^{n \times n}$ can be written in the block form

$$M = \begin{bmatrix} M_{11} & M_{12} & \dots & M_{1p} \\ M_{12} & M_{22} & \dots & M_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ M_{p1} & M_{p2} & \dots & M_{pp} \end{bmatrix}$$

with $M_{ij} \in \mathbb{R}^{\alpha_i \times \alpha_j}$ for all $i, j = 1, \dots, p$. For the finest partition $\alpha = \{1, \dots, 1\} = \mathbf{1}_n$, the block M_{ij} reduces to the entry (i, j) of M . As shown below and in Section 5.3, however, the freedom to consider a nontrivial partition offers considerable flexibility when devising decomposition strategies for a large matrix M . In particular, by refining or coarsening a partition one can in principle split a matrix into blocks of optimal size for the computational resources at one's disposal.

The block sparsity pattern of an $n \times n$ matrix M whose blocks are defined by a partition $\alpha = \{\alpha_1, \dots, \alpha_p\}$ of n can be described using a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with $\mathcal{V} = \{1, \dots, p\}$ and edge set such that $M_{ij} = 0$ if $(i, j) \notin \mathcal{E}$, where M_{ij} is the (i, j) -th block in M and 0 denotes a zero block of appropriate size. We call α a *chordal partition* if the corresponding block sparsity graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is chordal. The linear space of sparse symmetric block matrices with a prescribed block sparsity




Fig. 2.7. (a) Nonchordal sparsity graph of the 9×9 matrix M in Remark 2.3. Blue nodes and edges form a cycle of length 6 with no chord. (b) Chordal block sparsity graph of the same matrix with partition $\alpha_1 = \{2, 1, 1, 2, 2, 1\}$. (c) Chordal block sparsity graph of the same matrix with partition $\alpha_2 = \{3, 3, 3\}$.

pattern \mathcal{E} is then given by

$$\mathbb{S}_{\alpha}^n(\mathcal{E}, 0) := \{M \in \mathbb{S}^n \mid M_{ij} = 0 \text{ if } (i, j) \notin \mathcal{E}\}.$$

The block-sparse positive semidefinite cone and the block-sparse positive-semidefinite-completable cone are simply

$$\mathbb{S}_{\alpha,+}^n(\mathcal{E}, 0) := \mathbb{S}_{\alpha}^n(\mathcal{E}, 0) \cap \mathbb{S}_{+}^n, \quad (2.13a)$$

$$\mathbb{S}_{\alpha,+}^n(\mathcal{E}, ?) := \mathbb{P}_{\mathbb{S}_{\alpha}^n(\mathcal{E}, 0)}(\mathbb{S}_{+}^n). \quad (2.13b)$$

Remark 2.3 (Chordal Partitions and Chordal Extension). If M is a sparse matrix with nonchordal sparsity pattern, it is often possible to find one or more chordal partitions α . An example is the 9×9 symbolic matrix

$$M = \begin{bmatrix} \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \end{bmatrix} = \begin{bmatrix} \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \end{bmatrix} = \begin{bmatrix} \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \end{bmatrix}$$

where the partitions $\alpha_1 = \{2, 1, 1, 2, 2, 1\}$ and $\alpha_2 = \{3, 3, 3\}$ are both chordal (the corresponding block sparsity graphs are illustrated in Fig. 2.7). For a given chordal partition, in this example but also in general, completing all blocks of M that are not identically zero results in a chordal extension of M . For instance, the chordal extension of the 9×9 matrix above resulting from the partitions α_1 and α_2 are, respectively,

$$\begin{bmatrix} \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \end{bmatrix} \text{ and } \begin{bmatrix} \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \end{bmatrix},$$

where entries colored in red have been added by the block-completion process. Finding a chordal partition for a matrix, therefore, gives a way of performing a particular chordal extension of its sparsity pattern. The opposite, however, is not true: not all chordal extensions are obtained via a block-completion operation. One example for the 9×9 matrix above is the chordal extension

$$\begin{bmatrix} \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times & \times \end{bmatrix},$$

which is obtained by a symbolic Cholesky factorization with an approximate minimum degree ordering. ■

2.3.2. Chordal decomposition of sparse block matrices

As anticipated above, decomposition results similar to Theorems 2.1 and 2.2 hold for $\mathbb{S}_{\alpha,+}^n(\mathcal{E}, ?)$ and $\mathbb{S}_{\alpha,+}^n(\mathcal{E}, 0)$ when α is a chordal partition of n . Given a clique C_k of the chordal block sparsity graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$

subordinate to the chordal partition α , we define the block matrix $E_{C_k, \alpha} \in \mathbb{R}^{s(\alpha, k) \times n}$, where $s(\alpha, k) = \sum_{i \in C_k} \alpha_i$, as

$$(E_{C_k, \alpha})_{ij} = \begin{cases} I_{\alpha_i}, & \text{if } C_k(i) = j, \\ 0, & \text{otherwise.} \end{cases} \quad (2.14)$$

Here, I_{α_i} is an identity matrix of dimension α_i . When $\alpha = \{1, \dots, 1\}$ is the trivial partition, $E_{C_k, \alpha}$ reduces to the matrix E_{C_k} in (2.2). Similar to the case studied in Section 2.2, the operation $E_{C_k, \alpha} X E_{C_k, \alpha}^T \in \mathbb{S}^{s(\alpha, k)}$ extracts the principal block-submatrix of X whose blocks are indexed by C_k , while $E_{C_k, \alpha}^T Y E_{C_k, \alpha}$ “inflates” an $s(\alpha, k) \times s(\alpha, k)$ matrix into a sparse $n \times n$ block matrix.

We are now ready to extend Theorems 2.1 and 2.2 to the case of sparse block matrices.

Theorem 2.3 (Chordal Block-decomposition). Let $\mathcal{G}(\{1, \dots, p\}, \mathcal{E})$ be a chordal graph with maximal cliques C_1, \dots, C_t , and let $\alpha = \{\alpha_1, \dots, \alpha_p\}$ be a partition of n . Then, $Z \in \mathbb{S}_{\alpha,+}^n(\mathcal{E}, 0)$ if and only if there exist matrices $Z_k \in \mathbb{S}_{+}^{s(\alpha, k)}$ for $k = 1, \dots, t$ such that

$$Z = \sum_{k=1}^t E_{C_k, \alpha}^T Z_k E_{C_k, \alpha}. \quad (2.15)$$

Theorem 2.4 (Chordal Block-completion). Let $\mathcal{G}(\{1, \dots, p\}, \mathcal{E})$ be a chordal graph with maximal cliques C_1, \dots, C_t , and let $\alpha = \{\alpha_1, \dots, \alpha_p\}$ be a partition of n . Then, $X \in \mathbb{S}_{\alpha,+}^n(\mathcal{E}, ?)$ if and only if

$$E_{C_k, \alpha} X E_{C_k, \alpha}^T \in \mathbb{S}_{+}^{s(\alpha, k)} \quad \forall k = 1, \dots, t. \quad (2.16)$$

The proofs of Theorems 2.3 and 2.4 rely on the fact that the block sparsity graph of $X \in \mathbb{S}_{\alpha,+}^n(\mathcal{E}, 0)$ induces a chordal extension of the standard sparsity graph of X (cf. Remark 2.3) and, in fact, it is a hypergraph of the latter. The normal chordal decomposition and completion from Theorems 2.1 and 2.2 can then be applied to the chordal extension of X , and the hypergraph structure implies the two results above. Interested readers are referred to Zheng (2019, Chapter 2.4) for details.

Example 2.5. Consider the 9×9 matrices

$$X = \begin{bmatrix} 2 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 2 & 1 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 2 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 2 & 1 & 0 & 1 & 1 & 0 \\ 1 & 1 & 0 & 1 & 2 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 1 & 2 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 & 2 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 & 2 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 2 \end{bmatrix}$$

and

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

which have the same nonchordal sparsity pattern as the symbolic matrix considered in Remark 2.3. Readers can easily check that X is positive semidefinite, while Y admits a positive semidefinite completion (e.g., replace all zero entries with ones to obtain $\bar{Y} = \mathbf{1}\mathbf{1}^T \geq 0$). The partitions $\alpha_1 = \{2, 1, 1, 2, 2, 1\}$ and $\alpha_2 = \{3, 3, 3\}$ are both chordal, so while Theorem 2.1 cannot be directly applied to decompose X ,

Theorem 2.3 guarantees the existence of decompositions either in the symbolic form

$$X = \underbrace{\begin{bmatrix} \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \end{bmatrix}}_{\geq 0} + \underbrace{\begin{bmatrix} \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \end{bmatrix}}_{\geq 0} \\ + \underbrace{\begin{bmatrix} \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \end{bmatrix}}_{\geq 0} + \underbrace{\begin{bmatrix} \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \end{bmatrix}}_{\geq 0}$$

(corresponding to the partition α_1) or in the symbolic form

$$X = \begin{bmatrix} \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \end{bmatrix} + \begin{bmatrix} \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \end{bmatrix}$$

(corresponding to the partition α_2). These coincide with the classical chordal decompositions applied to the chordal extensions of X from the chordal partitions α_1 and α_2 . Thus, one can choose whether to decompose X as a sum of four matrices with 5×5 nonzero principal submatrices, or as a sum of two matrices with 6×6 nonzero principal submatrices. Similarly, one can apply [Theorem 2.4](#) to verify that the matrix Y admits a positive semidefinite completion by checking the positive semidefiniteness of either four 5×5 principal submatrices, or two 6×6 ones. ■

3. Sparse semidefinite optimization

The matrix decomposition and completion results in [Theorems 2.1](#) and [2.2](#) can be used to reduce the complexity of algorithms for sparse semidefinite optimization. A semidefinite program (SDP) in standard primal form takes the form

$$\begin{aligned} \min_x \quad & \langle C, X \rangle \\ \text{subject to} \quad & \langle A_i, X \rangle = b_i, \quad i = 1, \dots, m, \\ & X \in \mathbb{S}_+^n, \end{aligned} \quad (3.1)$$

where $C, A_1, \dots, A_m \in \mathbb{S}^n, b \in \mathbb{R}^m$ are the problem data. The dual problem to (3.1) is also an SDP,

$$\begin{aligned} \max_{y, Z} \quad & b^\top y \\ \text{subject to} \quad & Z + \sum_{i=1}^m y_i A_i = C, \\ & Z \in \mathbb{S}_+^n. \end{aligned} \quad (3.2)$$

In this section, we describe decomposition techniques for SDPs that exploit the joint sparsity pattern of the coefficient matrices C, A_1, \dots, A_m , called the *aggregate sparsity pattern*. For simplicity, we assume that the matrices A_1, \dots, A_m are linearly independent and that there exist $X \succ 0, y \in \mathbb{R}$ and $Z \succ 0$ satisfying the equality constraints in (3.1) and (3.2). This ensures that the primal and dual optimal values are finite, equal, and attained. SDPs that are infeasible or have unbounded objective can be tackled using homogeneous self-dual embeddings ([O'Donoghue, Chu, Parikh, & Boyd, 2016](#); [Permenter, Friberg, & Andersen, 2017](#); [Ye, 2011](#); [Ye, Todd, & Mizuno, 1994](#)) or by analyzing the divergence of the iterates produced by solution algorithms ([Banjac, Goulart, Stellato, & Boyd, 2019](#); [Liu, Ryu, & Yin, 2017](#)). Sparsity can be exploited within these frameworks, too, and we refer the interested reader to [Zheng, Fantuzzi, Papachristodoulou, Goulart, and Wynn \(2020, Section 5\)](#) and [Garstka, Cannon, and Goulart \(2021\)](#) for details.

3.1. Aggregate sparsity

The pair of SDPs (3.1)–(3.2) is said to have aggregate sparsity graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ if

$$C, A_1, \dots, A_m \in \mathbb{S}^n(\mathcal{E}, 0). \quad (3.3)$$

Of course, if \mathcal{E}' is an extension of \mathcal{E} , then $\mathcal{G}(\mathcal{V}, \mathcal{E}')$ is also a suitable aggregate sparsity graph. The minimal one, therefore, is simply the union of the individual sparsity graphs of C, A_1, \dots, A_m . Throughout this section, however, we consider a chordal extension of the minimal aggregate sparsity graph. We therefore assume from now on that the aggregate sparsity pattern \mathcal{E} is chordal and has t maximal cliques C_1, \dots, C_t .

It must be noted that an SDP may have a complete aggregate sparsity graph even if all coefficient matrices C and A_1, \dots, A_m are very sparse; [Zheng, Fantuzzi, and Papachristodoulou \(2018c\)](#) give explicit examples. The decomposition methods described below cannot be applied to such problems. However, there are broad classes of SDPs for which very sparse aggregate sparsity graphs can be expected, or even guaranteed.

One such family consists of SDPs arising from relaxations of graph optimization problems and control problems over networks, which typically inherit the structure of the underlying network or graph. Notable examples include SDP relaxations of combinatorial graph optimization problems, such as Max-Cut ([Goemans & Williamson, 1995](#)) and graph equipartition ([Karisch & Rendl, 1998](#)), eigenvalue optimization problems over graphs ([Boyd, Diaconis, & Xiao, 2004](#)), analysis of linear networked systems ([Derou, Meinel, Ulbrich, & Hirche, 2015](#); [Mason & Papachristodoulou, 2014](#); [Zheng, Kamgarpour et al., 2018](#); [Zheng, Mason et al., 2018](#)), sensor network localization ([Kim et al., 2009](#); [Nie, 2009](#); [So & Ye, 2007](#)), neural network verification in machine learning ([Batten et al., 2021](#); [Raghunathan, Steinhardt, & Liang, 2018](#)), and the optimal power flow problem in electricity networks ([Andersen, Hansson et al., 2014](#); [Bai, Wei, Fujisawa, & Wang, 2008](#); [Jabr, 2011](#)). We briefly discuss some of these applications in Section 6.

Another source of SDPs with aggregate sparsity is the reformulation of intractable constraints (either convex or nonconvex) as tractable LMIs using auxiliary variables ([Ben-Tal & Nemirovski, 2001](#); [Vandenberghe & Andersen, 2015](#)). For example, consider the uncountable family of “uncertain” convex quadratic constraints

$$x^\top A^\top A x - 2b^\top x - c \leq 0$$

on a variable $x \in \mathbb{R}^q$, to be imposed for all matrices $A \in \mathbb{R}^{p \times q}$, vectors $b \in \mathbb{R}^q$ and scalars $c \in \mathbb{R}$ in the form

$$A = A_0 + \sum_{i=1}^r u_i A_i, \quad b = b_0 + \sum_{i=1}^r u_i b_i, \quad c = c_0 + \sum_{i=1}^r u_i c_i$$

with $u^\top u \leq 1$. Here, A_0, b_0 and c_0 are nominal reference values, and $\{A_i, b_i, c_i\}$ are fixed perturbations. [Andersen, Vandenberghe and Dahl \(2010\)](#) showed that this family of constraints is equivalent to a sparse LMI in the form

$$\begin{bmatrix} f(x) - t & (A_0 x)^\top & h(x)^\top \\ A_0 x & I_p & G(x)^\top \\ h(x) & G(x) & t I_r \end{bmatrix} \geq 0, \quad (3.4)$$

where $t \in \mathbb{R}$ while $G : \mathbb{R}^q \rightarrow \mathbb{R}^{r \times p}$, $h : \mathbb{R}^q \rightarrow \mathbb{R}^r$ and $f : \mathbb{R}^q \rightarrow \mathbb{R}$ are known linear functions whose exact form is not important here. When $r \gg p$, this matrix has a “block-arrow” aggregate sparsity pattern analogous to that shown in [Fig. 2.3\(c\)](#) (that figure is recovered exactly when $p = 1, r = 6$, and q is arbitrary). This particular type of sparsity pattern is commonly encountered in robust optimization ([Andersen, Vandenberghe et al., 2010](#); [Ben-Tal & Nemirovski, 1998](#); [Goldfarb & Iyengar, 2003](#)).

Remark 3.1 (Promoting Aggregate Sparsity). Sometimes, it is possible to reformulate SDPs with no aggregate sparsity as equivalent SDPs with very sparse aggregate sparsity graphs through a carefully chosen transformation of variables (Fukuda et al., 2001, Section 6; Vandenberghe & Andersen, 2015, Chapter 14.1). For instance, the SDP relaxation of the graph equipartition problem studied by Fukuda et al. (2001, Section 6) has sparse data matrices C and A_1, \dots, A_{m-1} , but the m th constraint $\langle \mathbf{1}\mathbf{1}^\top, X \rangle = 0$ destroys the problem's aggregate sparsity because the matrix $A_m = \mathbf{1}\mathbf{1}^\top$ is dense. However, any matrix $X \in \mathbb{S}_+^n$ satisfying $\langle \mathbf{1}\mathbf{1}^\top, X \rangle = 0$ can be expressed as $X = VYV^\top$ for some matrix $Y \in \mathbb{S}_+^{n-1}$, where

$$V = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 & 0 \\ -1 & 1 & 0 & \dots & 0 & 0 \\ 0 & -1 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -1 & 1 \\ 0 & 0 & 0 & \dots & 0 & -1 \end{bmatrix} \in \mathbb{R}^{n \times (n-1)}.$$

Thus, the original SDP can be reformulated as

$$\begin{aligned} \min_Y \quad & \langle C', Y \rangle \\ \text{subject to} \quad & \langle A'_i, Y \rangle = b_i, \quad i = 1, \dots, m-1, \\ & Y \in \mathbb{S}_+^{n-1}, \end{aligned}$$

where $C' := V^\top C V$ and $A'_i := V^\top A_i V$. Since V is a sparse basis matrix and the original data matrices are sparse, this new SDP is characterized by aggregate sparsity (Fukuda et al., 2001, Section 6). Sparsity-promoting modeling strategies that generalize this example are discussed in Vandenberghe and Andersen (2015, Chapter 14.1). ■

3.2. Nonsymmetric formulation

The aggregate sparsity of the primal–dual pair of SDPs (3.1)–(3.2) can be exploited by reformulating them into a nonsymmetric pair of optimization problems, proposed by Fukuda et al. (2001) and later discussed extensively by Andersen, Dahland Vandenberghe (2010), Kim et al. (2011), Sun et al. (2014), Zheng et al. (2020).

Consider first the dual-standard-form SDP (3.2). Any feasible matrix Z must be at least as sparse as the aggregate sparsity pattern of the SDP. We can therefore restrict Z to the subspace $\mathbb{S}^n(\mathcal{E}, 0)$, where \mathcal{E} is the edge set of the aggregate sparsity graph, and rewrite (3.2) as

$$\begin{aligned} \max_{y, Z} \quad & \langle b, y \rangle \\ \text{subject to} \quad & Z + \sum_{i=1}^m A_i y_i = C, \\ & Z \in \mathbb{S}_+^n(\mathcal{E}, 0). \end{aligned} \quad (3.5)$$

The primal-standard-form SDP (3.1), instead, typically has a dense optimal matrix X . However, the value of the cost function and the equality constraints depend only on the entries X_{ij} with $(i, j) \in \mathcal{E}$, while the remaining ones simply guarantee that X is positive semidefinite. We can therefore pose (3.1) as an optimization problem over the cone $\mathbb{S}_+^n(\mathcal{E}, ?)$ of sparse matrices that admit a positive semidefinite completion,

$$\begin{aligned} \min_X \quad & \langle C, X \rangle \\ \text{subject to} \quad & \langle A_i, X \rangle = b_i, \quad i = 1, \dots, m, \\ & X \in \mathbb{S}_+^n(\mathcal{E}, ?). \end{aligned} \quad (3.6)$$

Problems (3.5) and (3.6) are a primal–dual pair of linear conic programs because the cones $\mathbb{S}_+^n(\mathcal{E}, ?)$ and $\mathbb{S}_+^n(\mathcal{E}, 0)$ are dual to each other (see Section 2.2 and Fig. 2.5). Even though the sparse matrix cones $\mathbb{S}_+^n(\mathcal{E}, ?)$ and $\mathbb{S}_+^n(\mathcal{E}, 0)$ are not self-dual (Andersen, 2011; Andersen, Dahl et al., 2010), so this sparse formulation is nonsymmetric, one can solve (3.6), (3.5), or both problems simultaneously using a variety of first-order or interior-point algorithms. The next two subsections discuss some of them.

Remark 3.2. A special type of aggregate sparsity arises when the data matrices C, A_1, \dots, A_m are block-diagonal. In this case, any feasible matrix X for (3.6) is automatically positive semidefinite and, consequently, can be restricted to $\mathbb{S}_+^n(\mathcal{E}, 0)$. Therefore, the nonsymmetric formulation described above becomes symmetric. In particular, problems (3.5) and (3.6) are simply SDPs with a Cartesian product $\mathbb{S}_+^{n_1} \times \mathbb{S}_+^{n_2} \times \dots \times \mathbb{S}_+^{n_l}$ of semidefinite cones, where n_i is the size of the i th diagonal block and l is the number of blocks. ■

3.3. First-order algorithms

First-order optimization algorithms rely only on gradient information and have iterations with low computational complexity, which can often be implemented in a distributed manner (Beck, 2017; Boyd, Parikh, Chu, Peleato, & Eckstein, 2011). For these reasons, the last decade has witnessed the development of a range of first-order methods to solve large-scale SDPs, many of which are listed in Table 1. Some of these methods (O'Donoghue et al., 2016; Wen et al., 2010; Yurtsever et al., 2021; Zhao et al., 2010) focus on generic SDPs and do not exploit aggregate sparsity. Others, instead, tackle the sparsity-exploiting nonsymmetric formulations (3.5)–(3.6) using so-called *domain space* or *range-space* conversion frameworks, which replace the matrix cones $\mathbb{S}_+^n(\mathcal{E}, ?)$ and $\mathbb{S}_+^n(\mathcal{E}, 0)$ with smaller positive semidefinite cones using the chordal decomposition and completion results in Theorems 2.1 and 2.2 (see, e.g., Dall'Anese et al., 2013; Garstka et al., 2021; Kalbat & Lavaei, 2015; Lam et al., 2012; Lu et al., 2007; Madani, Kalbat et al., 2017; Sun et al., 2014; Sun & Vandenberghe, 2015; Zheng, Fantuzzi, & Papachristodoulou, 2018b; Zheng et al., 2020). Many of these works combine this strategy with additional separability assumptions for the equality constraints, which are satisfied in optimal power flow problems (Dall'Anese et al., 2013; Eltvéd, Dahl, & Andersen, 2020; Kalbat & Lavaei, 2015; Lam et al., 2012) and the matrix nearness problems (Sun et al., 2014) but not in general. To the best of our knowledge, the only first-order methods that can currently handle general SDPs with aggregate sparsity (including infeasible or unbounded ones) are those developed by Zheng et al. (2020) and Garstka et al. (2021).

3.3.1. Domain- and range-space conversion

Consider problem (3.6). When the aggregate sparsity graph is chordal and has maximal cliques C_1, \dots, C_t , Theorem 2.2 allows one to replace the constraint $X \in \mathbb{S}_+^n(\mathcal{E}, ?)$ with

$$E_{C_k} X E_{C_k}^\top \in \mathbb{S}_+^{|C_k|} \quad \forall k = 1, \dots, t. \quad (3.7)$$

These constraints are coupled in general because the matrices $E_{C_p} X E_{C_p}^\top$ and $E_{C_q} X E_{C_q}^\top$ depend on the same entries of X if C_p intersects C_q . The works referenced above differ primarily in how these couplings are handled and, as discussed in Remarks 3.3 and 3.5 below, the choice of strategy can have a considerable impact on the overall complexity of the iterations in a first-order method.

A simple but powerful strategy was proposed recently by Zheng et al. (2020), who used “slack” matrices X_1, \dots, X_t to rewrite (3.7) as

$$\begin{cases} X_k = E_{C_k} X E_{C_k}^\top & \forall k = 1, \dots, t, \\ X_k \in \mathbb{S}_+^{|C_k|} & \forall k = 1, \dots, t. \end{cases} \quad (3.8)$$

The primal SDP (3.6) is then equivalent to

$$\begin{aligned} \min_{X, X_1, \dots, X_t} \quad & \langle C, X \rangle \\ \text{subject to} \quad & \langle A_i, X \rangle = b_i, \quad i = 1, \dots, m, \\ & X_k = E_{C_k} X E_{C_k}^\top, \quad k = 1, \dots, t, \\ & X_k \in \mathbb{S}_+^{|C_k|}, \quad k = 1, \dots, t. \end{aligned} \quad (3.9)$$

Following Fukuda et al. (2001) and Zheng et al. (2020), we refer to (3.9) as the *domain-space* decomposition of the primal SDP (3.1).

Table 1

Comparison of first-order algorithms for solving SDPs. “Chordal Sparsity”: whether the algorithm exploits chordal sparsity; “SDP Type”: the types of SDP problems the algorithm considers; “Algorithm”: the underlying first-order algorithm; “Infeas./unbounded”: whether the algorithm can detect infeasible or unbounded cases; “Solver”: whether the code is open-source.

Reference	Chordal Sparsity	SDP Type	Algorithm	Infeas./ Unbounded	Solver
Wen, Goldfarb, and Yin (2010)	✗	(3.2)	ADMM	✗	✗
Zhao, Sun, and Toh (2010)	✗	(3.2)	Augm. Lagrang.	✗	SDPNAL
O’Donoghue et al. (2016)	✗	(3.1)–(3.2)	ADMM	✓	SCS
Yurtsever, Tropp, Fercoq, Udell, and Cevher (2021)	✗	(3.1) ^a	SketchyCGAL	✗	CGAL
Lu, Nemirovski, and Monteiro (2007)	✓	(3.1)	Mirror-Prox	✗	✗
Lam, Zhang, and David (2012)	✓	OPF ^b	Primal–dual	✗	✗
Dall’Anese, Zhu, and Giannakis (2013)	✓	OPF ^b	ADMM	✗	✗
Sun, Andersen, and Vandenberghe (2014)	✓	(3.1)–(3.2)	Spingarn	✗	✗
Sun and Vandenberghe (2015)	✓	Special ^c	Gradient proj.	✗	✗
Kalbat and Lavaei (2015)	✓	Special ^d	ADMM	✗	✗
Madani, Kalbat and Lavaei (2017)	✓	General ^e	ADMM	✗	✗
Zheng et al. (2020)	✓	(3.1)–(3.2)	ADMM	✓	CDCS
Garstka et al. (2021)	✓	Quad. SDP ^f	ADMM	✓	COSMO

^aSpecial SDPs with an explicit trace constraint on X .

^bSpecial SDPs from the optimal power flow (OPF) problem.

^cSpecial SDPs from the matrix nearness problem.

^dSpecial SDPs with decoupled affine constraints.

^eGeneral SDPs with inequality constraints.

^fA dual SDP (3.2) with a quadratic objective function.




Fig. 3.1. Duality between the original primal and dual SDPs, and the decomposed primal and dual SDPs.

A *range-space* decomposition of the dual SDP (3.2) can be formulated in a very similar way. When the aggregate sparsity pattern \mathcal{E} is chordal, Theorem 2.1 implies that the constraint $Z \in \mathbb{S}_+^n(\mathcal{E}, 0)$ is equivalent to

$$\begin{cases} Z = \sum_{k=1}^t E_{C_k}^\top Z_k E_{C_k}, \\ Z_k \in \mathbb{S}_+^{|C_k|} \quad \forall k = 1, \dots, t. \end{cases} \quad (3.10)$$

Observe that, as before, the first of these conditions couples the positive semidefinite matrices Z_p and Z_q if the cliques C_p and C_q of the aggregate sparsity graph overlap. To decouple them, Zheng et al. (2020) introduced slack variables V_1, \dots, V_t and reformulated (3.10) as

$$\begin{cases} Z = \sum_{k=1}^t E_{C_k}^\top V_k E_{C_k}, \\ V_k = Z_k \quad \forall k = 1, \dots, t, \\ Z_k \in \mathbb{S}_+^{|C_k|} \quad \forall k = 1, \dots, t. \end{cases} \quad (3.11)$$

Using this to eliminate Z from (3.5) yields the range-space decomposition

$$\begin{aligned} & \max_{y, Z_1, \dots, Z_t, V_1, \dots, V_t} \langle b, y \rangle \\ & \text{subject to} \quad \sum_{i=1}^m A_i y_i + \sum_{k=1}^t E_{C_k}^\top V_k E_{C_k} = C, \\ & \quad Z_k - V_k = 0, \quad k = 1, \dots, t, \\ & \quad Z_k \in \mathbb{S}_+^{|C_k|}, \quad k = 1, \dots, t. \end{aligned} \quad (3.12)$$

While the domain- and range-space decompositions (3.9) and (3.12) have been derived independently, it is not difficult to verify that they are a primal–dual pair of SDPs. The duality between the original SDPs (3.1) and (3.2) is thus inherited by the decomposed SDPs (3.9) and (3.12) by virtue of the duality between Theorems 2.1 and 2.2. This elegant picture is illustrated in Fig. 3.1.

Remark 3.3. The introduction of variables X_k and V_k leads to redundancies in the affine constraints of (3.9) and (3.12), but is essential to obtain a decomposition framework that is suitable for the development of fast first-order SDP solvers. For example, as explained in Section 3.3.2 below, applying the alternating direction method of multipliers (ADMM) to (3.9) leads to an algorithm whose iterations have closed-form update rules that can be implemented efficiently. The same is usually not true if the redundant constraints in (3.9) are used to eliminate the matrix X : the iterations of the first-order method proposed by Sun et al. (2014), for instance, require the solution of a further SDP with quadratic objective function, which limits its scalability. However, the matrix X may be eliminated from (3.9) without compromising efficiency if the original primal SDP (3.1) has separable affine constraints. This observation was exploited to solve sparse SDPs arising from optimal power flow problems (Dall’Anese et al., 2013; Eltved et al., 2020; Kalbat & Lavaei, 2015) and matrix nearness problems (Sun & Vandenberghe, 2015). Similar observations apply to the seemingly redundant matrices V_k in the range-space decomposed SDP (3.12). ■

3.3.2. ADMM for decomposed SDPs

The alternating direction method of multipliers (ADMM) is a first-order operator-splitting method developed in the mid-1970s (Gabay & Mercier, 1976; Glowinski & Marroco, 1975) to solve general optimization problems in the form

$$\begin{aligned} & \min_{\mathcal{X} \in \mathbb{X}} f(\mathcal{X}) + g(\mathcal{Y}) \\ & \text{subject to} \quad \mathcal{A}(\mathcal{X}) + \mathcal{B}(\mathcal{Y}) = C, \end{aligned} \quad (3.13)$$

where f and g are proper convex (but not necessarily smooth) functions on finite-dimensional normed vector spaces \mathbb{X} and \mathbb{Y} , \mathcal{A} and \mathcal{B} are given linear operators from \mathbb{X} and \mathbb{Y} into a finite-dimensional normed vector space \mathbb{Z} , and $C \in \mathbb{Z}$ is given. Given a penalty parameter $\rho > 0$ and a dual variable $\mathcal{Z} \in \mathbb{Z}$ that acts as a Lagrange multiplier for

the equality constraint, ADMM finds a saddle point of the (scaled) augmented Lagrangian

$$\mathcal{L}_\rho(\mathcal{X}, \mathcal{Y}, \mathcal{Z}) := f(\mathcal{X}) + g(\mathcal{Y}) + \frac{\rho}{2} \|\mathcal{A}(\mathcal{X}) + \mathcal{B}(\mathcal{Y}) - C + \mathcal{Z}\|^2$$

by updating the primal variables \mathcal{X} , \mathcal{Y} and the dual variable \mathcal{Z} according to the following rules:

$$\mathcal{X}^{(q+1)} = \operatorname{argmin}_{\mathcal{X}} \mathcal{L}_\rho(\mathcal{X}, \mathcal{Y}^{(q)}, \mathcal{Z}^{(q)}), \quad (3.14a)$$

$$\mathcal{Y}^{(q+1)} = \operatorname{argmin}_{\mathcal{Y}} \mathcal{L}_\rho(\mathcal{X}^{(q+1)}, \mathcal{Y}, \mathcal{Z}^{(q)}), \quad (3.14b)$$

$$\mathcal{Z}^{(q+1)} = \mathcal{Z}^{(n)} + \mathcal{A}(\mathcal{X}^{(q+1)}) + \mathcal{B}(\mathcal{Y}^{(q+1)}) - C. \quad (3.14c)$$

The superscript (q) indicates that a variable is fixed to its value at the q th iteration. Under mild technical conditions (Boyd et al., 2011, Section 3.2), the method converges to an ϵ -approximate solution of (3.13) using at most $\mathcal{O}(1/\epsilon)$ iterations.

Given its slow convergence rate, ADMM is suitable only when (3.14a) and (3.14b) have closed-form expressions and/or can be solved efficiently. Below, we show that this is true when the method is applied to the decomposed SDPs (3.9) and (3.12).

Domain-space decomposition. Consider the domain-space decomposition (3.9). Let $\chi_{\mathcal{K}}(x)$ denote the characteristic function of a set \mathcal{K} , i.e.,

$$\chi_{\mathcal{K}}(x) := \begin{cases} 0 & \text{if } x \in \mathcal{K}, \\ +\infty & \text{otherwise.} \end{cases}$$

For simplicity, we write χ_0 when $\mathcal{K} \equiv \{0\}$. Problem (3.9) is equivalent to

$$\min_{X, X_1, \dots, X_t} \langle C, X \rangle + \sum_{i=1}^m \left[\chi_0(\langle A_i, X \rangle - b_i) + \chi_{\mathbb{S}_+^{|C_k|}}(X_k) \right]$$

subject to $X_k = E_{C_k} X E_{C_k}^\top$, $k = 1, \dots, t$.

Upon letting $\mathcal{X} := \{X\}$ and $\mathcal{Y} := \{X_1, \dots, X_t\}$, this problem may be written in the standard form (3.13) over the spaces $\mathbb{X} = \mathbb{S}^n$ and $\mathbb{Y} = \mathbb{Z} = \mathbb{S}^{|C_1|} \times \dots \times \mathbb{S}^{|C_t|}$, and can therefore be solved using ADMM. Introducing a penalty parameter $\rho > 0$ and a dual variable $\mathcal{Z} := \{\Lambda_1, \dots, \Lambda_t\}$, where each $\Lambda_k \in \mathbb{S}^{|C_k|}$ acts as a Lagrange multiplier for the corresponding constraint $X_k = E_{C_k} X E_{C_k}^\top$, it is not difficult to check that the ADMM step (3.14a) reduces to an equality-constrained quadratic program,

$$X^{(q+1)} = \operatorname{argmin}_{\substack{\langle A_i, X \rangle = b_i \\ i=1, \dots, m}} \left\{ \frac{\rho}{2} \sum_{k=1}^t \|X_k^{(q)} - E_{C_k} X E_{C_k}^\top + \Lambda_k^{(q)}\|_F^2 + \langle C, X \rangle \right\}.$$

Step (3.14b), instead, reduces to t independent positive semidefinite projections of the form

$$X_k^{(q+1)} = \operatorname{argmin}_{X_k \in \mathbb{S}_+^{|C_k|}} \|X_k - E_{C_k} X^{(q+1)} E_{C_k}^\top + \Lambda_k^{(q)}\|_F^2.$$

Finally, step (3.14c) updates the multipliers $\Lambda_1, \dots, \Lambda_t$ according to

$$\Lambda_k^{(q+1)} = \Lambda_k^{(q)} + X_k^{(q+1)} - E_{C_k} X^{(q+1)} E_{C_k}^\top.$$

These three steps have efficient closed-form solutions and can be implemented efficiently (Zheng et al., 2020, Section 4.1). In particular, the t independent projections onto the cones $\mathbb{S}_+^{|C_k|}$ required to compute $X_1^{(q+1)}, \dots, X_t^{(q+1)}$ can be performed through an eigenvalue decomposition with complexity of $\mathcal{O}(|C_k|^3)$ floating-point operations. This is not expensive when all cliques C_1, \dots, C_t of the aggregate sparsity graph \mathcal{E} are small, which is often true in many applications. In contrast, the first-order algorithms for generic SDPs developed by Wen et al. (2010) and by O'Donoghue et al. (2016) without chordal decomposition require a projection onto the semidefinite cone \mathbb{S}_+^n at each iteration, which becomes a bottleneck when $n \gg 1$. It is therefore clear that exploiting sparsity via chordal decomposition can bring significant

computational savings in ADMM algorithms. In principle, the algorithm by O'Donoghue et al. (2016) can be used to solve the homogeneous self-dual embedding of (3.9) and (3.12), but to maximize efficiency the underlying problem structure should be exploited when solving a large linear system of equations at each iteration. This can be done via a sequence of block eliminations, as explained by Zheng et al. (2020, Section 5.2).

Range-space decomposition. The range-space decomposition (3.12) of the dual-standard-form SDP (3.2) can be solved using an ADMM algorithm very similar to that presented above. First, observe that (3.12) is equivalent to

$$\min_{y, V_k, Z_k} -\langle b, y \rangle + \chi_0 \left(C - \sum_{i=1}^m A_i y_i - \sum_{k=1}^t E_{C_k}^\top V_k E_{C_k} \right) + \sum_{k=1}^t \chi_{\mathbb{S}_+^{|C_k|}}(Z_k)$$

subject to $Z_k - V_k = 0$, $k = 1, \dots, t$.

Grouping the variables as $\mathcal{X} := \{y, V_1, \dots, V_t\}$ and $\mathcal{Y} := \{Z_1, \dots, Z_t\}$, this problem can be written in the general form (3.13) over the spaces $\mathbb{X} = \mathbb{R}^m \times \mathbb{S}^{|C_1|} \times \dots \times \mathbb{S}^{|C_t|}$ and $\mathbb{Y} = \mathbb{Z} = \mathbb{S}^{|C_1|} \times \dots \times \mathbb{S}^{|C_t|}$. Given a penalty parameter $\rho > 0$ and a dual variable $\mathcal{Z} := \{\Lambda_1, \dots, \Lambda_t\}$, where each Λ_k acts as a Lagrange multiplier for the corresponding constraint $Z_k - V_k = 0$, one can easily verify that the ADMM step (3.14a) reduces to solving the equality-constrained quadratic program

$$\min_{y, V_1, \dots, V_t} -\langle b, y \rangle + \frac{\rho}{2} \sum_{k=1}^t \|Z_k^{(q)} - V_k + \Lambda_k^{(q)}\|_F^2$$

subject to $C - \sum_{i=1}^m A_i y_i - \sum_{k=1}^t E_{C_k}^\top V_k E_{C_k} = 0$.

Step (3.14b), instead, reduces to t independent positive semidefinite projections of the form

$$Z_k^{(q+1)} = \operatorname{argmin}_{Z_k \in \mathbb{S}_+^{|C_k|}} \|Z_k - V_k^{(q+1)} + \Lambda_k^{(q)}\|_F^2.$$

Finally, the dual variables $\Lambda_1, \dots, \Lambda_t$ are updated through step (3.14c) as

$$\Lambda_k^{(q+1)} = \Lambda_k^{(q)} + Z_k^{(q+1)} - V_k^{(q+1)}.$$

Again, these iterations admit inexpensive closed-loop expressions. Moreover, it is not difficult to see that the ADMM iterations for the range-space decomposition (3.12) and for the domain-space decomposition (3.9) have similar leading-order complexity. In fact, Zheng et al. (2020, Section 4.3) showed that the ADMM algorithms for the primal and dual decomposed SDPs are scaled versions of each other. The duality picture of Figs. 2.5 and 3.1 is therefore reflected also at the algorithmic level.

Remark 3.4. For any fixed penalty parameter $\rho > 0$, the primal and dual ADMM algorithms outlined above converge to a solution of (3.9) and (3.12), respectively, provided that strict primal–dual feasibility conditions are satisfied (Boyd et al., 2011, Section 3.2). An efficient ADMM algorithm that can handle primal or dual infeasible problems was developed by Zheng et al. (2020, Section 5), who considered the homogeneous self-dual embedding (O'Donoghue et al., 2016; Ye et al., 1994) of the domain-space decomposition (3.9) and the range-space decomposition (3.12). ■

Remark 3.5. As anticipated in Remark 3.3, considering the variables X_k and the constraints $X_k = E_{C_k} X E_{C_k}^\top$ without eliminating any redundant variables is essential to obtain efficient ADMM iterations. This is because the conic constraints separate completely from the affine ones in (3.9) when applying the splitting strategy of ADMM, making it easy to update each X_k via simple projections onto positive semidefinite cones. Similarly, the redundant variables V_k and the constraints $Z_k = V_k$ in (3.12) are essential to decouple the conic constraints from the affine ones, which enables one to handle positive semidefinite constraints via simple projections. ■

3.4. Interior-point algorithms

Interior-point algorithms for convex optimization problems with equality and inequality constraints employ Newton's method to solve a sequence of modified equality-constrained problems, obtained by replacing any inequality constraints with barrier functions in the objective (Nesterov, 2003; Ye, 2011). These barrier functions approximate the characteristic function of the set defined by the original inequality constraints and ensure that the optimal solution of each modified problem is strictly feasible for the original problem, meaning that it is an interior point of the original feasible set.

Since Newton's method relies on second-order (Hessian) information, interior-point algorithms do not share the slow convergence of first-order methods. Instead, they converge to an ϵ -approximate solution using at most $\mathcal{O}(\log(1/\epsilon))$ Newton iterations (Nesterov, 2003; Ye, 2011). In practice, convergence often occurs within tens of iterations. Therefore, interior-point methods are typically preferred when solving (3.1)–(3.2) to high accuracy. The general-purpose SDP solvers SeDuMi (Sturm, 1999), SDPT3 (Tütüncü, Toh, & Todd, 2003), SDPA (Yamashita, Fujisawa, Fukuda, Kobayashi, Nakata, & Nakata, 2012), and MOSEK (Mosek, 2015) are all based on primal–dual interior-point methods, and they can very reliably solve small and medium-sized SDPs (e.g., when n is less than a few hundreds and m is less than a few thousands in (3.1)–(3.2)) on regular computers. However, they become impractical for large SDPs because the CPU time and memory requirements for each interior-point iteration increase as $\mathcal{O}(n^3m + n^2m^2 + m^3)$ and $\mathcal{O}(n^2 + m^2)$, respectively (Nesterov, 2003, Section 4.3.3).

Chordal graph techniques can be exploited to improve the efficiency of interior-point methods when solving large-scale SDPs with chordal aggregate sparsity (Andersen, 2011; De Klerk, 2010; Fukuda et al., 2001). This section reviews two general approaches for doing so. The first one, similar to the conversion methods in Section 3.3.1, reformulates problems (3.5) and (3.6) as SDPs with small positive semidefinite cones, which are often easier to solve with general-purpose interior-point solvers (Fukuda et al., 2001; Kim et al., 2011; Nakata et al., 2003; Zhang & Lavaei, 2020b). The second approach, instead, directly solves (3.6)–(3.5) using an interior-point method for nonsymmetric conic optimization (Andersen, Dahl et al., 2010; Coey, Kapelevich, & Vielma, 2020; Nesterov, 2012; Skajaa & Ye, 2015). For other ways to exploit chordal sparsity in the computation of interior-point search directions, we refer the reader to the works by (Benson, Ye, & Zhang, 2000), Pakazad, Hansson, Andersen and Rantzer (2017) and Fukuda et al. (2001, Section 5).

3.4.1. Conversion methods

Starting from the domain-space decomposed SDP (3.9), Fukuda et al. (2001) and Kim et al. (2011) suggested to eliminate the global matrix X and rewrite the SDP (3.6) only in terms of variables $X_k \in \mathbb{S}_+^{|C_k|}$, $k = 1, \dots, t$. To rewrite the cost function and the first set of equality constraints, one must choose matrices C_k and A_{ik} that satisfy

$$\sum_{k=1}^t \langle C_k, X_k \rangle = \langle C, X \rangle$$

and

$$\sum_{k=1}^t \langle A_{ik}, X_k \rangle = \langle A_i, X \rangle, \quad i = 1, \dots, m.$$

These affine relations do not usually determine C_k and A_{ik} uniquely, and some choices may be more convenient than others from the point of view of computations (Sun et al., 2014, Section 3.1, Zhang & Lavaei, 2020b, Section 6). The second set of constraints in (3.9), instead, can be enforced via consistency constraints on the entries of X_1, \dots, X_t that correspond to the same elements of X . Such consistency constraints can be formulated as

$$E_{C_j \cap C_k} \left(E_{C_k}^T X_k E_{C_k} - E_{C_j}^T X_j E_{C_j} \right) E_{C_j \cap C_k}^T = 0 \quad \forall j, k : C_j \cap C_k \neq \emptyset. \quad (3.18)$$

The primal SDP (3.6) can therefore be rewritten as

$$\begin{aligned} \min_{X_1, \dots, X_t} \quad & \sum_{k=1}^t \langle C_k, X_k \rangle \\ \text{subject to} \quad & \sum_{k=1}^t \langle A_{ik}, X_k \rangle = b_i, \quad i = 1, \dots, m, \\ & (3.18), \quad X_k \in \mathbb{S}_+^{|C_k|} \quad \forall k = 1, \dots, t. \end{aligned} \quad (3.19)$$

This conversion process, first proposed in Fukuda et al. (2001), is known as the *domain-space decomposition* (Kim et al., 2011). The reformulated problem (3.19) has more constraints than the original SDP (3.1), but the large matrix constraint $X \in \mathbb{S}_+^n$ is replaced by t smaller ones, $X_k \in \mathbb{S}_+^{|C_k|}$ for $k = 1, \dots, t$. In certain cases, the decomposed problem (3.19) is easier to solve than the original SDP (3.1) using general-purpose interior-point solvers; see Nakata et al. (2003) and Fujisawa, Kim, Kojima, Okamoto, and Yamashita (2009) for numerical examples. However, the decomposed problem can be degenerate if the original SDP has a low-rank optimal solution (Raghunathan & Knyazev, 2016, Theorem 1), which may complicate its numerical solution. Three other variants of this conversion method, including *range-space decompositions*, have been studied by Kim et al. (2011).

The main drawback of these conversion methods is that, sometimes, the additional consistency constraints (3.18) significantly increase the size of the Schur complement system that needs to be solved at each interior-point iteration. Even though this system is typically sparse, the increase in its size can offset the benefits of the clique-based matrix decomposition. As shown recently by Zhang and Lavaei (2020b), this issue can be mitigated using a dualization technique (Löfberg, 2009a).

Remark 3.6 (Removing Redundant Constraints). Since the maximal cliques in a chordal graph satisfy the running intersection property (Blair & Peyton, 1993; Fukuda et al., 2001) (see also Appendix C), it is in fact sufficient to enforce the consistency between pairs C_j, C_k that correspond to the parent–child pairs in a clique tree. Redundant constraints in (3.18) can therefore be removed using the running intersection property. Interested readers are referred to Kim et al. (2011) and Vandenberghe and Andersen (2015) for details. ■

Remark 3.7 (Dropping or Fixing Consistency Constraints). In some applications, the SDP (3.1) comes from a semidefinite relaxation of a non-convex optimization problem. Dropping some consistency constraints in (3.18) leads to a valid weaker relaxation with a lower computational complexity. This idea was successfully applied to semidefinite relaxations for optimal power flow problems (Andersen, Hansson et al., 2014) and neural network verification (Batten et al., 2021). Other times, one can enforce some of the consistency conditions *a priori* and look for feasible but suboptimal points for an SDP at a low computational cost. This idea was used by Zheng, Mason et al. (2018) to develop a scalable approach for solving distributed control problems. ■

3.4.2. Nonsymmetric interior-point algorithms

Chordal graph techniques can be exploited to speed up interior-point methods for the nonsymmetric pair of sparse SDPs (3.6)–(3.5) without appealing to the matrix decomposition and conversion frameworks described above. Since the cones $\mathbb{S}_+^n(\mathcal{E}, ?)$ and $\mathbb{S}_+^n(\mathcal{E}, 0)$ are not self-dual, such sparsity-exploiting methods cannot enjoy a complete primal–dual symmetry (Andersen, Dahl et al., 2010). Instead, one must resort to purely primal, purely dual, or nonsymmetric primal–dual path-following methods (Andersen, Dahl et al., 2010; Burer, 2003; Coey et al., 2020; Nesterov, 2012; Skajaa & Ye, 2015).

To construct nonsymmetric interior-point methods, Dahl et al. (2008) and Andersen, Dahl et al. (2010) introduced barrier functions $\phi : \mathbb{S}_+^n(\mathcal{E}, 0) \rightarrow \mathbb{R}$ and $\phi_* : \mathbb{S}_+^n(\mathcal{E}, 0) \rightarrow \mathbb{R}$ for the cones $\mathbb{S}_+^n(\mathcal{E}, 0)$ and $\mathbb{S}_+^n(\mathcal{E}, ?)$, defined as

$$\phi(Z) = \begin{cases} -\log \det Z & Z \in \text{int}(\mathbb{S}_+^n(\mathcal{E}, 0)), \\ +\infty & \text{otherwise,} \end{cases} \quad (3.20a)$$

and

$$\phi_*(X) = \sup_{Z \in \mathbb{S}^n(\mathcal{E}, 0)} (-\langle X, Z \rangle - \phi(Z)). \quad (3.20b)$$

Note that ϕ (resp. ϕ_*) is finite only on the interior of $\mathbb{S}_+^n(\mathcal{E}, 0)$ (resp. $\mathbb{S}_+^n(\mathcal{E}, ?)$) and tends to $+\infty$ as Z (resp. X) approaches the boundary of this cone. Observe also that ϕ_* is simply the Legendre transform of ϕ evaluated at $-X$.

Thanks to the properties of the barrier functions, a minimizing sequence $\{X^\mu\}_{\mu>0}$ for (3.6) can be computed by solving the regularized primal problem

$$\begin{aligned} \min_X \quad & \langle C, X \rangle + \mu\phi_*(X) \\ \text{subject to} \quad & \langle A_i, X \rangle = b_i, \quad i = 1, \dots, m, \end{aligned} \quad (3.21)$$

and letting $\mu \rightarrow 0$. Similarly, a minimizing sequence $\{y^\mu, Z^\mu\}_{\mu>0}$ for (3.5) is found upon solving the regularized dual problem

$$\begin{aligned} \max_{y, Z} \quad & \langle b, y \rangle - \mu\phi(Z) \\ \text{subject to} \quad & Z + \sum_{i=1}^m A_i y_i = C \end{aligned} \quad (3.22)$$

for $\mu \rightarrow 0$. Solutions of the regularized problems for fixed finite μ are usually found using Newton's method, leading to so-called *primal scaling* and *dual scaling* interior-point methods. Other methods can also be used; for instance, Jiang and Vandenberghe (2021) recently suggested solving (3.21) with a Bregman first-order method, where the complexity of evaluating the Bregman proximal operator can be reduced using a sparse Cholesky factorization.

When Newton's method is applied to (3.21), the KKT optimality conditions are

$$\langle A_i, X^\mu \rangle = b_i, \quad i = 1, \dots, m, \quad (3.23a)$$

$$\sum_{i=1}^m y_i A_i + Z = C, \quad (3.23b)$$

$$\mu \nabla \phi_*(X^\mu) + Z = 0, \quad (3.23c)$$

where $y \in \mathbb{R}^m$ is a Lagrange multiplier for the equality constraint in (3.21) and Z is an auxiliary variable arising from the definition of ϕ_* via the Legendre transform. Solutions $X^\mu \in \mathbb{S}_+^n(\mathcal{E}, ?)$ as μ is varied define the so-called *central path* for (3.6). Similarly, the KKT optimality conditions for (3.22) are

$$\langle A_i, X \rangle = b_i, \quad i = 1, \dots, m, \quad (3.24a)$$

$$\sum_{i=1}^m y_i^\mu A_i + Z^\mu = C, \quad (3.24b)$$

$$\mu \nabla \phi(Z^\mu) + X = 0, \quad (3.24c)$$

where X is a Lagrange multiplier for the equality constraint in (3.22). Solutions $\{y^\mu, Z^\mu\} \in \mathbb{R}^m \times \mathbb{S}_+^n(\mathcal{E}, 0)$ as μ is varied define the central path for (3.5). It is possible to show that (3.23) and (3.24) are equivalent (Andersen, 2011, Chapter 3), so the set of points $\{X^\mu, y^\mu, Z^\mu\}_{\mu>0}$ in $\mathbb{S}_+^n(\mathcal{E}, ?) \times \mathbb{R}^m \times \mathbb{S}_+^n(\mathcal{E}, 0)$ define a primal–dual central path.

The rest of this section briefly outlines how the chordality of the sparsity pattern \mathcal{E} can be exploited in the context of dual-scaling interior point methods. Similar ideas can be used to formulate primal-scaling methods, and we refer interested readers to the work by Andersen, Dahl et al. (2010, Section 4.2) for details.

Dual-scaling interior-point methods. Search directions in a dual-scaling interior-point method are obtained by linearizing (3.24) around the current interior iterate $X \in \text{int}(\mathbb{S}_+^n(\mathcal{E}, ?))$, $y \in \mathbb{R}^m$ and $Z \in \text{int}(\mathbb{S}_+^n(\mathcal{E}, 0))$. Replacing X , y and Z with $X + \Delta X$, $y + \Delta y$, $Z + \Delta Z$ in (3.24), linearizing (3.24c), and eliminating ΔZ yields the Newton equations

$$\begin{aligned} \langle A_i, \Delta X \rangle &= r_i, \quad i = 1, \dots, m, \\ \sum_{i=1}^m \Delta y_i A_i - \frac{1}{\mu} \nabla^2 \phi(Z)^{-1} [\Delta X] &= R, \end{aligned} \quad (3.25)$$

where $\nabla^2 \phi(Z)^{-1}$ is the inverse Hessian of ϕ at Z , $r_i = b_i - \langle A_i, X \rangle$ and $R = C - \sum_{i=1}^m y_i A_i - 2Z + \frac{1}{\mu} \nabla^2 \phi(Z)^{-1} [X]$. Further elimination of ΔX leads to the Schur complement equation

$$H \Delta y = g, \quad (3.26)$$

where $g \in \mathbb{R}^m$ is a vector and H is an $m \times m$ positive definite matrix, both depending only on the current (known) iterates X , y and Z . Explicit expression for these quantities are given by Andersen, Dahl et al. (2010, Section 4.3).

Finding the dual-scaling search direction $\Delta X, \Delta y, \Delta Z$ requires solving the Newton equation (3.25) or the Schur complement equation (3.26). To do this using a direct method, one must first calculate the Hessian of the barrier function $\phi(Z)$ in (3.20a), and then form and factorize the matrix H . This is typically the most computationally expensive part of any interior-point method. One must also be able to apply the inverse Hessian to form the right-hand side g of (3.26). It is in these computations that one can exploit the chordality of the sparsity pattern \mathcal{E} (Andersen, Dahl et al., 2010).

Fast calculations involving the barrier functions. The value, gradient, Hessian, and inverse Hessian of the dual barrier $\phi(Z)$ in (3.20a) can be computed efficiently if the sparsity pattern \mathcal{E} of Z is chordal. Similar fast algorithms exist for the primal barrier $\phi_*(X)$ in (3.20b), but we do not review them here and refer interested readers to Andersen, Dahl et al. (2010, Section 3.2) for details.

The key ingredient of these efficient algorithms is a sparse Cholesky factorization with zero fill-in (Blair & Peyton, 1993; Rose, 1970; Vandenberghe & Andersen, 2015): as reviewed in Appendix A, for any positive definite matrix Z in $\text{int}(\mathbb{S}_+^n(\mathcal{E}, 0))$ with chordal sparsity there exists a permutation matrix P and a lower triangular matrix L such that

$$P Z P^\top = L L^\top, \quad P^\top (L + L^\top) P \in \mathbb{S}^n(\mathcal{E}, 0). \quad (3.27)$$

This factorization can be computed efficiently by following a recursion on a clique tree (Vandenberghe & Andersen, 2015, Chapter 9.3).

Now, to evaluate $\phi(Z)$ it suffices to substitute $Z = P^\top L L^\top P$ into (3.20a) and observe that

$$\phi(Z) = -2 \sum_{i=1}^n \log L_{ii}$$

because determinants distribute over products and permutation matrices have unit determinant. Thus, $\phi(Z)$ can be evaluated efficiently once the Cholesky factorization (3.27) has been computed.

The gradient of $\phi(Z)$, instead, is given by the following negative projected inverse

$$\nabla \phi(Z) = -\mathbb{P}_{\mathbb{S}^n(\mathcal{E}, 0)}(Z^{-1}).$$

Despite the fact that Z^{-1} is in general dense, the projection onto $\mathbb{S}^n(\mathcal{E}, 0)$ can be computed from its sparse Cholesky factorization (3.27) without computing any other entries of Z^{-1} (Vandenberghe & Andersen, 2015, Chapter 9.5).

The Hessian of ϕ at Z applied to a matrix $Y \in \mathbb{S}^n(\mathcal{E}, 0)$ is computed as

$$\nabla^2 \phi(Z)[Y] = \frac{d}{dt} \nabla \phi(Z + tY) |_{t=0} = \mathbb{P}_{\mathbb{S}^n(\mathcal{E}, 0)}(Z^{-1} Y Z^{-1}).$$

Again, this quantity can be evaluated knowing only the sparse Cholesky factorization of Z and its projected inverse $\mathbb{P}_{\mathbb{S}^n(\mathcal{E}, 0)}(Z^{-1})$, without having to compute explicitly the inverse Z^{-1} or the matrix product $Z^{-1} Y Z^{-1}$ (Andersen, Dahl et al., 2010; Andersen, Dahl, & Vandenberghe, 2013).

Finally, thanks to the chordal structure, solving the linear equation $\nabla^2 \phi(Z)[U] = Y$ for U in order to apply the inverse Hessian to Y has the same cost as calculating $\nabla^2 \phi(Z)[Y]$; see Andersen, Dahl et al. (2010, Section 3.2) and Andersen et al. (2013).

3.5. Algorithm implementations

We conclude this section by providing a list of numerical packages that implement some of the approaches reviewed above. This list is not exhaustive, and the goal here is to give the interested reader a starting point for numerical experiments. First-order solvers based on augmented Lagrangian methods and ADMM for generic SDPs include SDPNAL/SDPNAL+ (Sun, Toh, Yuan, & Zhao, 2020; Zhao et al., 2010) and SCS (O'Donoghue, Chu, Parikh, & Boyd, 2019). CDCS (Zheng, Fantuzzi, Papachristodoulou, Goulart, & Wynn, 2016) and COSMO (Garstka et al., 2021) are two open-source first-order solvers that exploit chordal sparsity in SDPs. The MATLAB package CDCS implements the algorithms described in Section 3.3.2 and has interfaces with the optimization toolboxes YALMIP (Löfberg, 2004) and SOSTOOLS (Prajna, Papachristodoulou, & Parrilo, 2002). The Julia package COSMO solves SDPs with quadratic objective functions.

The conversion methods in Section 3.4.1 are implemented in SparseCoLO (Fujisawa et al., 2009) and CHOMPAC (Andersen & Vandenberghe, 2015). We note that CHOMPAC also provides useful implementation of many other chordal matrix computations, including maximum determinant positive definite completion and minimum rank positive semidefinite completion. Another MATLAB package Dual-CTC (Zhang & Lavaei, 2020a) implements a dualized clique tree conversion (Zhang & Lavaei, 2020b). The reformulated SDPs after conversion can be solved using general-purpose interior-point solvers, such as SeDuMi (Sturm, 1999), SDPT3 (Tütüncü et al., 2003), SDPA (Yamashita et al., 2012), and MOSEK (Mosek, 2015). SMCP (Andersen & Vandenberghe, 2014) is a nonsymmetric interior-point solver that provides a Python implementation of the algorithms in Section 3.4.2. Finally, SDPA-C (Fujisawa, Fukuda, Kojima, Nakata, & Yamashita, 2004) is a primal-dual interior-point solver that exploits chordal sparsity using the maximum-determinant positive definite completion.

4. Sparse polynomial optimization

We have seen in Section 3 that the chordal decomposition of large semidefinite matrices allows for significant efficiency gains in the solution of sparse SDPs. The same ideas can often be leveraged to replace SDP relaxations of intractable optimization problems, which generally have no inherent sparsity or other computationally advantageous structure, with SDPs that do.

This section describes how sparsity (primarily chordal, but also nonchordal) can be exploited in the context of sum-of-squares (SOS) relaxation techniques for polynomial optimization. As mentioned in the introduction, SOS methods are at the heart of many recent tractable frameworks for the analysis and optimal control of nonlinear systems with polynomial dynamics; see Ahmadi and Gunluk (2018), Ahmadi, Valmorbida, Gayme and Papachristodoulou (2019), Fantuzzi and Goluskin (2020), Fantuzzi et al. (2016), Goluskin (2020), Han and Tedrake (2018), Henrion and Korda (2014), Jones and Peet (2019), Korda et al. (2021), Lasagna et al. (2016), Lasserre et al. (2008), Majumdar et al. (2014), Miller et al. (2021), Papachristodoulou and Prajna (2005), Prajna et al. (2004), Valmorbida, Ahmadi, and Papachristodoulou (2015), Valmorbida and Anderson (2017) to name but a few contributions.

Our goal is not to offer an exhaustive review of all sparsity-exploiting methods that have been proposed in this field, but rather to introduce the key ideas underpinning most of these methods from a general perspective, in the hope that this can guide further developments. For this reason, we concentrate mainly on two basic problems. The first, discussed in Section 4.2, is to prove that an n -variate polynomial of even degree $2d$ is a sum of squares and, therefore, globally nonnegative. In this case, we seek to exploit the structure of polynomials that depend only on a small subset of all possible degree- $2d$ monomials—a property often referred to as *term sparsity*. The second problem, discussed in Section 4.3, is to check whether a sparse and symmetric n -variate polynomial matrix $P(x)$ is SOS, and therefore

positive semidefinite for all $x \in \mathbb{R}^n$. In this case, our goal is to leverage the *structural sparsity* of P , meaning that many of its entries are zero.

Although we focus only on global nonnegativity, all of the sparsity-exploiting techniques discussed in this section can be extended to prove polynomial (matrix) nonnegativity locally on basic semialgebraic sets. Such extensions, which have been studied extensively in order to build hierarchies of sparse SDP relaxations for polynomial optimization problems (Lasserre, 2006; Waki, Kim, Kojima, & Muramatsu, 2006; Waki, Kim, Kojima, Muramatsu, & Sugimoto, 2008; Wang, Magron, & Lasserre, 2021a, 2021b; Wang, Magron, Lasserre & Mai, 2020; Zheng & Fantuzzi, 2020), require some careful technical adjustments, but the underlying strategy is the same as for the global nonnegativity setting. We outline some of these adjustments in Sections 4.2.5 and 4.3.2, and refer readers to the excellent literature on this topic for full details.

4.1. Background

Let $\mathbb{R}[x]_{n,d}$ be the $\binom{n+d}{d}$ -dimensional space of polynomials with independent variables $x = (x_1, \dots, x_n)$ and degree no larger than d . The n -variate monomial with exponent $\beta = (\beta_1, \dots, \beta_n) \in \mathbb{N}^n$ and degree $|\beta| = \beta_1 + \dots + \beta_n$ is denoted by $x^\beta = x_1^{\beta_1} x_2^{\beta_2} \dots x_n^{\beta_n}$. Given a finite set of exponents $\mathbb{B} \subseteq \mathbb{N}^n$, we write $x^\mathbb{B} = (x^\beta)_{\beta \in \mathbb{B}}$ for the (column) vector of monomials with exponents in \mathbb{B} . The cardinality of \mathbb{B} is denoted by $|\mathbb{B}|$. We also define

$$\mathbb{B} + \mathbb{B} := \{\beta + \gamma : \beta, \gamma \in \mathbb{B}\}, \quad (4.1a)$$

$$2\mathbb{B} := \{2\beta : \beta \in \mathbb{B}\}. \quad (4.1b)$$

If $\mathbb{N}_d^n = \{\beta \in \mathbb{N}^n : |\beta| \leq d\}$ is the set of all n -variate exponents of degree d or less, the vector $x^{\mathbb{N}_d^n}$ is a basis for $\mathbb{R}[x]_{n,d}$ and any polynomial $f \in \mathbb{R}[x]_{n,d}$ can be written as $f(x) = \sum_{\beta \in \mathbb{N}_d^n} f_\beta x^\beta$ for some coefficients $f_\beta \in \mathbb{R}$. The set of exponents with nonzero coefficient,

$$\text{supp}(f) = \{\beta \in \mathbb{N}_d^n : f_\beta \neq 0\}, \quad (4.2)$$

is called the *support* of f . Its convex hull is called the *Newton polytope* of f and is denoted by $\text{New}(f)$.

4.1.1. SOS polynomials and SDPs

A polynomial $f \in \mathbb{R}[x]_{n,2d}$ of even degree $2d$ is SOS if there exist degree- d polynomials $f_1, \dots, f_k \in \mathbb{R}[x]_{n,d}$ such that

$$f = f_1^2 + \dots + f_k^2. \quad (4.3)$$

The set of n -variate degree- $2d$ SOS polynomials, denoted by $\Sigma_{n,2d}$, is a proper cone in $\mathbb{R}[x]_{n,2d}$ (Blekherman, Parrilo, & Thomas, 2012, Theorem 3.26). Given an exponent set $\mathbb{A} \subseteq \mathbb{N}_{2d}^n$, we define the subcone of SOS polynomials supported on \mathbb{A} as

$$\Sigma[\mathbb{A}] := \{f \in \Sigma_{n,2d} : \text{supp}(f) \subseteq \mathbb{A}\}. \quad (4.4)$$

It is well known (see, e.g., Parrilo, 2003, 2013) that a polynomial $f \in \mathbb{R}[x]_{n,2d}$ is SOS if and only if there exist a set of exponents $\mathbb{B} \subseteq \mathbb{N}_d^n$ and a positive semidefinite matrix $Q \in \mathbb{S}_+^{|\mathbb{B}|}$ such that

$$f(x) = (x^\mathbb{B})^\top Q x^\mathbb{B}. \quad (4.5)$$

In particular, if f is SOS, the so-called *Gram matrix representation* in (4.5) is guaranteed to exist with (Reznick, 1978)

$$\mathbb{B} = \frac{1}{2} \text{New}(f) \cap \mathbb{N}_d^n. \quad (4.6)$$

The exponent set obtained with this Newton polytope reduction can be simplified further using more general *facial reduction* techniques (Löfberg, 2009b; Permenter & Parrilo, 2014a, 2014b; Waki & Muramatsu, 2010). These techniques analyze the support of f in order to remove redundant elements from \mathbb{B} , and construct a smaller exponent set for which (4.5) is guaranteed to hold as long as f is SOS.

It is clear that SOS polynomials are nonnegative globally. The converse is true only for univariate polynomials ($n = 1$, d arbitrary),

quadratic polynomials ($d = 1$, n arbitrary), and bivariate quartics ($n = 2$, $d = 2$) (Hilbert, 1888). In general, therefore, being SOS is only a sufficient condition for global nonnegativity, and there are well-known examples of nonnegative polynomials that are not SOS, such the Motzkin polynomial (Motzkin, 1967). However, while verifying polynomial nonnegativity is an NP-hard problem (Murty & Kabadi, 1987), checking whether a polynomial f is SOS can be done in polynomial time by solving an SDP. Specifically, for each exponent $\alpha \in \mathbb{B} + \mathbb{B}$, let $A_\alpha \in \mathbb{S}^{|\mathbb{B}|}$ be the symmetric binary matrix satisfying

$$[A_\alpha]_{\beta, \gamma} := \begin{cases} 1, & \beta + \gamma = \alpha, \\ 0, & \text{otherwise,} \end{cases} \quad (4.7)$$

and observe that

$$(x^\mathbb{B})^\top Q x^\mathbb{B} = \langle Q, x^\mathbb{B}(x^\mathbb{B})^\top \rangle = \sum_{\alpha \in \mathbb{B} + \mathbb{B}} \langle Q, A_\alpha \rangle x^\alpha. \quad (4.8)$$

Then, condition (4.5) holds if and only if $\langle Q, A_\alpha \rangle = f_\alpha$ for all $\alpha \in \mathbb{B} + \mathbb{B}$ and we conclude that

$$f \in \Sigma_{n,2d} \iff \exists Q \in \mathbb{S}_+^{|\mathbb{B}|} \text{ such that } \langle Q, A_\alpha \rangle = f_\alpha \forall \alpha \in \mathbb{B} + \mathbb{B}. \quad (4.9)$$

The condition on the right-hand side defines an SDP, so a positive semidefinite Gram matrix Q certifying that f is SOS can (in principle) be constructed in polynomial time.

4.1.2. SOS polynomial matrices and SDPs

Let $\mathbb{R}[x]_{n,d}^{r \times s}$ be the space of $r \times s$ matrices whose entries are n -variate polynomials of degree d . We say that a symmetric polynomial matrix $P \in \mathbb{R}[x]_{n,2d}^{r \times r}$ is positive semidefinite (resp. definite) globally if $P(x) \geq 0$ (resp. $P(x) > 0$) for all $x \in \mathbb{R}^n$. We also say that P is positive semidefinite locally on a set \mathbb{K} if the same conditions hold for $x \in \mathbb{K}$, but not necessarily otherwise.

A symmetric polynomial matrix $P \in \mathbb{R}[x]_{n,2d}^{r \times r}$ is called SOS if there exist an integer s and a polynomial matrix $M \in \mathbb{R}[x]_{n,d}^{s \times r}$ such that

$$P(x) = M(x)^\top M(x). \quad (4.10)$$

The set of $r \times r$ SOS polynomial matrices with entries in $\mathbb{R}[x]_{n,2d}$ will be denoted by $\Sigma_{n,2d}^r$. All SOS polynomial matrices are clearly positive semidefinite globally, and the converse is true in the univariate case ($n = 1$); see Aylward, Itani, and Parrilo (2007) for a recent proof.

It is well known (see, e.g., Gatermann & Parrilo, 2004; Kojima, 2003; Parrilo, 2013) that a symmetric polynomial matrix $P \in \mathbb{R}[x]_{n,2d}^{r \times r}$ is SOS if and only if it admits a Gram matrix representation in the form

$$P(x) = (I_r \otimes x^\mathbb{B})^\top Q (I_r \otimes x^\mathbb{B}) \quad (4.11)$$

for some exponent set $\mathbb{B} \subseteq \mathbb{N}_d^n$ and some positive semidefinite symmetric matrix $Q \in \mathbb{S}_+^{r|\mathbb{B}|}$. One may always take $\mathbb{B} = \mathbb{N}_d^n$, and smaller exponent sets can be constructed with the same reduction techniques used for SOS polynomials. As in the scalar case ($r = 1$), condition (4.11) defines a set of affine constraints on Q , so verifying that a polynomial matrix is SOS amounts to solving an SDP.

4.2. Sparse SOS decompositions

A major obstacle to constructing SOS certificates of global polynomial nonnegativity via semidefinite programming is that the matrix Q is both dense and very large. If $f \in \mathbb{R}[x]_{n,2d}$ has dense support $\text{supp}(f) = \mathbb{N}_{2d}^n$, then one must take $\mathbb{B} = \mathbb{N}_d^n$ and Q is a $\binom{n+d}{d} \times \binom{n+d}{d}$ dense matrix. Often, however, the support of f is small, i.e., $|\text{supp}(f)|$ is much smaller than $\binom{n+d}{2d}$. This property, called *term sparsity* (Wang, Li, & Xia, 2019; Wang et al., 2021a, 2021b; Wang, Magron et al., 2020), can be exploited in various ways to reduce the computational complexity of the SDP in (4.9).

The facial reduction techniques mentioned above, which replace the full exponent set \mathbb{N}_d^n with a (sometimes significantly) smaller subset, are arguably the simplest way to exploit term sparsity. However, as the next example demonstrates, they are often not sufficient.

Example 4.1. Fix $n = 50$ and $d = 2$. The support of

$$f(x) = \sum_{i=2}^{49} (x_{i-1} + x_i + x_{i+1})^4 \quad (4.12)$$

contains only 485 out of the $\binom{50+4}{4} = 316251$ possible monomials, so f is term sparse. However, it is not hard to check that the Newton polytope $\text{New}(f)$ consists of all points $\xi \in \mathbb{R}_+^{50}$ with $\|\xi\|_1 = 4$, so the Newton-reduced exponent set $\mathbb{B} = \frac{1}{2} \text{New}(f) \cap \mathbb{N}_2^{50} = \mathbb{N}_2^{50} \setminus \mathbb{N}_1^{50}$ contains all homogeneous exponents of degree 2. Therefore, Newton polytope reduction removes only $\binom{50+1}{1} = 51$ of the possible $\binom{50+2}{2} = 1326$ in the full set \mathbb{N}_2^{50} , and the SDP in (4.9) still involves a 1275×1275 Gram matrix Q . ■

Techniques to exploit term sparsity beyond what can be achieved with facial reduction methods alone are clearly desirable. Section 4.2.1 describes a general strategy to search for *sparse* SOS decompositions, which is based on the same matrix decomposition approach used to tackle large-scale sparse SDPs in Section 3. Sections 4.2.2–4.2.4 show that different types of sparse SOS decompositions proposed in the literature are particular cases of this general approach. Section 4.2.5 outlines how these methods can be extended to prove polynomial nonnegativity on basic semialgebraic sets, rather than globally. Throughout, \mathbb{B} will denote a fixed set of candidate exponents for the SOS decomposition of a polynomial f , generated from \mathbb{N}_d^n using facial reduction or any other exponent selection technique.

4.2.1. General approach

Let \mathbb{A} be a small subset of \mathbb{N}_{2d}^n and f be a term-sparse polynomial supported on \mathbb{A} . To reduce the cost of testing if f is SOS, a natural idea is to check whether f belongs to a subset of the sparse SOS cone $\Sigma[\mathbb{A}]$ that admits a semidefinite representation with low computational complexity. Such a subset can be constructed using a simple strategy: *prescribe a sparsity graph $\mathcal{G}(\mathbb{B}, \mathcal{E})$ for the Gram matrix Q and impose its positive semidefiniteness through matrix decomposition.*

Precisely, let $\mathcal{G}(\mathbb{B}, \mathcal{E})$ be a graph with maximal cliques C_1, \dots, C_t and with edge set $\mathcal{E} \subseteq \mathbb{B} \times \mathbb{B}$ satisfying

$$\mathbb{A} \subseteq \{\beta + \gamma : (\beta, \gamma) \in \mathcal{E}\}. \quad (4.13)$$

Consider the cone of sparse SOS polynomial whose Gram matrix Q has sparsity graph \mathcal{G} and admits the clique-based positive semidefinite decomposition

$$Q = \sum_{k=1}^t E_{C_k}^\top S_k E_{C_k}, \quad S_k \in \mathbb{S}_+^{|C_k|}. \quad (4.14)$$

We denote this cone by

$$\Sigma[\mathbb{A}; \mathcal{E}] := \{f \in \Sigma[\mathbb{A}] : f(x) = (x^\mathbb{B})^\top Q x^\mathbb{B}, Q \text{ satisfies (4.14)}\}. \quad (4.15)$$

Conditions (4.13) and (4.14) imply that $\Sigma[\mathbb{A}; \mathcal{E}] \subseteq \Sigma[\mathbb{A}]$. Moreover, inserting the clique-based decomposition (4.14) of Q into (4.9) one finds that $f \in \Sigma[\mathbb{A}; \mathcal{E}]$ if and only if

$$\begin{aligned} \exists S_1 \in \mathbb{S}_+^{|C_1|}, \dots, S_t \in \mathbb{S}_+^{|C_t|} \text{ such that} \\ \sum_{k=1}^t \langle S_k, E_{C_k} A_\alpha E_{C_k}^\top \rangle = f_\alpha \quad \forall \alpha \in \mathbb{B} + \mathbb{B}. \end{aligned} \quad (4.16)$$

If the cliques of the prescribed sparsity graph are small, condition (4.16) defines an SDP with small semidefinite cones and can be solved more efficiently than (4.9).

Remark 4.1 (Chordality of the Sparsity Graph). The Gram matrix decomposition (4.14) is motivated by the chordal decomposition result in Theorem 2.1. However, we do not assume here that the sparsity graph $\mathcal{G}(\mathbb{B}, \mathcal{E})$ is chordal, so (4.14) is generally *not* equivalent to requiring $Q \in \mathbb{S}_+^{|\mathbb{B}|}(\mathcal{E}, 0)$. The lack of chordality makes searching for the maximal cliques C_1, \dots, C_t an NP-hard problem (Tomita, Tanaka, & Takahashi, 2006). Allowing for nonchordal graphs with small cliques that can




Fig. 4.1. (a) Sparsity graph $\mathcal{G}(\mathbb{B}, \mathcal{E})$ for Example 4.2. The vertices \mathbb{B} are such that $x^{\mathbb{B}} = (x_1, x_1 x_2, x_2, x_3, x_2 x_3)$. (b) Sparsity pattern of the Gram matrix of SOS polynomials in $\Sigma[\mathbb{A}; \mathcal{E}]$.

be determined analytically, however, can be extremely useful when a chordal extension leads to unacceptably large cliques, even if it is approximately minimal. Examples of this situation can be found in works by Nie and Demmel (2009) and Kočvara (2020). ■

Remark 4.2 (Sparse SOS Decompositions). Given a sparsity graph $\mathcal{G}(\mathbb{B}, \mathcal{E})$, the cone $\Sigma[\mathbb{A}; \mathcal{E}] \subset \Sigma[\mathbb{A}]$ contains special SOS polynomials that admit a *sparse SOS decomposition*, i.e., a decomposition into a sum of sparse SOS polynomials. Indeed, substituting (4.14) into (4.5) yields

$$f(x) = \sum_{k=1}^t (x^{\mathbb{B}})^{\top} E_{C_k}^{\top} S_k E_{C_k} x^{\mathbb{B}} = \sum_{k=1}^t \underbrace{(E_{C_k} x^{\mathbb{B}})^{\top} S_k (E_{C_k} x^{\mathbb{B}})}_{=: \sigma_k(x)}. \quad (4.17)$$

Each polynomial $\sigma_k(x)$ is SOS because S_k is positive semidefinite, and is sparse because the operation $E_{C_k} x^{\mathbb{B}}$ extracts a subset of the full monomial vector $x^{\mathbb{B}}$. ■

It is important to observe that the reduction in computational complexity granted by the clique-based decomposition (4.14) usually comes at the expense of conservatism. This is because sparsity in the support set \mathbb{A} does not guarantee the existence of a sparse Gram matrix Q . For a given support set \mathbb{A} , special choices of the sparsity graph $\mathcal{G}(\mathbb{B}, \mathcal{E})$ may ensure that $\Sigma[\mathbb{A}; \mathcal{E}] = \Sigma[\mathbb{A}]$ (Zheng & Fantuzzi, 2020, Corollaries 4.1 & 4.2; Mai, Magron, & Lasserre, 2020, Theorem 2.1; Wang et al., 2019, Theorem 4.1; Wang et al., 2021b, Theorem 3.3). In general, however, sparse SOS polynomials need not admit a sparse SOS decomposition, so the inclusion $\Sigma[\mathbb{A}; \mathcal{E}] \subset \Sigma[\mathbb{A}]$ is *strict*. The next example illustrates this.

Example 4.2. (Klep, Magron, & Povh, 2019, Lemma 5.2) Consider the polynomial $f(x) = x_1^2 - 2x_1 x_2 + 3x_2^2 - 2x_1^2 x_2 + 2x_1^2 x_2^2 - 2x_2 x_3 + 6x_3^2 + 18x_2^2 x_3 - 54x_2 x_3^2 + 142x_2^2 x_3^2$ and set $\mathbb{A} = \text{supp}(f)$. Let $\mathbb{B} \subset \mathbb{N}_3^2$ be the exponent set such that $x^{\mathbb{B}} = (x_1, x_1 x_2, x_2, x_3, x_2 x_3)$, which is obtained via Newton polytope reduction. Consider also the (chordal) sparsity graph $\mathcal{G}(\mathbb{B}, \mathcal{E})$ shown in Fig. 4.1, which satisfies (4.13). We claim that f belongs to $\Sigma[\mathbb{A}]$ but not to $\Sigma[\mathbb{A}; \mathcal{E}]$. To see this, observe that any Gram matrix representation of f must take the form

$$f(x) = \begin{pmatrix} x_1 \\ x_1 x_2 \\ x_2 \\ x_3 \\ x_2 x_3 \end{pmatrix}^{\top} \underbrace{\begin{pmatrix} 1 & -1 & -1 & 0 & \alpha \\ -1 & 2 & 0 & -\alpha & 0 \\ -1 & 0 & 3 & -1 & 9 \\ 0 & -\alpha & -1 & 6 & -27 \\ \alpha & 0 & 9 & -27 & 142 \end{pmatrix}}_Q \begin{pmatrix} x_1 \\ x_1 x_2 \\ x_2 \\ x_3 \\ x_2 x_3 \end{pmatrix},$$

where $\alpha \in \mathbb{R}$ can be chosen arbitrarily. Setting $\alpha = 1$ makes the Gram matrix Q positive semidefinite, so $f \in \Sigma[\mathbb{A}]$. However, f cannot be in $\Sigma[\mathbb{A}; \mathcal{E}]$ because this would require $\alpha = 0$, for which Q is not positive semidefinite. ■

4.2.2. Correlative sparsity

The general approach presented in Section 4.2.1 requires specifying the sparsity graph for the Gram matrix Q in (4.5). A natural strategy to do this, pioneered by Waki et al. (2006) and Lasserre (2006), is to consider the couplings between any two independent variables x_i and x_j in a polynomial f supported on \mathbb{A} . Two variables x_i and x_j are considered coupled if a monomial in the vector $x^{\mathbb{A}}$ depends on both simultaneously, i.e., if there exists $\alpha \in \mathbb{A}$ with $\alpha_i \alpha_j > 0$. These couplings can be described using the *correlative sparsity (csp) graph* of the support set \mathbb{A} (or, alternatively, of the polynomial f), which has vertices $\{1, \dots, n\}$ and edge set

$$S_{\text{csp}}(\mathbb{A}) := \{(i, j) : \exists \alpha \in \mathbb{A} \text{ with } \alpha_i \alpha_j > 0\}. \quad (4.18)$$

Correlatively sparse SOS decompositions are obtained upon imposing that the entry $Q_{\gamma, \beta}$ of the Gram matrix in (4.5) vanishes if the monomial $x^{\beta+\gamma}$ introduces couplings between variables that are not consistent with the csp graph of f . This amounts to requiring that Q has sparsity graph $\mathcal{G}_{\text{csp}}(\mathbb{B}, \mathcal{E}_{\text{csp}})$ with edge set

$$\mathcal{E}_{\text{csp}} := \{(\beta, \gamma) \in \mathbb{B} \times \mathbb{B} : (\beta_i + \gamma_i)(\beta_j + \gamma_j) > 0 \Rightarrow (i, j) \in S_{\text{csp}}(\mathbb{A})\}. \quad (4.19)$$

One may consider $\mathcal{G}(\mathbb{B}, \mathcal{E}_{\text{csp}})$ a “hypergraph” with $|\mathbb{B}|$ nodes, built from the csp graph of f (which has n nodes) to ensure that polynomials $(x^{\mathbb{B}})^{\top} Q x^{\mathbb{B}}$ with $Q \in \mathbb{S}^{|\mathbb{B}|}(\mathcal{E}_{\text{csp}}, 0)$ inherit the correlative sparsity of the original support set \mathbb{A} . Unsurprisingly, therefore, the properties of $\mathcal{G}(\mathbb{B}, \mathcal{E}_{\text{csp}})$ can be inferred from those of the (usually much smaller) csp graph. In the following statement, which can be proved using arguments similar to those given by Zheng (2019, Section 2.4.3), $\text{nnz}(\beta)$ denotes the indices of the nonzero entries of an exponent β .

Proposition 4.1. Suppose that the csp graph of the support set \mathbb{A} has maximal cliques J_1, \dots, J_t . Then, $\mathcal{G}(\mathbb{B}, \mathcal{E}_{\text{csp}})$ has maximal cliques $C_k = \{\beta \in \mathbb{B} : \text{nnz}(\beta) \subseteq J_k\}$ for $k = 1, \dots, t$. Moreover, if the csp graph of \mathbb{A} is chordal, then so is $\mathcal{G}(\mathbb{B}, \mathcal{E}_{\text{csp}})$.

Proposition 4.1 considerably simplifies the construction of the “inflation” matrices E_{C_k} in (4.14), because it suffices to find the maximal cliques of the csp graph of \mathbb{A} without building the (much larger) graph $\mathcal{G}(\mathbb{B}, \mathcal{E}_{\text{csp}})$. In addition, it is not difficult to check that the operation $E_{C_k} x^{\mathbb{B}}$ extracts monomials that depend only on variables indexed by J_k . Using (4.17), one concludes that exploiting correlative sparsity amounts to searching for a sparse SOS decomposition in the form

$$f(x) = \sum_{k=1}^t \sigma_k(x_{J_k}), \quad \sigma_k \text{ is SOS}, \quad (4.20)$$

where x_{J_k} denotes the subset of variables x indexed by J_k (cf. Zheng, Fantuzzi et al., 2019, Theorem 2).

Remark 4.3. Example 4.2 shows that correlatively sparse SOS polynomials need not admit the sparse SOS decomposition (4.20), even if the csp graph is chordal. Thus, the inclusion $\Sigma[\mathbb{A}; \mathcal{E}_{\text{csp}}] \subset \Sigma[\mathbb{A}]$ is generally strict. For further discussion on the existence of sparse SOS decompositions for polynomials with chordal correlative sparsity, see Mai et al. (2020) and Zheng and Fantuzzi (2020). ■

Example 4.3. The quartic polynomial f in (4.12) is correlatively sparse, and the csp graph of its support is chordal with maximal cliques $J_i = \{i, i+1, i+2\}$ for $i = 1, \dots, n-2$. It is clear that f admits a sparse SOS decomposition (4.20) and this can be searched for by solving the SDP in (4.16). Since, for each clique J_i , only six elements in $x^{\mathbb{B}} = x^{\mathbb{N}_n^4 \setminus \mathbb{N}_i^4}$ can be multiplied together without introducing spurious couplings to different cliques, this SDP has semidefinite matrix variables $S_1, \dots, S_{n-2} \in \mathbb{S}_+^6$. Its computational complexity is clearly much lower than the corresponding dense formulation in Example 4.1, and a sparse SOS decomposition for f can be found in less than one second on a standard laptop. ■

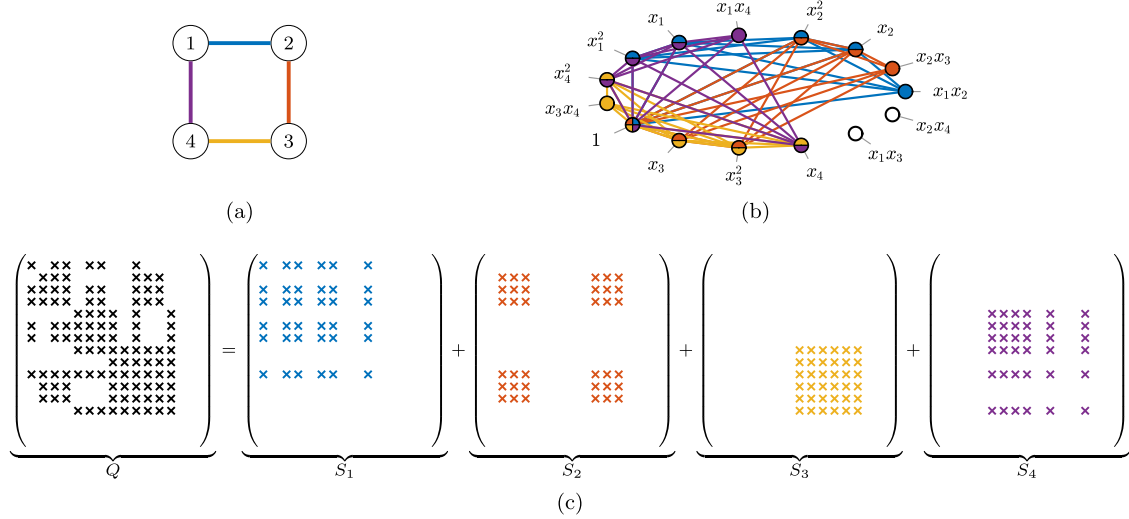


Fig. 4.2. (a) Correlative sparsity graph of the polynomial in Example 4.4. (b) The corresponding sparsity graph $G(\mathbb{B}, \mathcal{E}_{\text{csp}})$. Graph vertices are labeled by monomials in $x^{\mathbb{B}}$ instead of the corresponding exponents in \mathbb{B} to ease the visualization. Filled vertices have a self-loop (not shown), empty ones do not. Colors mark the maximal cliques C_1 (—), C_2 (—), C_3 (—) and C_4 (—). (c) Sparsity pattern of the Gram matrix Q induced by $G(\mathbb{B}, \mathcal{E}_{\text{csp}})$ and its clique-based decomposition. The vertices of $G(\mathbb{B}, \mathcal{E}_{\text{csp}})$ are ordered anticlockwise starting from x_1x_2 . The last two rows and columns of all matrices are empty.

Example 4.4. Consider the quartic polynomial

$$f(x) = 2 + x_1^2x_4^2(x_1^2x_4^2 - 1) - x_1^2 + x_1^4 + \sum_{i=2}^4 (x_i^2x_{i-1}^2(x_i^2x_{i-1}^2 - 1) - x_i^2 + x_i^4).$$

Its csp graph, shown in Fig. 4.2(a), is nonchordal and has maximal cliques $\mathcal{J}_1 = \{1, 2\}$, $\mathcal{J}_2 = \{2, 3\}$, $\mathcal{J}_3 = \{3, 4\}$ and $\mathcal{J}_4 = \{4, 1\}$. The corresponding graph $G(\mathbb{B}, \mathcal{E}_{\text{csp}})$, where the set of exponents obtained with Newton polytope reduction is $\mathbb{B} = \mathbb{N}_2^4$, is shown in Fig. 4.2(b) and has cliques C_1, \dots, C_4 containing 6 elements each, which are determined using Proposition 4.1. The sparsity pattern of the Gram matrix Q induced by $G(\mathbb{B}, \mathcal{E}_{\text{csp}})$ and the clique-based matrix decomposition in (4.14) are also illustrated in the figure. Through this decomposition, a 15×15 positive semidefinite constraint on Q is replaced with four semidefinite constraints on 6×6 matrices S_1, \dots, S_4 . According to (4.20), searching for these matrices is equivalent to looking for a sparse SOS decomposition $f = \sigma_1(x_1, x_2) + \sigma_2(x_2, x_3) + \sigma_3(x_3, x_4) + \sigma_4(x_4, x_1)$. Such a decomposition is not guaranteed to exist even if f were SOS, but it does for this example with

$$\begin{aligned} \sigma_1(x_1, x_2) &= \frac{1}{2} \left(x_1^2 - \frac{1}{2} \right)^2 + \left(x_1x_2 - \frac{1}{2} \right)^2 + \frac{1}{2} \left(x_2^2 - \frac{1}{2} \right)^2 \\ \sigma_2(x_2, x_3) &= \frac{1}{2} \left(x_2^2 - \frac{1}{2} \right)^2 + \left(x_2x_3 - \frac{1}{2} \right)^2 + \frac{1}{2} \left(x_3^2 - \frac{1}{2} \right)^2 \\ \sigma_3(x_3, x_4) &= \frac{1}{2} \left(x_3^2 - \frac{1}{2} \right)^2 + \left(x_3x_4 - \frac{1}{2} \right)^2 + \frac{1}{2} \left(x_4^2 - \frac{1}{2} \right)^2 \\ \sigma_4(x_4, x_1) &= \frac{1}{2} \left(x_4^2 - \frac{1}{2} \right)^2 + \left(x_4x_1 - \frac{1}{2} \right)^2 + \frac{1}{2} \left(x_1^2 - \frac{1}{2} \right)^2. \end{aligned}$$

This proves that $f \in \Sigma[\text{supp}(f); \mathcal{E}_{\text{csp}}]$. ■

4.2.3. TSSOS, chordal-TSSOS and related hierarchies

Fix an exponent set $\mathbb{A} \subset \mathbb{N}_{2d}^n$ and a polynomial f with $\text{supp}(f) \subseteq \mathbb{A}$. Correlative sparsity exploits only the sparse couplings between variables as encoded by the csp graph of \mathbb{A} , but does not take into account any further structure of \mathbb{A} . This is not efficient when $|\mathbb{A}|$ is much smaller than $\binom{n+2d}{2d}$, so f is term-sparse, but the csp graph is fully connected or nearly so.

For this reason, Wang et al. (2019, 2021a, 2021b) introduced the term-sparse-SOS (TSSOS) and the chordal-TSSOS decomposition hierarchies, which exploit term sparsity irrespective of whether f is correlatively sparse. These are two particular examples of a broader family of possible sparsity-exploiting SOS decomposition hierarchies, each of which is obtained upon imposing the clique-based Gram matrix decomposition (4.14) for a sequence $\{G(\mathbb{B}, \mathcal{E}_k)\}_{k \geq 1}$ of increasingly connected sparsity graphs ($\mathcal{E}_k \subseteq \mathcal{E}_{k+1}$).

Irrespective of the particular hierarchy being considered (TSSOS, chordal-TSSOS, or another), the construction of such sparsity graphs begins with the observation that, in order to ensure (4.13), each edge set \mathcal{E}_k should contain at least all edges (β, γ) with $\beta + \gamma \in \mathbb{A}$. This guarantees that $\mathbb{A} \subseteq \text{supp}((x^{\mathbb{B}})^T Q x^{\mathbb{B}})$ for any Gram matrix Q defined via the clique-based decomposition (4.14), which is necessary for the feasibility of the SDP in (4.16). One should also not force diagonal entries $Q_{\beta\beta}$ of the Gram matrix to vanish, because this amounts to saying that the monomial x^β is redundant and β could be removed from the exponent set \mathbb{B} . For these reasons, we define an initial exponent set \mathbb{B}_0 and an initial edge set \mathcal{E}_0 as

$$\mathbb{B}_0 := 2\mathbb{B} \cup \mathbb{A}, \quad (4.21a)$$

$$\mathcal{E}_0 = \{(\beta, \gamma) \in \mathbb{B} \times \mathbb{B} : \beta + \gamma \in \mathbb{B}_0\}. \quad (4.21b)$$

Next, consider an *extension operator* $\mathbb{E} : \mathbb{B} \times \mathbb{B} \rightarrow \mathbb{B} \times \mathbb{B}$, which extends a given edge set $\mathcal{E} \subset \mathbb{B} \times \mathbb{B}$ according to a given rule. The edge sets $\mathcal{E}_1 \subseteq \mathcal{E}_2 \subseteq \dots \subseteq \mathcal{E}_k \subseteq \dots$ and their corresponding support sets \mathbb{B}_k are defined using the iterative rule

$$\mathcal{E}_k := \mathbb{E}(\{(\beta, \gamma) \in \mathbb{B} \times \mathbb{B} : \beta + \gamma \in \mathbb{B}_{k-1}\}), \quad (4.22a)$$

$$\mathbb{B}_k := \{\beta + \gamma : (\beta, \gamma) \in \mathcal{E}_k\}. \quad (4.22b)$$

Note that $\mathcal{E}_k \subseteq \{(\beta, \gamma) \in \mathbb{B} \times \mathbb{B} : \beta + \gamma \in \mathbb{B}_k\}$, so the extension operator guarantees that $\mathcal{E}_k \subseteq \mathcal{E}_{k+1}$. Moreover, the sequence $\{\mathcal{E}_k\}_{k \geq 1}$ must converge to an edge set \mathcal{E}^* in a finite number of iterations because \mathcal{E}_k cannot be extended beyond the complete edge set $\mathbb{B} \times \mathbb{B}$. The sequence of sparsity graphs $\{G(\mathbb{B}, \mathcal{E}_k)\}_{k \geq 1}$ obtained in this way is therefore finite, and yields the (finite) hierarchy of nested sparse SOS cones

$$\Sigma[\mathbb{A}; \mathcal{E}_1] \subseteq \Sigma[\mathbb{A}; \mathcal{E}_2] \subseteq \dots \subseteq \Sigma[\mathbb{A}; \mathcal{E}^*] \subseteq \Sigma[\mathbb{A}]. \quad (4.23)$$

Here, $\Sigma[\mathbb{A}; \mathcal{E}]$ is as defined in (4.15) and all inclusions are strict in general.

Different extension operators produce different types of sparse SOS decomposition hierarchies. In particular:

- If \mathbb{E} is a *block-completion* operator that completes all connected components of the edge set $\{(\beta, \gamma) \in \mathbb{B} \times \mathbb{B} : \beta + \gamma \in \mathbb{B}_{k-1}\}$, one recovers the TSSOS hierarchy (Wang et al., 2019, 2021b). At each step of the hierarchy, Q has chordal sparsity (specifically, a block-diagonal structure) and (4.14) is equivalent to imposing $Q \in \mathbb{S}_+^{|\mathbb{B}|}(\mathcal{E}_k, 0)$.

- If \mathbb{B} is an *approximately minimal chordal extension* operator that extends the edge sets $\{(\beta, \gamma) \in \mathbb{B} \times \mathbb{B} : \beta + \gamma \in \mathbb{B}_{k-1}\}$ such that $\mathcal{G}(\mathbb{B}, \mathcal{E}_k)$ is chordal, one recovers the chordal-TSSOS hierarchy (Wang et al., 2021a). At each step of the hierarchy, \mathcal{Q} has chordal sparsity and (4.14) is equivalent to requiring $\mathcal{Q} \in \mathbb{S}_+^{|\mathbb{B}|}(\mathcal{E}_k, 0)$.

In both cases, the edge extensions are performed on a graph with $|\mathbb{B}|$ nodes and the maximal cliques of $\mathcal{G}(\mathbb{B}, \mathcal{E}_k)$ must be found at each iteration. This is unlike the correlative sparsity strategy in Section 4.2.2, where the maximal cliques of $\mathcal{G}(\mathbb{B}, \mathcal{E}_{\text{csp}})$ are built from those in the csp graph of \mathbb{A} , which has only n nodes (cf. Proposition 4.1).

The choice of extension operator determines the computational complexity of the resulting sparse SOS decomposition hierarchy, as well as the gap between $\Sigma[\mathbb{A}, \mathcal{E}^*]$ and $\Sigma[\mathbb{A}]$. For example, the chordal-TSSOS hierarchy has a lower complexity than the TSSOS one in general, as its sparsity graphs have fewer edges (see Wang et al., 2021a, 2021b for detailed complexity estimates). However, the TSSOS hierarchy has a higher representation power because $\Sigma[\mathbb{A}, \mathcal{E}^*] = \Sigma[\mathbb{A}]$, which is generally not true for the chordal-TSSOS hierarchy.

Theorem 4.1 (Wang et al., 2021b). *If $\mathcal{E}_{\text{TSSOS}}^*$ is the stabilized edge set of the TSSOS hierarchy, then $\Sigma[\mathbb{A}, \mathcal{E}_{\text{TSSOS}}^*] = \Sigma[\mathbb{A}]$, i.e., f is SOS if and only if $f \in \Sigma[\mathbb{A}, \mathcal{E}_{\text{TSSOS}}^*]$.*

Remark 4.4. Theorem 4.1 follows from a stronger result (Wang et al., 2021b, Theorem 6.5) which reveals that the constraint $\mathcal{Q} \in \mathbb{S}_+^{|\mathbb{B}|}(\mathcal{E}_{\text{TSSOS}}^*, 0)$ imposes the well-known block-diagonal structure implied by the sign symmetries of f (see, e.g., Löfberg, 2009b). ■

Example 4.5. The trivariate quartic polynomial

$$f(x) = 1 + x_1^4 + x_2^4 + x_3^4 + x_1^2 x_2^2 + x_1^2 x_3^2 + x_2^2 x_3^2 + x_2 x_3$$

is term sparse but not correlatively sparse, since its csp graph is a complete graph with three nodes. The candidate exponent set to search for an SOS decomposition of f is $\mathbb{B} = \mathbb{N}_2^3$, as Newton polytope reduction removes no terms. For convenience, we order \mathbb{B} such that

$$x^{\mathbb{B}} = (x_3^2, x_2^2, x_1^2, x_2 x_3, 1, x_1, x_1 x_3, x_1 x_2, x_3, x_2)^T.$$

The TSSOS hierarchy yields the sparsity graphs shown in Fig. 4.3, which stabilize at the second iteration ($k = 2$). The corresponding sparsity patterns of the Gram matrix \mathcal{Q} are also shown in that figure. Observe how the connected components of the initial graph $\mathcal{G}(\mathbb{B}, \mathcal{E}_0)$ are completed at the first iteration to obtain the graph $\mathcal{G}(\mathbb{B}, \mathcal{E}_1)$. As discusses in Remark 4.4, the stabilized block-diagonal structure of \mathcal{Q} coincides with the partition of $x^{\mathbb{B}}$ into the groups $\{x_3^2, x_2^2, x_1^2, 1, x_2 x_3\}$, $\{x_1\}$, $\{x_3, x_2\}$ and $\{x_1 x_2, x_1 x_3\}$ implied by the sign symmetries of f , which is invariant under the transformations $(x_1, x_2, x_3) \mapsto (-x_1, -x_2, -x_3)$ and $(x_1, x_2, x_3) \mapsto (x_1, -x_2, -x_3)$ (the four groups of monomials are invariant under both, the first, the second, and none of these transformations). In this example, the SDP in (4.16) is feasible at all iterations of the TSSOS hierarchy because f admits the positive semidefinite Gram matrix representation

$$f(x) = \frac{1}{8} (x^{\mathbb{B}})^T \begin{pmatrix} 8 & 3 & 4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 3 & 8 & 4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 4 & 4 & 8 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2 & 4 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 4 & 8 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} x^{\mathbb{B}}$$

and the Gram matrix is consistent with the sparsity graphs in Fig. 4.3. Thus, all steps of the TSSOS hierarchy are able to prove that f is SOS.

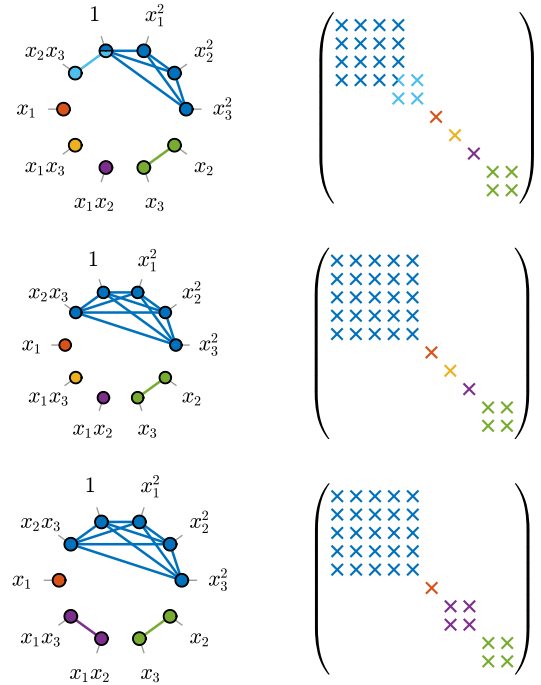


Fig. 4.3. Sparsity graphs and corresponding matrix sparsity patterns for the TSSOS hierarchy in Example 4.5 at initialization (top; edge set \mathcal{E}_0), at the first iteration (middle; edge set \mathcal{E}_1) and at the second iteration (bottom; edge set \mathcal{E}_2). After that, the hierarchy stabilizes. Graph vertices are labeled by monomials in $x^{\mathbb{B}}$ instead of the corresponding exponents in \mathbb{B} to ease the visualization. Colors mark the maximal cliques.




Fig. 4.4. Sparsity graph (left) and corresponding matrix sparsity pattern (right) for the chordal-TSSOS hierarchy in Example 4.5, which stabilizes at the first iteration ($k = 1$). Graph vertices are labeled by monomials in $x^{\mathbb{B}}$ instead of the corresponding exponents in \mathbb{B} to ease the visualization. Colors mark the maximal cliques; multicolored vertices and matrix entries belong to multiple cliques.

Note that this can be guaranteed a priori only for the last step by virtue of Theorem 4.1.

For the same polynomial f , the chordal-TSSOS hierarchy stabilizes at the first iteration ($k = 1$) and yields the sparsity graph shown in Fig. 4.4. The corresponding Gram matrix \mathcal{Q} is sparser than those encountered in the TSSOS hierarchy, leading to smaller semidefinite constraints in (4.16). Again, this SDP is feasible in light of the Gram matrix decomposition given above, so the chordal-TSSOS hierarchy is able to prove that f is SOS. This, however, cannot be guaranteed. ■

Remark 4.5. The explicit Gram matrix decomposition in Example 4.5 reveals that the smaller monomial basis $x^{\mathbb{B}} = (x_3^2, x_2^2, x_1^2, x_2 x_3, 1)$ would suffice to construct an SOS decomposition of f . It remains to be seen whether this reduced basis can be identified using strategies that are more sophisticated than the Newton polytope reduction. ■

4.2.4. Correlatively term-sparse hierarchies

The sparse SOS decomposition hierarchies described in Section 4.2.3 can be combined with the correlative sparsity techniques outlined in

Section 4.2.2 in a natural way. Let \mathcal{E}_k be the edge set obtained using the iterations in (4.22) for a given extension operator, and let \mathcal{E}_{csp} be the edge set in (4.19) constructed using correlative sparsity. Then, the sequence of sparsity graphs

$$\mathcal{G}(\mathbb{B}, \mathcal{E}_k \cap \mathcal{E}_{\text{csp}}), \quad k \geq 1 \quad (4.24)$$

yields a hierarchy of “correlatively term-sparse” SOS decompositions, which exploit simultaneously term and correlative sparsity. Since $\mathcal{E}_k \subseteq \mathcal{E}_{k+1}$ by construction, and since the sequence $\{\mathcal{E}_k\}_{k \geq 1}$ stabilizes onto an edge set \mathcal{E}^* in a finite number of steps, the sparse SOS cones corresponding to this hierarchy satisfy

$$\begin{aligned} \Sigma[\mathbb{A}; \mathcal{E}_1 \cap \mathcal{E}_{\text{csp}}] &\subseteq \Sigma[\mathbb{A}; \mathcal{E}_2 \cap \mathcal{E}_{\text{csp}}] \subseteq \dots \\ &\dots \subseteq \Sigma[\mathbb{A}; \mathcal{E}^* \cap \mathcal{E}_{\text{csp}}] \subseteq \begin{cases} \Sigma[\mathbb{A}; \mathcal{E}^*] \\ \Sigma[\mathbb{A}; \mathcal{E}_{\text{csp}}]. \end{cases} \end{aligned}$$

All inclusions are generally strict. Note also that one may remove from \mathbb{B} all exponents that violate the correlative sparsity *before* constructing the edge sets \mathcal{E}_k , because the intersection with \mathcal{E}_{csp} eliminates all edges between such exponents (including self-loops).

When the extension operator used to build \mathcal{E}_k is the block-completion operation used in the TSSOS hierarchy, the sparsity graphs in (4.24) yield exactly the CS-TSSOS hierarchy introduced by Wang, Magron et al. (2020). In this case, by Theorem 4.1, the stabilized sparsity graph $\mathcal{G}(\mathbb{B}, \mathcal{E}^* \cap \mathcal{E}_{\text{csp}})$ simply encodes sign symmetries and correlative sparsity. Since exploiting sign symmetries in SOS decompositions brings no conservatism (Löfberg, 2009b), one immediately obtains the following corollary.

Proposition 4.2. *If $\mathcal{E}_{\text{TSSOS}}^*$ is the stabilized edge set of the TSSOS hierarchy, then $\Sigma[\mathbb{A}; \mathcal{E}_{\text{TSSOS}}^* \cap \mathcal{E}_{\text{csp}}] = \Sigma[\mathbb{A}; \mathcal{E}_{\text{csp}}]$ for any exponent set $\mathbb{A} \subseteq \mathbb{N}_{2d}^n$.*

Example 4.6. (Wang, Magron et al., 2020, Example 3.4) Let

$$f(x) = 1 + x_1 x_2 x_3 + x_3 x_4 x_5 + x_3 x_4 x_6 + x_3 x_5 x_6 + x_4 x_5 x_6 + \sum_{i=1}^6 x_i^4. \quad (4.25)$$

This polynomial is both term and correlative sparse, and its csp graph has two maximal cliques $\mathcal{J}_1 = \{1, 2, 3\}$ and $\mathcal{J}_2 = \{3, 4, 5, 6\}$. It is also invariant under the sign symmetry transformation $(x_1, x_2, x_3, x_4, x_5, x_6) \mapsto (-x_1, -x_2, x_3, x_4, x_5, x_6)$. To search for an SOS decomposition of f using the CS-TSSOS hierarchy, we let \mathbb{B} be the exponent set defining the monomial vector

$$x^{\mathbb{B}} = (x_2, x_1 x_3, x_2 x_3, x_1, x_1 x_2, x_3, x_5 x_6, x_4 x_6, x_3 x_6, x_5, x_4 x_5, x_3 x_5, x_4, x_3 x_4, x_6, x_1^2, x_2^2, x_3^2, 1, x_6^2, x_5^2, x_4^2).$$

This is obtained upon removing from the full basis $x^{\mathbb{N}_6^6}$ all monomials that violate the correlative sparsity of f (these would be removed anyway by the CS-TSSOS hierarchy).

The CS-TSSOS sparsity graphs obtained with (4.24) and the corresponding sparsity patterns for the Gram matrix Q of f , illustrated in Fig. 4.5, are chordal. The hierarchy stabilizes after three steps. At the first step, the clique-based decomposition (4.14) replaces the 22×22 semidefinite constraint on the Gram matrix Q with six semidefinite constraints of size 2, 2, 2, 10, 4 and 5. The first step of the TSSOS hierarchy, instead, leads to an SDP with five semidefinite constraints of size 2, 2, 2, 10, 7 (Wang, Magron et al., 2020, Example 3.4). The second iteration of the CS-TSSOS hierarchy produces significant fill-in, and the size of the largest semidefinite constraint increases to 15. The third iteration brings only minimal additional fill-in. At this final stage, the connected components of the sparsity graph correspond to a partition of the monomials $x^{\mathbb{B}}$ according to the sign symmetry of f (the first four monomials in $x^{\mathbb{B}}$ are not invariant under the symmetry transformation, while the rest are), but the correlative sparsity prevents the completion

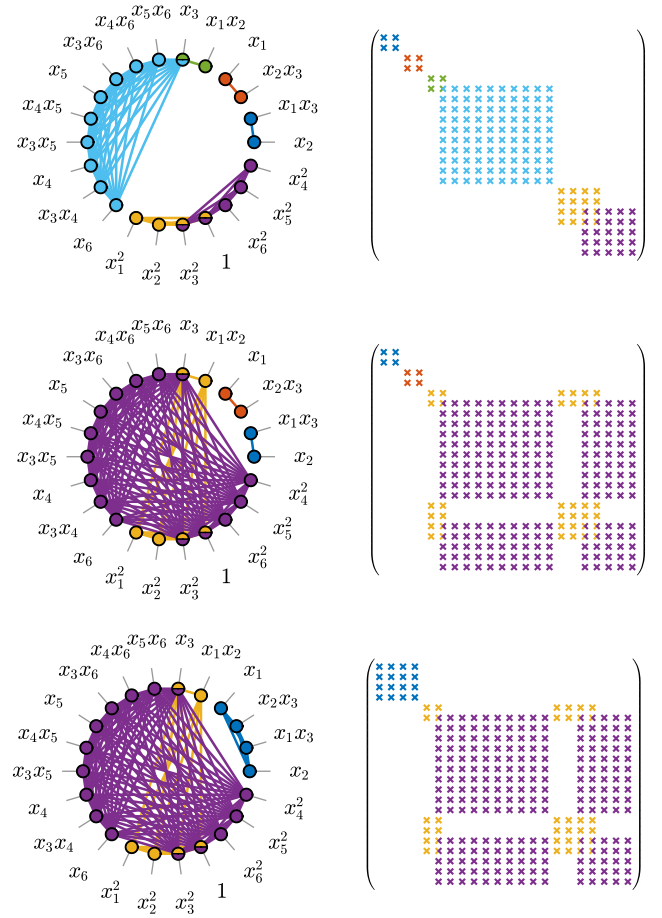


Fig. 4.5. Sparsity graphs and corresponding matrix sparsity patterns for the first (top), second (middle) and third (bottom) iterations of the CS-TSSOS hierarchy in Example 4.6. After that, the hierarchy stabilizes. Graph vertices are labelled by monomials in $x^{\mathbb{B}}$ instead of the corresponding exponents in \mathbb{B} to ease the visualization. Colors mark the maximal cliques; multicolor vertices and matrix entries belong to multiple cliques.

of the largest connected component, i.e., of the bottom-right connected matrix block in Fig. 4.5(c).

Numerical solution of the SDP (4.16) shows that all steps of the CS-TSSOS hierarchy are feasible, so an SOS decomposition of the polynomial f in (4.25) can be constructed at a lower computational cost than any other hierarchy discussed in this work. Note that feasibility cannot be guaranteed a priori at any step of the hierarchy, even the last (stabilized) one, due to the conservative nature of correlative sparse SOS decomposition (see Remark 4.3).

4.2.5. Sparse SOS decompositions on semialgebraic sets

The sparsity-exploiting methods to construct SOS decompositions described so far prove global polynomial nonnegativity, but can be extended to establish local nonnegativity on a basic semialgebraic set defined by m polynomial inequalities,

$$\mathbb{K} := \{x \in \mathbb{R}^n : g_1(x) \geq 0, \dots, g_m(x) \geq 0\}. \quad (4.26)$$

Set $g_0(x) \equiv 1$ for convenience. To verify that $f \in \mathbb{R}[x]_{n,d}$ (not necessarily of even degree) is nonnegative on \mathbb{K} , it suffices to find an integer ω (known as the *relaxation order*) and exponent sets $\mathbb{B}_0, \dots, \mathbb{B}_m \subseteq \mathbb{N}_{\omega}^n$ such that

$$f(x) = \sum_{i=0}^m g_i(x) (x^{\mathbb{B}_i})^T Q_i x^{\mathbb{B}_i}, \quad Q_i \in \mathbb{S}_+^{|\mathbb{B}_i|}. \quad (4.27)$$

As before, these conditions define the feasible set of an SDP. Generally, one chooses ω such that

$$2\omega \geq \max\{\deg(f), \deg(g_1), \dots, \deg(g_m)\}$$

and then takes

$$\mathbb{B}_i = \mathbb{N}_{\omega_i}^n, \quad \omega_i := \omega - \lceil \frac{1}{2} \deg(g_i) \rceil,$$

where $\lceil a \rceil$ is the smallest integer greater than or equal to a . This ensures that each term in the sum in (4.27) is a polynomial of degree at most 2ω . One can allow for $2\omega > \deg(f)$ because cancellations may occur when summing all terms.

Since each polynomial $\sigma_i(x) := (x^{\mathbb{B}_i})^\top Q_i x^{\mathbb{B}_i}$ in (4.27) is SOS, this condition gives a weighted SOS decomposition of f —that is, a representation of f as a weighted sum of SOS polynomials, where the weights are $g_0 = 1$ and the polynomials g_1, \dots, g_m appearing in the semialgebraic definition of \mathbb{K} . Remarkably, this sufficient condition for local nonnegativity is also necessary if f is strictly positive on \mathbb{K} and this is a compact set satisfying the so-called Archimedean condition.

Assumption 1 (Archimedean Condition). There exist an integer $\nu \geq 0$, SOS polynomials $\sigma_0, \dots, \sigma_m \in \Sigma_{n, 2\nu}$, and a constant $r \in \mathbb{R}$ such that $r^2 - \|x\|_2^2 = \sum_{i=0}^m g_i(x) \sigma_i(x)$.

Theorem 4.2 (Putinar, 1993). Suppose that $f \in \mathbb{R}[x]_{n,d}$ is strictly positive on a basic semialgebraic set \mathbb{K} defined as in (4.26) that satisfies the Archimedean condition. Then, there exists a relaxation order ω such that f admits the weighted SOS decomposition (4.27).

If f and the polynomials g_1, \dots, g_m are term sparse, one can proceed as in Section 4.2.1 and attempt to reduce the computational complexity of the generic weighted SOS decomposition (4.27) by requiring the matrices Q_0, \dots, Q_m to be sparse and to admit a clique-based positive semidefinite matrix decomposition. The only new aspect is that one must consider how the sparse polynomial $(x^{\mathbb{B}_i})^\top Q_i x^{\mathbb{B}_i}$ interacts with the corresponding g_i in order to determine the overall structure of the sum on the right-hand side of (4.27). This requires some care, especially if one hopes to recover sparse versions of Theorem 4.2.

To give an example of this general strategy, let us explain how to extend the correlative sparsity technique outlined in Section 4.2.2. In this case, one replaces the csp graph of f constructed with the joint csp graph of the polynomials f, g_1, \dots, g_m , which has vertices $\{1, \dots, n\}$ and an edge between vertices i and j if at least one of the following conditions hold:

- (a) The variables x_i and x_j are multiplied together in f ;
- (b) At least one of g_1, \dots, g_m depends on both x_i and x_j , even if these variables are not multiplied together.

The different treatment of f and g_1, \dots, g_m reflects the asymmetric role these polynomials play in (4.27). Then, one imposes that each matrix Q_i in (4.27) is the densest possible matrix such that the support of $g_i(x)(x^{\mathbb{B}_i})^\top Q_i x^{\mathbb{B}_i}$ is consistent with the joint csp graph. Precisely, let $\mathcal{J}_1, \dots, \mathcal{J}_t$ be the maximal cliques of the joint csp graph and, for each $i = 0, \dots, m$, let $\text{var}(g_i) \subset \{1, \dots, n\}$ be the set of indices of the variables on which g_i depends, with the convention that $\text{var}(g_0) = \text{var}(1) = \emptyset$. By condition (b) above, there is at least one clique \mathcal{J}_k such that $\text{var}(g_i) \subseteq \mathcal{J}_k$ and we denote the set of clique indices k for which this holds by

$$\mathcal{N}_i := \{k \in \{1, \dots, t\} : \text{var}(g_i) \subseteq \mathcal{J}_k\}. \quad (4.28)$$

Observe in particular that $\mathcal{N}_0 = \{1, \dots, t\}$ since $\text{var}(g_0) = \emptyset \subseteq \mathcal{J}_k$ for all $k = 1, \dots, t$. The sparsity graph $\mathcal{G}_i(\mathbb{B}_i, \mathcal{E}_i)$ of Q_i is defined to have the edge set

$$\mathcal{E}_i := \bigcup_{k \in \mathcal{N}_i} \{(\beta, \gamma) \in \mathbb{B}_i \times \mathbb{B}_i : \text{nnz}(\beta + \gamma) \subseteq \mathcal{J}_k\}. \quad (4.29)$$

One can check that $\mathcal{G}_i(\mathbb{B}_i, \mathcal{E}_i)$ is chordal if so is the joint csp graph of f, g_1, \dots, g_m . Moreover, it has maximal cliques $C_{i,1}, \dots, C_{i,|\mathcal{N}_i|}$ with

$C_{i,k} := \{\beta \in \mathbb{B}_i : \text{nnz}(\beta) \subseteq \mathcal{J}_k\}$. Consequently, the clique-based positive semidefinite decomposition of Q_i reads

$$Q_i = \sum_{k=1}^{|\mathcal{N}_i|} E_{C_{i,k}}^\top S_k E_{C_{i,k}}, \quad S_k \in \mathbb{S}_+^{|C_{i,k}|}. \quad (4.30)$$

If the cliques $C_{i,k}$ are small, imposing this clique-based decomposition for each $i = 0, \dots, m$ in (4.27) allows one to search for a weighted SOS decomposition of f by solving an SDP with low computational complexity. Moreover, arguing as in Remark 4.2, one concludes that this process yields the representation

$$f(x) = \sum_{i=0}^m \sum_{k \in \mathcal{N}_i} g_i(x) \sigma_{i,k}(x_{\mathcal{J}_k}), \quad (4.31)$$

where each SOS polynomial $\sigma_{i,k}$ depends only on variables indexed by a single clique of the joint csp graph. Crucially, the following sparse version of Theorem 4.2 guarantees that such a sparse weighted SOS decomposition exists if the joint csp graph is chordal, the semialgebraic definition of the set \mathbb{K} in (4.26) includes inequalities of the form $r_k^2 - \|x_{\mathcal{J}_k}\|_2^2 \geq 0$ for all $k = 1, \dots, t$, and $f > 0$ on \mathbb{K} .

Theorem 4.3 (Grimm, Netzer, & Schweighofer, 2007; Lasserre, 2006). Let f be a polynomial that is strictly positive on a basic semialgebraic set $\mathbb{K} = \{x \in \mathbb{R}^n : g_1(x) \geq 0, \dots, g_m(x) \geq 0\}$, whose definition includes the inequalities $r_k^2 - \|x_{\mathcal{J}_k}\|_2^2 \geq 0$ for some constants r_1, \dots, r_t and all $k = 1, \dots, t$. If the joint csp graph of f, g_1, \dots, g_m is chordal, then f has a sparse weighted SOS decomposition in the form (4.31).

Remark 4.6. The assumption that the semialgebraic definition of \mathbb{K} includes the inequalities $r_k^2 - \|x_{\mathcal{J}_k}\|_2^2 \geq 0$ can be weakened by requiring that the $|\mathcal{J}_k|$ -dimensional set

$$\mathbb{K}_k := \{\hat{x} \in \mathbb{R}^{|\mathcal{J}_k|} : g_i(\hat{x}) \geq 0 \ \forall i \text{ s.t. } \text{var}(g_i) \subseteq \mathcal{J}_k\}$$

satisfies the Archimedean condition for each $k = 1, \dots, t$. Moreover, the assumption is mild when \mathbb{K} is compact because, in principle, the inequalities $r_k^2 - \|x_{\mathcal{J}_k}\|_2^2 \geq 0$ can be added with values of r_k large enough not to change the set \mathbb{K} . Proving that \mathbb{K} remains unchanged for candidate r_k , however, may not be easy in practice. ■

The TSSOS, chordal-TSSOS and CS-TSSOS hierarchies can also be extended to produce weighted SOS decomposition on basic semialgebraic sets (Wang et al., 2021a, 2021b; Wang, Magron et al., 2020). Interested readers are referred to these works for the details. Here we simply observe that, just like their global counterparts described in Sections 4.2.3 and 4.2.4, these extended hierarchies stabilize after a finite number of steps. Upon stabilization, moreover, the extended TSSOS and CS-TSSOS hierarchies recover the block-diagonal structure of the matrices Q_0, \dots, Q_m implied by joint sign symmetries of the polynomials f, g_1, \dots, g_m (see Wang et al., 2021b, Theorem 6.5 and Corollary 6.8; Wang, Magron et al., 2020, Proposition 3.10). This observation can be combined with a symmetry-exploiting version of Theorem 4.2 (Riener, Theobald, Andr n, & Lasserre, 2013, Theorem 3.5) and with Theorem 4.3 to conclude that the TSSOS and CS-TSSOS hierarchies are guaranteed to work for term-sparse polynomials that are strictly positive on compact sets, provided the semialgebraic definition of the latter satisfies suitable versions of the Archimedean condition.

4.3. Decomposition of sparse polynomial matrices

Having studied sparsity-exploiting techniques to reduce the complexity of searching for SOS representations of term-sparse polynomials, we now switch gear and review how chordal sparsity can be exploited when looking for SOS representations of sparse polynomial matrices. Section 4.3.1 presents results by Zheng and Fantuzzi (2020) that partially extend the classical chordal decomposition theorem (Theorem 2.1) to SOS polynomial matrices with chordal sparsity. Decomposition results giving SOS certificates of matrix positivity on semialgebraic

sets are briefly outlined in Section 4.3.2. All of these results are useful for static output controller design (Henrion & Lasserre, 2006), robust stability region analysis (Henrion & Lasserre, 2011), and stability analysis of time-delay systems (Peet, Papachristodoulou, & Lall, 2009). Note that all results presented in this section consider the *structural* sparsity of polynomial matrices, *not* their term sparsity. In principle, one could exploit both structural and term sparsity by combining the results reviewed below with those of Section 4.2.

4.3.1. Global decomposition

Consider a symmetric n -variate polynomial matrix $P \in \mathbb{R}[x]_{n,2d}^{r \times r}$ of degree $2d$ whose (structural) sparsity pattern is described by a chordal graph $\mathcal{G}(\{1, \dots, r\}, \mathcal{E})$, i.e.,

$$(i, j) \notin \mathcal{E} \implies P_{ij}(x) = 0 \quad \forall x \in \mathbb{R}^n. \quad (4.32)$$

Since checking whether P is positive semidefinite globally via the SOS certificates described in Section 4.1.2 is expensive when r is large, we seek to exploit the sparsity of P and replace one large matrix SOS constraint with multiple smaller ones.

Let C_1, \dots, C_t be the maximal cliques of \mathcal{G} . If P is positive semidefinite globally, then applying Theorem 2.1 for each $x \in \mathbb{R}^n$ reveals that there exists x -dependent positive semidefinite matrices $S_k : \mathbb{R}^n \rightarrow \mathbb{S}_+^{|C_k|}$ such that

$$P(x) = \sum_{k=1}^t E_{C_k}^\top S_k(x) E_{C_k}. \quad (4.33)$$

However, this decomposition is not immediately useful in practice because the matrices S_k need not be polynomial, so they cannot be searched for using SOS methods. As an example, consider

$$P(x) = \begin{pmatrix} 2+x^2 & x+x^2 & 0 \\ x+x^2 & 1+2x^2 & x-x^2 \\ 0 & x-x^2 & 2+x^2 \end{pmatrix},$$

whose sparsity graph is a simple three-node chain graph with two maximal cliques, $C_1 = \{1, 2\}$ and $C_2 = \{2, 3\}$. Zheng and Fantuzzi (2020) proved that this matrix is positive definite globally, but does not admit a chordal decomposition (4.33) with polynomial S_1 and S_2 . Using this example, and recalling that all positive semidefinite univariate polynomial matrices are also SOS, one can prove the following general statement.

Proposition 4.3 (Zheng & Fantuzzi, 2020). *Let \mathcal{G} be a connected and not complete chordal graph with $r \geq 3$ vertices and maximal cliques C_1, \dots, C_t . For any positive integers n and d , there exists a positive definite SOS matrix $P \in \Sigma_{n,2d}^r$ with sparsity graph \mathcal{G} that does not admit a decomposition (4.33) with polynomial matrices S_1, \dots, S_t .*

On the other hand, the direct proof of Theorem 2.1 given by Kakimura (2010) can be combined with a diagonalization procedure for polynomial matrices due to Schmüdgen (2009) to show that (4.33) holds with SOS matrices S_1, \dots, S_k for all positive semidefinite polynomial matrices, up to multiplication by an SOS polynomial.

Theorem 4.4 (Zheng & Fantuzzi, 2020). *Let $P \in \mathbb{R}[x]_{n,2d}^{r \times r}$ be positive semidefinite and let C_1, \dots, C_t be the maximal cliques of its sparsity graph. There exist $v \in \mathbb{N}$, an SOS polynomial $\sigma \in \Sigma_{n,2d+2v}$, and SOS polynomial matrices $S_k \in \Sigma_{n,2d+2v}^{|C_k|}$ for $k = 1, \dots, t$, such that*

$$\sigma(x)P(x) = \sum_{k=1}^t E_{C_k}^\top S_k(x) E_{C_k}. \quad (4.34)$$

When the maximal cliques of the sparsity graph of P are small, this result enables one to construct an SOS certificate of global positive semidefiniteness using small matrix SOS constraints, which have a much lower computational complexity than simply requiring σP to be SOS. Implementation of the chordal SOS decomposition in Theorem 4.4

using SDPs requires the matrix P to be fixed, because the SOS weight σ must be determined alongside the SOS matrices S_1, \dots, S_t . Often, however, P depends on a vector of parameters $\lambda \in \mathbb{R}^\ell$ that must be optimized whilst ensuring that P is positive semidefinite. In these cases, condition (4.34) is not jointly convex in λ and σ , so the latter must be fixed *a priori*. This is generally restrictive because, when σ is fixed arbitrarily, Proposition 4.3 implies that the decomposition (4.34) may not exist. However, one can prove a sparse-matrix version of Reznick's Positivstellensatz (Reznick, 1995) to conclude that the weight $\sigma(x) = \|x\|_2^{2v}$ is guaranteed to work at least when P is a homogeneous positive definite matrix.

Theorem 4.5 (Zheng & Fantuzzi, 2020). *Let $P \in \mathbb{R}[x]_{n,2d}^{r \times r}$ be homogeneous of degree $2d$ and positive definite on $\mathbb{R}^n \setminus \{0\}$. Let C_1, \dots, C_t be the maximal cliques of the sparsity graph of P . There exist $v \in \mathbb{N}$ and SOS polynomial matrices $S_k \in \Sigma_{n,2d+2v}^{|C_k|}$ for $k = 1, \dots, t$ such that*

$$\|x\|_2^{2v} P(x) = \sum_{k=1}^t E_{C_k}^\top S_k(x) E_{C_k}.$$

Decomposition results such as this, where the SOS weight σ is fixed, are of considerable interest because they enable the construction of convergent hierarchies of sparsity-exploiting SOS relaxations for optimization problems with global polynomial matrix inequalities (see Henrion & Lasserre, 2006, 2011 and Peet et al., 2009 for particular examples). To illustrate the idea, let us consider the generic convex minimization problem

$$\begin{aligned} b^* &:= \min_{\lambda \in \mathbb{R}^\ell} b(\lambda) \\ \text{subject to } & \underbrace{P_0(x) + \sum_{i=1}^\ell \lambda_i P_i(x)}_{=: P(x, \lambda)} \geq 0 \quad \forall x \in \mathbb{R}^n, \end{aligned} \quad (4.35)$$

where $b : \mathbb{R}^\ell \rightarrow \mathbb{R}$ is a convex cost function and $P_0, \dots, P_\ell \in \mathbb{R}[x]_{n,2d}^{r \times r}$ are symmetric polynomial matrices whose sparsity graph is chordal and has maximal cliques C_1, \dots, C_t . Given any integer $v \geq 0$, a feasible vector λ and an upper bound on the optimal cost b^* may be found by solving the SOS relaxation

$$\begin{aligned} b_v^* &:= \min_{\lambda \in \mathbb{R}^\ell} b(\lambda) \\ \text{subject to } & \|x\|_2^{2v} P(x; \lambda) = \sum_{k=1}^t E_{C_k}^\top S_k(x) E_{C_k}, \\ & S_k(x) \in \Sigma_{n,2d+2v}^{|C_k|} \text{ for } k = 1, \dots, t, \end{aligned} \quad (4.36)$$

which can be reformulated as a standard-form SDP. If the polynomial matrices P_0, \dots, P_ℓ are homogeneous of even degree, and there exists $\lambda_0 \in \mathbb{R}^\ell$ such that $P(x, \lambda_0)$ is positive definite, then one can use Theorem 4.5 to prove that $b_v^* \rightarrow b^*$ from above as $v \rightarrow \infty$; see Zheng and Fantuzzi (2020) for more details and numerical examples. Under further technical assumptions (see Zheng & Fantuzzi, 2020 for details), asymptotic convergence when P_0, \dots, P_ℓ are not homogeneous is preserved by replacing the SOS multiplier $\|x\|_2^{2v}$ with $(1 + \|x\|_2^2)^v$.

4.3.2. Decomposition on a semialgebraic set

We now turn our attention to sparse polynomial matrix inequalities on a semialgebraic set \mathbb{K} defined as in (4.26) by m polynomial inequalities $g_i(x) \geq 0$, $i = 1, \dots, m$. A sufficient condition for a symmetric polynomial matrix $P \in \mathbb{R}[x]_{n,d}^{r \times r}$ to be positive semidefinite on \mathbb{K} is that there exist an integer $v \in \mathbb{N}$ and SOS matrices $S_0, \dots, S_m \in \Sigma_{n,2v}^m$ such that

$$P(x) = S_0(x) + \sum_{i=1}^m g_i(x) S_i(x). \quad (4.37)$$

A matrix version of Putinar's Positivstellensatz proved by Scherer and Hol (2006) states that condition (4.37) is also necessary when P is positive definite on \mathbb{K} and this set satisfies the Archimedean condition.

The weighted matrix SOS decomposition (4.37) can be searched for with semidefinite programming, but this is prohibitively expensive when P is large. If it has chordal structural sparsity, however, one can show that the SOS matrices S_i admit a clique-based decomposition. This yields the following sparse matrix version of Putinar's Positivstellensatz.

Theorem 4.6 (Zheng & Fantuzzi, 2020). *Let \mathbb{K} be a semialgebraic set defined as in (4.26) that satisfies the Archimedean condition. Suppose that the symmetric polynomial matrix $P \in \mathbb{R}[x]_{n,d}^{r \times r}$ is positive definite on \mathbb{K} and that its sparsity graph has maximal cliques C_1, \dots, C_t . There exist an integer $v \in \mathbb{N}$ and SOS matrices $S_{i,k} \in \Sigma_{n,2v}^{|C_k|}$ for $i = 0, \dots, m$ and $k = 1, \dots, t$ such that*

$$P(x) = \sum_{k=1}^t E_{C_k}^T \left(S_{0,k}(x) + \sum_{i=1}^m g_i(x) S_{i,k}(x) \right) E_{C_k}. \quad (4.38)$$

This result can be used to construct sparsity-exploiting SOS relaxations of optimization problems with polynomial matrix inequalities on compact semialgebraic sets that satisfy the Archimedean condition. For example, consider an optimization problem analogous to (4.35), where the polynomial matrix inequality is enforced on \mathbb{K} rather than on the full space \mathbb{R}^n , and denote its optimal value by b^* . If there exists $\lambda_0 \in \mathbb{R}^{\ell}$ such that the inequality is strict on \mathbb{K} and this set satisfies the Archimedean condition, then the optimal value of the SOS problem

$$\begin{aligned} \min_{\lambda \in \mathbb{R}^{\ell}} \quad & b(\lambda) \\ \text{subject to} \quad & P(x, \lambda) \text{ satisfies (4.38)} \\ & S_{i,k} \in \Sigma_{n,2v}^{|C_k|} \text{ for } i = 0, \dots, m \text{ and } k = 1, \dots, t \end{aligned}$$

converges to b^* from above as $v \rightarrow \infty$. Interested readers are referred to Zheng and Fantuzzi (2020) for more details and computational examples.

4.4. Other approaches

The scalability of SOS approaches to polynomial inequalities and polynomial optimization problems can be improved using techniques beyond those described in this section. One example is to replace semidefinite conditions on a large Gram matrix with stronger conditions based on factor-width- k decompositions, which are discussed in Section 5.2. For the particular case of $k = 2$, one obtains *scaled diagonally dominant SOS* (SDSOS) certificates of nonnegativity (Ahmadi & Majumdar, 2019). Another approach is to use *bounded-degree SOS conditions* (Lasserre, Toh, & Yang, 2017), in which (loosely speaking) one restricts the degree of the monomial basis $x^{\mathbb{B}}$ used in the Gram matrix representation and handles monomials of higher degree using positivity certificates that can be reformulated as linear programs. Term sparsity can be exploited in these frameworks, too: the relation between correlative sparsity and SDSOS conditions is discussed by Zheng, Fantuzzi et al. (2019), while Weisser, Lasserre, and Toh (2018) develop sparsity-exploiting bounded-degree SOS hierarchies.

Finally, when working with polynomials that are invariant under groups of symmetry transformations, a large Gram matrix can be replaced with one that has a block-diagonal structure using symmetry reduction techniques (Gatermann & Parrilo, 2004; Löfberg, 2009b; Riener et al., 2013). The block-diagonalization based on sign-symmetries, recovered by the TSSOS and CS-TSSOS hierarchies discussed in Sections 4.2.3 and 4.2.4, is only one particular example; more sophisticated strategies require using a “symmetry-adapted” basis for the space of polynomials in lieu of the monomial basis $x^{\mathbb{B}}$.

4.5. Open-source software implementations

Many of the sparsity-exploiting techniques for polynomial optimization described in this section are implemented in open-source software. The Newton polytope reduction technique is implemented in almost all parsers for SOS optimization, including SOSTOOLS (Prajna et al., 2002), YALMIP (Löfberg, 2004), GloptiPoly (Henrion, Lasserre, & Löfberg, 2009) and SumOfSquares.jl (Legat, Coey, Deits, Huchette, & Perry, 2017; Weisser, Legat, Coey, Kapelevich, & Vielma, 2019). Correlative sparsity techniques are implemented in the MATLAB toolboxes SparsePOP (Waki et al., 2008) and aeroimperial-yalmip (Fantuzzi, 2020). The recent Julia package TSSOS (Magron & Wang, 2021) implements the TSSOS, chordal-TSOS and CS-TSSOS hierarchies. Term sparsity and symmetries in polynomial optimization can also be exploited through SumOfSquares.jl (Legat et al., 2017; Weisser et al., 2019).

5. Factor-width decomposition

We have seen that the matrix decomposition approach can lead to significant efficiency improvements in the solution of sparse SDPs (cf. Section 3) and sparse polynomial optimization problems (cf. Section 4). We now turn our attention to the problem of testing the positive-semidefiniteness of matrices that are not necessarily sparse, for which similar matrix decomposition ideas can also be leveraged using approximation methods. This class of methods is known as *factor-width decomposition* (Boman et al., 2005). We will highlight its connections and differences with the chordal decomposition reviewed above.

After reviewing some background in Section 5.1, we discuss how a hierarchy of inner and outer approximations for positive semidefinite matrices can be constructed based on *factor-width- k matrices* in Section 5.2. We then discuss in Section 5.3 how this can be extended further using the notion of *block factor-width-two matrices* (Zheng, Sootla & Papachristodoulou, 2019), which aims to strike a balance between numerical computation and approximation quality. Applications to semidefinite and SOS optimization are discussed in Sections 5.4 and 5.5.

5.1. Background

As emphasized in the previous sections, solving large-scale semidefinite programs is at the center of many problems in control engineering and beyond, and the development of fast and reliable solution methods (mainly focusing on the exploitation of sparsity exploiting and low-rank structure) has recently attracted significant attention (De Klerk, 2010; Majumdar et al., 2020). Some of these methods attempt to solve the problem exactly, for instance by using chordal decomposition when sparsity is present (cf. Sections 3 and 4), while others try to provide approximate solutions when these problems are large and dense. This section focuses on the latter case and shows that, similar to Sections 3 and 4, certain matrix decomposition strategies allow for significant progress when dealing with large-scale dense SDPs.

One basic way to approach such SDPs is to approximate the positive semidefinite cone \mathbb{S}_+^n with the cone of factor-width- k matrices (Boman et al., 2005), which allows for a certain matrix decomposition discussed in Section 5.2. We will denote the cone of factor-width- k matrices by \mathcal{FW}_k^n , where n is the matrix dimension. The case $k = 2$ is of special interest: \mathcal{FW}_2^n is the cone of symmetric *scaled diagonally dominant* matrices, which can be characterized using second-order cone constraints. Thus, linear functions can be optimized over \mathcal{FW}_2^n by solving second-order cone programs (SOCPs). SOCP strengthenings of SDPs are much more scalable, but can be very conservative: the restricted problem may even become infeasible. Unfortunately, attempting to reduce this conservativeness by approximating \mathbb{S}_+^n with \mathcal{FW}_3^n will result into an SDP with $\mathcal{O}(n^3)$ positive semidefinite constraints, which may not strike a good balance between approximation and computational efficiency.

For this reason, most work has focused on the case of factor-width-two matrices and on some closely related extensions (Ahmadi, Dash, Hall, 2017; Ahmadi & Hall, 2017; Wang, Tanaka & Yoshise, 2021).

The notion of factor-width-two matrices was recently extended to *block-partitioned* matrices by Zheng, Sootla et al. (2019), who showed that the approximation of \mathbb{S}_+^n can be improved significantly compared to using the cone \mathcal{FW}_2^n without sacrificing computational feasibility. This is in contrast to using \mathcal{FW}_3^n . Block factor-width-two matrices also offer a new hierarchy of approximations through a “coarsening” of the decomposition results (cf. Definition 2.2). An alternative approach that results in an improved approximation is based on the use of *decomposed structured subsets* (Miller, Zheng, Sznajder & Papachristodoulou, 2019), but it will not be reviewed here.

5.2. Factor-width- k decompositions

We now introduce the concept of *factor-width- k matrices*, originally defined in Boman et al. (2005).

Definition 5.1. The *factor width* of a matrix $Z \in \mathbb{S}_+^n$ is the smallest integer k for which there exists a matrix V such that $Z = VV^\top$ and each column of V has at most k nonzeros.

The factor width of Z is also the smallest integer k for which Z is the sum of positive semidefinite matrices that are non-zero at most on a $k \times k$ principal submatrix:

$$Z = \sum_{i=1}^s E_{C_i}^\top Z_i E_{C_i} \quad (5.1)$$

for some matrices $Z_i \in \mathbb{S}_+^k$, where C_i is a set of k distinct integers from 1 to n and $s = \binom{n}{k}$. We use \mathcal{FW}_k^n to denote the set of $n \times n$ matrices with factor-width at most k . The dual of \mathcal{FW}_k^n with respect to the normal trace inner product is

$$(\mathcal{FW}_k^n)^* = \left\{ X \in \mathbb{S}^n \mid E_{C_i} X E_{C_i}^\top \in \mathbb{S}_+^k \forall i = 1, \dots, s \right\}.$$

The following hierarchies of inner and outer approximations of \mathbb{S}_+^n follow directly from these definitions:

$$\mathcal{FW}_1^n \subseteq \mathcal{FW}_2^n \subseteq \dots \subseteq \mathcal{FW}_n^n = \mathbb{S}_+^n. \quad (5.2a)$$

$$\mathbb{S}_+^n = (\mathcal{FW}_n^n)^* \subseteq \dots \subseteq (\mathcal{FW}_2^n)^* \subseteq (\mathcal{FW}_1^n)^*. \quad (5.2b)$$

The set \mathcal{FW}_2^n is of particular interest because it is equivalent to the set of symmetric scaled diagonally dominant matrices (Boman et al., 2005). Precisely, a symmetric matrix $A = [a_{ij}]$ is diagonally dominant if $a_{ii} \geq \sum_{j \neq i} |a_{ij}|$, $i = 1, \dots, n$, and it is scaled diagonally dominant if there exists a diagonal positive semidefinite matrix D such that DAD is diagonally dominant. Denote the set of $n \times n$ diagonally dominant matrices as DD^n and the set of $n \times n$ scaled diagonally dominant matrices as SDD^n . We then have the following relationship (Boman et al., 2005):

$$DD^n \subset SDD^n \equiv \mathcal{FW}_2^n.$$

Furthermore, linear optimization over \mathcal{FW}_2^n can be converted into an SOCP, for which efficient algorithms exist. The better scalability of SOCPs compared to SDPs makes inner approximations of positive semidefinite cones based on \mathcal{FW}_2^n very attractive, and forms the basis of the SDSOS framework for polynomial optimization proposed by Ahmadi and Majumdar (2019).

Remark 5.1 (Factor-Width Decomposition vs Chordal Decomposition). The decomposition (5.1) is formally the same as the chordal decomposition in Theorem 2.1, and the two differ only in the choice of “cliques” C_1, \dots, C_s . For chordal decomposition, they are the maximal cliques of (a chordal extension of) the sparsity graph of Z . For factor-width- k decomposition, instead, they are all $\binom{n}{k}$ sets of k distinct indices from $\{1, \dots, n\}$. These two different choices, however, have considerably different implications: while chordal decomposition is necessary

and sufficient for a sparse matrix to be positive semidefinite, factor-width- k decomposition is only sufficient unless $k = n$. The quality of the approximation of positive semidefinite cones by $(\mathcal{FW}_k^n)^*$ was recently investigated by Song and Parrilo (2021) and Blekherman, Dey, Molinaro, and Sun (2020). ■

5.3. Block factor-width-two decomposition

The representation (5.1) reveals that checking whether a matrix Z belongs to \mathcal{FW}_k^n for any values of n and k is equivalent to an SDP. When $k < n$, this SDP has smaller semidefinite cones than \mathbb{S}_+^n , but may be more expensive than checking whether $Z \in \mathbb{S}_+^n$ directly because of the combinatorial number of cones, $\binom{n}{k}$. Setting $k = 2$ does lead to efficiency gains, but the gap between \mathcal{FW}_2^n and \mathbb{S}_+^n might be unacceptably large in some applications. For this reason, *block factor-width-two matrices* are of interest.

Recall from Section 2.3.1 the notion of a block-partition of a matrix $Z \in \mathbb{S}^n$ subordinate to a partition α of n . Recall also the definition of the index matrix $E_{C_{k,\alpha}}$ in (2.14). Here, we further define

$$E_{i,\alpha} := [0 \quad \dots \quad I_{\alpha_i} \quad \dots \quad 0] \in \mathbb{R}^{\alpha_i \times n}, \quad (5.3a)$$

$$E_{ij,\alpha} := [(E_{i,\alpha})^\top \quad (E_{j,\alpha})^\top]^\top \in \mathbb{R}^{(\alpha_i + \alpha_j) \times n}. \quad (5.3b)$$

The set of block factor-width-two matrices, denoted by $\mathcal{FW}_{\alpha,2}^n$, is defined as follows (Zheng, Sootla et al., 2019).

Definition 5.2. For any partition $\alpha = \{\alpha_1, \dots, \alpha_p\}$ of n , a symmetric matrix $Z \in \mathbb{S}^n$ belongs to the class $\mathcal{FW}_{\alpha,2}^n$ of block factor-width-two matrices if and only if

$$Z = \sum_{i=1}^{p-1} \sum_{j=i+1}^p (E_{ij,\alpha})^\top X_{ij} E_{ij,\alpha} \quad (5.4)$$

for some $X_{ij} \in \mathbb{S}_+^{\alpha_i + \alpha_j}$, where $E_{ij,\alpha}$ is defined in (5.3b).

It is clear that (5.4) is a direct block extension of (5.1) when $k = 2$. Also, it is not hard to check that $\mathcal{FW}_{\alpha,2}^n$ is a cone. Its dual with respect to the trace inner product is characterized by the following proposition.

Proposition 5.1 (Zheng, Sootla et al., 2019). For any partition $\alpha = \{\alpha_1, \dots, \alpha_p\}$ of n , the dual of $\mathcal{FW}_{\alpha,2}^n$ is

$$(\mathcal{FW}_{\alpha,2}^n)^* = \{X \in \mathbb{S}^n \mid E_{ij,\alpha} X (E_{ij,\alpha})^\top \geq 0, 1 \leq i < j \leq p\}.$$

Furthermore, both $\mathcal{FW}_{\alpha,2}^n$ and $(\mathcal{FW}_{\alpha,2}^n)^*$ are proper cones, i.e., they are convex, closed, solid, and pointed cones.

It should be clear from Definition 5.2 and Proposition 5.1 that semidefinite programming can be used to verify whether a matrix belongs to $\mathcal{FW}_{\alpha,2}^n$ or to $(\mathcal{FW}_{\alpha,2}^n)^*$. While a gap between these cones and the positive semidefinite cone \mathbb{S}_+^n remains, the next theorem states that the size of the gap can be reduced by coarsening the partition α (cf. Definition 2.2), generally at the expense of increasing the computational complexity of the semidefinite representations of $\mathcal{FW}_{\alpha,2}^n$ and $(\mathcal{FW}_{\alpha,2}^n)^*$. This tradeoff between approximation gap and complexity is the main advantage of using block factor-width-two cones. Recalling the order relation $\alpha \sqsubset \beta$ between two partitions from Definition 2.2, we have the following result.

Theorem 5.1 (Zheng, Sootla et al., 2019). Let $\alpha \sqsubset \beta \sqsubset \gamma$ be partitions of n with $\gamma = \{\gamma_1, \gamma_2\}$, and let $\mathbf{1} = \{1, \dots, 1\}$ denote the uniform unit partition. Then, we have

$$\begin{aligned} \mathcal{FW}_2^n &= \mathcal{FW}_{1,2}^n \subseteq \mathcal{FW}_{\alpha,2}^n \subseteq \mathcal{FW}_{\beta,2}^n \subseteq \mathcal{FW}_{\gamma,2}^n \equiv \mathbb{S}_+^n \\ &\equiv (\mathcal{FW}_{\gamma,2}^n)^* \subseteq (\mathcal{FW}_{\beta,2}^n)^* \subseteq (\mathcal{FW}_{\alpha,2}^n)^* \subseteq (\mathcal{FW}_{1,2}^n)^*. \end{aligned}$$




Fig. 5.1. Regions of the (x, y) plane for which the 5×5 matrix in Example 5.1 belongs to the block factor-width-two cones $F\mathcal{W}_{1,2}^5 \subset F\mathcal{W}_{\alpha,2}^5 \subset F\mathcal{W}_{\beta,2}^5 \subset F\mathcal{W}_{\gamma,2}^5 \equiv \mathbb{S}_+^5$ (top panel), and the dual cones $\mathbb{S}_+^5 \equiv (F\mathcal{W}_{\gamma,2}^5)^* \subset (F\mathcal{W}_{\beta,2}^5)^* \subset (F\mathcal{W}_{\alpha,2}^5)^*$ (bottom panel). The partitions are $\mathbf{1} = \{1, 1, 1, 1, 1\}$, $\alpha = \{2, 1, 1, 1\}$, $\beta = \{2, 1, 2\}$ and $\gamma = \{2, 3\}$. The inclusions of the plotted regions reflect the inclusions of the cones and the order relation $\mathbf{1} \sqsubset \alpha \sqsubset \beta \sqsubset \gamma$.

This result does not quantify how well $F\mathcal{W}_{\alpha,2}^n$ and $(F\mathcal{W}_{\alpha,2}^n)^*$ approximate the positive semidefinite cone. Such information is clearly not only of theoretical interest, but also of practical importance, especially for dense positive semidefinite cone that cannot be studied using chordal decomposition. Some progress in this direction was recently made by Zheng, Sootla et al. (2019), who leveraged results by Blekherman et al. (2020) to show that the normalized distance between either $F\mathcal{W}_{\alpha,2}^n$ or $(F\mathcal{W}_{\alpha,2}^n)^*$ and \mathbb{S}_+^n is at most $\frac{p-2}{p}$, where $p \geq 2$ is the number of blocks in the partition α .

Compared to (5.2), one main advantage of the hierarchy of inner/outer approximations of \mathbb{S}_+^n via block factor-width-two cones in Theorem 5.1 is that the number of basis matrices in the representation (5.4) remains $\mathcal{O}(p^2)$, instead of being a combinatorial number $\binom{n}{k}$. Moreover, the value of p decreases when coarsening the partition. Therefore, the cone $F\mathcal{W}_{\alpha,2}^n$ is often computationally more tractable than the cone $F\mathcal{W}_k^n$ with $k \geq 3$.

Example 5.1. Consider the 5×5 matrix

$$\begin{bmatrix} 1+6x+4y & 3x+y & 2x+y & x+4y & 3x+3y \\ 3x+y & 1+6y & 5x+3y & y & 2x+2y \\ 2x+y & 5x+3y & 1+2x+2y & x+2y & 5x+6y \\ x+4y & y & x+2y & 1+2x & 3x+3y \\ 3x+3y & 2x+2y & 5x+6y & 3x+3y & 1+6x+2y \end{bmatrix}$$

and the progressively coarser partitions $\mathbf{1} = \{1, 1, 1, 1, 1\}$, $\alpha = \{2, 1, 1, 1\}$, $\beta = \{2, 1, 2\}$ and $\gamma = \{2, 3\}$. The regions of the (x, y) plane for which the matrix is in the cones $F\mathcal{W}_{1,2}^5 \subset F\mathcal{W}_{\alpha,2}^5 \subset F\mathcal{W}_{\beta,2}^5 \subset F\mathcal{W}_{\gamma,2}^5 \equiv \mathbb{S}_+^5$ are shown in the top panel in Fig. 5.1. The bottom panel of the same figure, instead, shows the regions of the plane for which the matrix is in the dual cones $\mathbb{S}_+^5 \equiv (F\mathcal{W}_{\gamma,2}^5)^* \subset (F\mathcal{W}_{\beta,2}^5)^* \subset (F\mathcal{W}_{\alpha,2}^5)^*$. It is evident from these figures that all of the inclusions are strict. However, the block factor-width-two cones approximate well the positive semidefinite one along some directions. ■

5.4. Applications to semidefinite programming

Recall from Theorem 5.1 that the cones $F\mathcal{W}_{\alpha,2}^n$ and $(F\mathcal{W}_{\alpha,2}^n)^*$ approximate the positive semidefinite cone \mathbb{S}_+^n from the inside and from the

outside, respectively, and that the approximation improves as the partition α is coarsened. This allows one to compute convergent sequences of upper and lower bounds on the optimal value of an SDP in the primal standard form (3.1) using optimization problems of increasing computational complexity that are always simpler to solve than (3.1) itself. Precisely, since $F\mathcal{W}_{\alpha,2}^n \subseteq \mathbb{S}_+^n$ for any partition α of n , the optimal value of the block factor-width cone program

$$\begin{aligned} U_\alpha &:= \min_X \langle C, X \rangle \\ \text{subject to } & \langle A_i, X \rangle = b_i, \quad i = 1, \dots, m, \\ & X \in F\mathcal{W}_{\alpha,2}^n \end{aligned} \quad (5.5)$$

bounds the optimal value of the SDP (3.1) from above. A complementary lower bound is given by

$$\begin{aligned} L_\alpha &:= \min_X \langle C, X \rangle \\ \text{subject to } & \langle A_i, X \rangle = b_i, \quad i = 1, \dots, m, \\ & X \in (F\mathcal{W}_{\alpha,2}^n)^* \end{aligned} \quad (5.6)$$

because $\mathbb{S}_+^n \subseteq (F\mathcal{W}_{\alpha,2}^n)^*$. By Theorem 5.1, replacing α with a coarser partition can only improve these upper and lower bounds, and we have the following corollary.

Corollary 5.1. Let J^* denote the optimal value of the SDP (3.1) and let $\alpha_1 \sqsubseteq \alpha_2 \sqsubseteq \dots \sqsubseteq \alpha_k = \{\alpha_{k_1}, \alpha_{k_2}\}$ be a sequence of partitions of n . Then, $L_{\alpha_1} \leq \dots \leq L_{\alpha_k} = J^* = U_{\alpha_k} \leq \dots \leq U_{\alpha_1}$.

When $\alpha = \mathbf{1} = \{1, \dots, 1\}$ is the finest possible partition, problems (5.5) and (5.6) can be reformulated as SOCPs. This case was studied extensively by Ahmadi and Majumdar (2019), and numerical experiments show that the optimal values L_1 and U_1 can often be very poor bounds for J^* . To obtain better results using coarser partitions, one can leverage the definition of $F\mathcal{W}_{\alpha,2}^n$ and rewrite the upper bound problem (5.5) as

$$\begin{aligned} \min_{X_{jl}} & \sum_{j=1}^{p-1} \sum_{l=j+1}^p \langle C_{jl,\alpha}, X_{jl} \rangle \\ \text{subject to } & \sum_{j=1}^{p-1} \sum_{l=j+1}^p \langle A_{ijl,\alpha}, X_{jl} \rangle = b_i, \quad i = 1, \dots, m, \\ & X_{jl} \in \mathbb{S}_+^{a_j+a_l}, \quad 1 \leq j < l \leq p, \end{aligned} \quad (5.7)$$

where $C_{jl,\alpha} := E_{jl,\alpha} C (E_{jl,\alpha})^\top$ and $A_{ijl,\alpha} := E_{jl,\alpha} A_i (E_{jl,\alpha})^\top$. This is a standard-form SDP and can be solved with general-purpose solvers. Observe that the number of equality constraints in this SDP is the same as for the original problem (3.1), but the dimension of semidefinite cones has been reduced. Since general-purpose SDP solvers can handle multiple small semidefinite cones much more efficiently than a single large one, problem (5.7) can often be solved much faster than (3.1). For instance, the numerical experiments in Zheng, Sootla et al. (2019) show that useful upper bounds U_α on the optimal value of SDP relaxations of polynomial optimization problems can be found with a reduction of up to 80% in CPU time.

5.5. Applications to SOS optimization

Block factor-width-two decompositions can also be applied to reduce the computational cost of SOS optimization. As discussed in Section 4, an n -variate polynomial $p \in \mathbb{R}[x]_{n,2d}$ of even degree $2d$ is SOS if and only if there exist an exponent set $\mathbb{B} \subseteq \mathbb{N}_d^n$ and a positive semidefinite matrix Q such that (Parrilo, 2000)

$$p(x) = (x^\mathbb{B})^\top Q x^\mathbb{B}. \quad (5.8)$$

The fundamental computational challenge in optimization over the cone $\Sigma_{n,2d}$ of n -variate SOS polynomials of degree at most $2d$ is that the parameterization (5.8) requires in general an $N \times N$ positive

semidefinite matrix with $N = \binom{n+d}{d}$. This may be prohibitive even for moderate values of n and d .

For polynomials characterized by term sparsity, the computational complexity can be reduced dramatically using the approaches reviewed in Section 4, which are based on chordal decomposition. To handle polynomials that are not term sparse, Ahmadi and Majumdar (2019) introduced the notion of *scaled diagonally dominant sum-of-squares* (SDSOS). These are special SOS polynomials whose Gram matrix Q in (5.8) belongs to the factor-width-two cone $\mathcal{FW}_2^{|\mathbb{B}|}$. As in the case of semidefinite programming, defining block-SDSOS polynomials by replacing $\mathcal{FW}_2^{|\mathbb{B}|}$ with its superset $\mathcal{FW}_{\alpha,2}^{|\mathbb{B}|}$ for any partition α of $|\mathbb{B}|$ offers an improved inner approximation of $\Sigma_{n,2d}$.

Definition 5.3. Given a partition $\alpha = \{\alpha_1, \dots, \alpha_g\}$ of $|\mathbb{B}|$, a polynomial $p \in \mathbb{R}[x]_{n,2d}$ is said to be α -SDSOS if and only if there exists coefficient vectors $f_{ij,t} \in \mathbb{R}^{\alpha_i + \alpha_j}$ and exponent sets $\mathbb{B}_{ij} \subseteq \mathbb{N}_d^n$ such that

$$p(x) = \sum_{1 \leq i < j \leq g} \left(\sum_{t=1}^{\alpha_i + \alpha_j} \left(f_{ij,t}^\top x^{\mathbb{B}_{ij}} \right)^2 \right). \quad (5.9)$$

The set of all α -SDSOS polynomials in n independent variables and degree no larger than $2d$ will be denoted by α -SDSOS $_{n,2d}$. It is not difficult to check that it is a cone. Moreover, since definition (5.9) is considerably more structured than the definition (4.3) of general SOS polynomials, the inclusion α -SDSOS $_{n,2d} \subseteq \Sigma_{n,2d}$ is immediate.

For the uniform unit partition $\alpha = \{1, \dots, 1\}$ of $\binom{n+d}{d}$, the cone α -SDSOS $_{n,2d}$ reduces to the normal SDSOS cone studied by Ahmadi and Majumdar (2019). At the other hand of the spectrum, for any partition in the form $\alpha = \{\alpha_1, \alpha_2\}$ one has α -SDSOS $_{n,2d} = \Sigma_{n,2d}$. This second statement is a direct consequence of the following result, which reveals a connection between the polynomial cone α -SDSOS $_{n,2d}$ and the block factor-width-two cone $\mathcal{FW}_{\alpha,2}^{|\mathbb{B}|}$.

Theorem 5.2 (Zheng, Sootla et al., 2019). A polynomial $p \in \mathbb{R}[x]_{n,2d}$ belongs to the cone α -SDSOS $_{n,2d}$ if and only if it admits a Gram matrix representation (5.8) with $\mathbb{B} \subseteq \mathbb{N}_d^n$ and $Q \in \mathcal{FW}_{\alpha,2}^{|\mathbb{B}|}$.

Similar to Theorem 5.1, we can build a hierarchy of inner approximations for the SOS cone $\Sigma_{n,2d}$.

Corollary 5.2. Let $\mathbf{1} = \{1, \dots, 1\}$, $\alpha = \{\alpha_1, \dots, \alpha_g\}$, $\beta = \{\beta_1, \dots, \beta_h\}$ and $\gamma = \{\gamma_1, \gamma_2\}$ be partitions of $|\mathbb{B}|$ such that $\alpha \subseteq \beta$. Then,

$$\begin{aligned} \text{SDSOS}_{n,2d} &= \mathbf{1}\text{-SDSOS}_{n,2d} \subseteq \alpha\text{-SDSOS}_{n,2d} \\ &\subseteq \beta\text{-SDSOS}_{n,2d} \subseteq \gamma\text{-SDSOS}_{n,2d} = \Sigma_{n,2d}. \end{aligned} \quad (5.10)$$

Consider now an optimization problem of the form

$$\begin{aligned} w^* &:= \min_u w^\top u \\ \text{subject to } p(x) &:= p_0(x) + \sum_{i=1}^t u_i p_i(x) \geq 0 \quad \forall x \in \mathbb{R}^n, \end{aligned} \quad (5.11)$$

where $p_0, \dots, p_t \in \mathbb{R}[x]_{n,2d}$ are given polynomials, $w \in \mathbb{R}^t$ is a given cost vector, and $u \in \mathbb{R}^t$ is the decision variable. Let α be any partition of $\binom{n+d}{d}$. To compute an upper bound on the optimal cost w^* , one can strengthen the nonnegativity constraint on p with the SOS constraints $p \in \Sigma_{n,2d}$, the SDSOS constraint $p \in \text{SDSOS}_{n,2d}$, or the block-SDSOS constraint $p \in \alpha$ -SDSOS $_{n,2d}$. The first approach replaces (5.11) with an SDP, the second one leads to an SOCP, and the third yields a block-factor-width cone program that can be reformulated as a standard-form SDP. According to Corollary 5.2, the SOS constraint provides the best upper bound on w^* , but is the most computationally expensive. At the other extreme is the SDSOS constraint, which offers the fastest computations but may be too restrictive—in fact, the corresponding SOCP may even be infeasible. The block-SDSOS constraint $p \in \alpha$ -SDSOS $_{n,2d}$, instead, can balance the computational speed and upper bound quality thanks to the freedom one has in choosing the partition α . This expectation

is confirmed by the numerical experiments of Zheng, Sootla et al. (2019), but the problem of choosing an optimal partition for given computational resources remains an open problem.

6. Applications

The matrix decomposition techniques reviewed in the previous sections can be used to reduce the computational complexity of a wide variety of analysis and control problems that can be formulated as SDPs or SOS programs. As anticipated in Section 3.1, complex large-scale dynamical systems at the heart of modern technology often possess a natural graph-like structure, due for example to sparse interactions between subsystems in a network (Andersen, Hansson et al., 2014; Dall’Anese et al., 2013; Riverso et al., 2014; Zheng, Mason et al., 2018; Zheng et al., 2020). The key to enabling efficient numerical treatment of control problems for such systems is to devise SDP or SOS relaxations that preserve this graph structure as much as possible. Precisely, one aims to obtain SDPs with aggregate sparsity (cf. Section 3) or polynomial optimization problems with term sparsity (cf. Section 4). If this can be done, then the sparsity exploiting techniques discussed in Sections 3 and 4 can bring considerable computational gains and enable the study of very large systems.

This section describes how chordal sparsity can be exploited for a small selection of problems in control and machine learning. Section 6.1 focuses on stability analysis for linear and nonlinear systems, and on decentralized control of networked linear systems. In Section 6.2, we review sparsity-promoting relaxations of nonconvex quadratically constrained quadratic programs (QCQPs) and apply them to the well-known Max-Cut problem from graph theory, as well as to a network sensor location problem. Finally, Section 6.3 shows how chordal sparsity allows for efficient verification of neural networks in machine learning. We stress that these are only a few of the application domains in which chordal decomposition has enabled considerable progress in recent years; other fields include, for instance, fluid mechanics, model predictive control, and optimal power flow. Table 2 provides a (non-exhaustive) list of references.

6.1. Stability analysis and decentralized control

Stability analysis and control synthesis problems for dynamical systems governed by ordinary differential equations can often be reformulated as SDPs or SOS programs using Lyapunov functions (Boyd et al., 1994; Lasserre, 2010; Papachristodoulou & Prajna, 2005; Parrilo, 2000; Zhou et al., 1996). If the interactions between individual components of the system have a sparse graph structure, considering Lyapunov functions with a separable or nearly-separable structure can lead to sparse SDPs and SOS programs, which can be solved efficiently using the techniques in Sections 3 and 4. Here, we give three simple examples of this fact.

6.1.1. Stability of linear networked systems

Consider a continuous-time linear autonomous system

$$\dot{x}(t) = Ax(t), \quad (6.1)$$

where $x(t) \in \mathbb{R}^n$ is the system state at time t and $A \in \mathbb{R}^{n \times n}$ is the system matrix. It is well known (Boyd et al., 1994; Zhou et al., 1996) that the equilibrium state $x(t) = 0$ is asymptotically stable if and only if all eigenvalues of A have negative real part. Classical Lyapunov stability theory guarantees that this is true if and only if there exists a positive definite matrix P such that the (positive definite) Lyapunov function $V(x) = x^\top P x$ decays monotonically along all system trajectories $x(t)$. Equivalently, P must satisfy the strict LMIs

$$P > 0, \quad A^\top P + PA < 0. \quad (6.2)$$

Table 2

Applications of exploiting chordal sparsity in control, machine learning, relaxation of QCQP (Quadratically-constrained quadratic program), fluid dynamics, and beyond.

Area	Topic	References
Control	Linear system analysis	Andersen, Pakazad et al. (2014), Deroo et al. (2015), Mason and Papachristodoulou (2014), Pakazad, Hansson, Andersen and Rantzer (2017), Zheng, Kamgarpour et al. (2018)
	Decentralized control	Deroo, Meinel, Ulbrich, and Hirche (2014), Heinke, Schug, and Werner (2020), Zheng, Kamgarpour, Sootla, and Papachristodoulou (2020), Zheng, Mason et al. (2018)
	Nonlinear system analysis	Schlosser and Korda (2020), Tacchi, Cardozo et al. (2019), Zheng, Fantuzzi et al. (2019); Mason (2015, Chapter 5)
	Model predictive control	Ahmadi, Hansson and Pakazad (2019), Hansson and Pakazad (2018)
Machine learning	Verification of neural networks	Batten et al. (2021), Dvijotham, Stanforth, Goyal, Qin, De, and Kohli (2020), Newton and Papachristodoulou (2021), Zhang (2020)
	Lipschitz constant estimation	Chen, Lasserre et al. (2020), Latorre et al. (2020)
	Training of support vector machine	Andersen and Vandenberghe (2010)
	Geometric perception & coarsening	Chen, Liu, Jacobson and Levin (2020), Liu, Jacobson, and Ovsjanikov (2019), Yang and Carlone (2020)
	Covariance selection	Dahl et al. (2008), Zhang et al. (2018)
	Subspace clustering	Miller, Zheng, Roig-Solvas, Sznaiar and Papachristodoulou (2019)
Relaxation of QCQP and POPs	Sensor network locations	Jing, Wan, and Dai (2019), Kim et al. (2009), Nie (2009)
	Max-Cut problem	Andersen, Dahl et al. (2010), Garstka et al. (2021), Zheng et al. (2020)
	Optimal power flow (OPF)	Andersen, Hansson et al. (2014), Dall'Anese et al. (2013), Jabr (2011), Jiang (2017), Molzahn and Hiskens (2014), Molzahn et al. (2013)
	State estimation in power systems	Weng, Li, Negi, and Ilić (2013), Zhang, Madani, and Lavaei (2017), Zhu and Giannakis (2014)
Others	Fluid dynamics	Arsilan, Fantuzzi, Craske, and Wynn (2021), Fantuzzi, Pershin, and Wynn (2018)
	Partial differential equations	Mevisen (2010), Mevisen, Kojima, Nie, and Takayama (2008), Mevisen, Lasserre, and Henrion (2011), Mevisen, Yokoyama, and Takayama (2009)
	Robust quadratic optimization	Andersen, Vandenberghe et al. (2010)
	Binary signal recovery	Fosson and Abuabiah (2019)
	Solving polynomial systems	Cifuentes and Parrilo (2016, 2017), Li, Xia, Zhang, and Zheng (2021), Mou, Bai, and Lai (2021), Tacchi, Weisser, Lasserre and Henrion (2019)
	Other problems	Baltea-Lugojan, Bonami, Misener, and Tramontani (2019), Jeyakumar, Kim, Lee, and Li (2016), Madani, Sojoudi, Fazelnia and Lavaei (2017), Pakazad, Hansson, Andersen and Nielsen (2017), Yang and Deng (2020)

Now, suppose that (6.1) is a compact representation of a network of l linear subsystems with states $x_1 \in \mathbb{R}^{n_1}, \dots, x_l \in \mathbb{R}^{n_l}$, whose interactions can be represented by a static undirected graph $\mathcal{G}_d(\{1, \dots, l\}, \mathcal{E}_d)$ with $(i, j) \in \mathcal{E}_d$ if and only if systems i and j are directly coupled. In particular, the dynamics of each subsystem are given explicitly by

$$\dot{x}_i = A_{ii}x_i + \sum_{j \in \mathcal{N}_i} A_{ij}x_j, \quad i = 1, \dots, l, \quad (6.3)$$

where $\mathcal{N}_i := \{j : (j, i) \in \mathcal{E}_d\}$ denotes the neighbors of system i . Systems of this type are encountered, for example, when modeling power grids (Riverso et al., 2014) and traffic systems (Wang, Zheng, Chen, Xu & Li, 2020; Zheng et al., 2020).

If the matrix P in (6.2) is assumed to be block-diagonal with l blocks of size n_1, \dots, n_l , meaning that we consider a quadratic Lyapunov function in the separable form (Boyd & Yang, 1989; Geromel, Bernussou, & Peres, 1994; Zheng et al., 2020; Zheng, Mason et al., 2018)

$$V(x) = \sum_{i=1}^l x_i^\top P_i x_i, \quad (6.4)$$

then it is not hard to see that the block-sparsity graph of the matrix $A^\top P + PA$ in (6.2) is the same as the system graph \mathcal{G}_d . When this graph is chordal with small maximal cliques, or admits a chordal extension with the same property, a feasible block-diagonal matrix P satisfying (6.2) can be constructed for significantly larger networks than those that can be handled without sparsity exploitation. Equivalently, for a given network size, CPU time requirements can be reduced dramatically.

As an example, consider a network with a master node and $l - 1$ independent subsystems connected to it, sketched in Fig. 6.1(a). For simplicity, suppose that the subsystems have size $n_1 = \dots = n_l = 10$. With a block-diagonal P , the second LMI in (6.2) has the chordal




Fig. 6.1. (a) Graph \mathcal{G}_d for the network of 10-dimensional linear systems used to generate the results reported in Table 3. (b) Block sparsity pattern of the matrix $PA + A^\top P$ when P is block-diagonal; each block has size 10×10 , and there are l diagonal blocks.

“arrow-type” block sparsity shown in Fig. 6.1(b). Table 3 reports the CPU time required to construct a feasible P with MOSEK as a function of the number l of subsystems when the sparsity of this LMI is and is not exploited.³ It is evident that exploiting chordal sparsity using the methods described in Section 3 leads to a significant reduction in CPU time. Similar results are obtained for systems with more realistic network graphs if their maximal cliques are small; see Deroo et al. (2015), Mason and Papachristodoulou (2014) and Zheng, Kamgarpour et al. (2018), Zheng, Mason et al. (2018).

³ Computations were performed using the MATLAB toolboxes YALMIP and SparseCoLo on a laptop with 16GB RAM and an Intel i7 processor. The nonzero system matrices A_{ij} were generated randomly whilst ensuring the existence of a feasible block-diagonal P .

Table 3

CPU time (in seconds) required by the SDP solver MOSEK to construct a block-diagonal Lyapunov matrix P satisfying the LMIs in (6.2) for a network of l 10-dimensional systems with connectivity graph G_d shown in Fig. 6.1(a).

l	No sparsity exploitation	Sparsity exploitation
10	0.55	0.26
50	14.92	0.90
100	86.09	1.21
125	113.06	1.17
150	185.42	1.96
175	334.13	2.69
200	498.49	3.55

Remark 6.1 (*Separable Lyapunov Functions*). Searching for a Lyapunov function $V(x)$ with the separable structure (6.4) is convenient to ensure that the sparsity of the system matrix A is inherited by the LMI $A^\top P + PA < 0$. The existence of such a separable Lyapunov function can be guaranteed for special classes of stable linear systems (Carlson, Hershkowitz, & Shasha, 1992; Sootla, Zheng, & Papachristodoulou, 2017, 2019), but not in general. When a separable Lyapunov function $V(x)$ fails to exist, the structure of the network graph G_d may be still be leveraged to promote sparsity in (6.2); for instance, the case of banded graphs, cycles and trees was studied by Mason and Papachristodoulou (2014). Determining a suitable structure for $V(x)$ (equivalently, for the matrix P) *a priori* for general graph structures, however, remains a challenging problem. ■

6.1.2. Stability of sparse polynomial systems

Structured Lyapunov functions can bring computational advantages also when studying the asymptotic stability of sparse nonlinear systems with polynomial dynamics. As an example, consider a nonlinear system with the structure (Zheng, Fantuzzi et al., 2019, Section VI.D)

$$\begin{aligned}\dot{x}_1 &= f_1(x_1, x_2), \\ \dot{x}_i &= f_i(x_{i-1}, x_i, x_{i+1}), \quad i = 2, \dots, l-1, \\ \dot{x}_l &= f_l(x_{l-1}, x_l),\end{aligned}\quad (6.5)$$

where each vector field f_i depends polynomially on its arguments and $x_i \in \mathbb{R}^{n_i}$. Let $x = (x_1, \dots, x_l)$ be the collection of all system states and write $f = (f_1, \dots, f_l)$. Suppose the system has an equilibrium at the origin. This equilibrium is locally asymptotically stable if there exists a region $D \subset \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_l}$ containing the origin, a constant $\epsilon > 0$, and a Lyapunov function $V : \mathbb{R}^{n_1} \times \dots \times \mathbb{R}^{n_l} \rightarrow \mathbb{R}$ such that

$$V(0) = 0, \quad (6.6a)$$

$$V(x) \geq \|x\|_2^2 \quad \forall x \in D, \quad (6.6b)$$

$$-f(x) \cdot \nabla V(x) \geq \epsilon \|x\|_2^2 \quad \forall x \in D. \quad (6.6c)$$

Upon fixing $D = \{x : r_i^2 - \|x_i\|^2 \geq 0 \ \forall i = 1, \dots, l\}$, which has a fully separable structure, and requiring V to be a polynomial, the last two inequalities become polynomial inequalities on a basic semialgebraic set. One can therefore search for V using SOS optimization. Moreover, the structure of V can be chosen to ensure that these polynomial inequalities are correlatively sparse (cf. Sections 4.2.2 and 4.2.5), enabling efficient implementation.

For example, if one takes

$$V(x) = \sum_{i=1}^l V_i(x_i) \quad (6.7)$$

to have a fully separable structure as in the case of linear systems considered previously, then the correlative sparsity graph of inequalities (6.6b) is a graph with no edges, while that of (6.6c) is the same chain graph characterizing the cascaded interactions between the state




Fig. 6.2. Correlative sparsity graphs for the polynomial inequality in (6.6c) when the Lyapunov function V has (a) the separable form (6.7), and (b) the partially separable form (6.8).

Table 4

CPU time, in seconds, required by MOSEK to construct a structured quadratic Lyapunov function (6.8) for a locally asymptotically stable, degree-3 polynomial system of the form (6.5). Entries marked oom indicate memory errors.

l	10	15	20	30	40	50
Standard SOS	1.4	21.3	262.1	oom	oom	oom
Sparse SOS	0.6	0.7	0.8	1.0	1.2	1.4

vectors x_1, \dots, x_l , shown in Fig. 6.2(a) for $l = 8$. If this choice for V is insufficient, one can try the structured choice

$$V(x) = \sum_{i=2}^{l-1} V_i(x_{i-1}, x_i, x_{i+1}). \quad (6.8)$$

In this case, the correlative sparsity graph of (6.6b) is the chain graph mentioned above, while that of (6.6c) is a chordal graph with maximal cliques $\{i, i+1, i+2, i+3\}$ for $i = 1, \dots, l-3$, which is shown in Fig. 6.2(b) for $l = 8$. One can of course build an entire hierarchy of structured Lyapunov functions with increasing degree of couplings between subsystem variables, at the expense of increasing the number of edges in the correlative sparsity graph of the polynomial inequalities (6.6b) and (6.6c). Numerical experiments by Zheng, Fantuzzi et al. (2019) for the structured Lyapunov function in (6.8), which we report in Table 4, show that this approach can significantly reduce the computation time and resources required to prove stability for nonlinear systems compared to standard SOS techniques.

Similar ideas can be used to partition nonlinear systems into subsystems (Anderson & Papachristodoulou, 2011) and can be adapted to problems beyond stability analysis, such as the estimation of region of attractions, positively invariant sets, and global attractors (Schlosser & Korda, 2020; Tacchi, Cardozo et al., 2019).

6.1.3. Decentralized control of linear networked systems

Consider a network of linear system with control inputs and disturbances,

$$\dot{x}_i = A_{ii}x_i + \sum_{j \in \mathcal{N}_i} A_{ij}x_j + B_i u_i + M_i d_i, \quad i = 1, \dots, l,$$

where $x_i \in \mathbb{R}^{n_i}$, $u_i \in \mathbb{R}^{m_i}$ and $d_i \in \mathbb{R}^{d_i}$ denote the local state, input, and disturbance of subsystem i , respectively, and \mathcal{N}_i is the index set of all systems connected to system i . Setting $x = (x_1, \dots, x_l)$, $u = (u_1, \dots, u_l)$ and $d = (d_1, \dots, d_l)$, the system can be written compactly as

$$\dot{x}(t) = Ax(t) + Bu(t) + Md(t),$$

where A has block sparsity induced by the system graph (cf. Section 6.1.1), while $B = \text{diag}(B_1, \dots, B_l)$ and $M = \text{diag}(M_1, \dots, M_l)$ are block-diagonal.

The *optimal decentralized control* problem (Geromel et al., 1994) seeks to design static state feedback laws,

$$u_i(t) = -K_{ii}x_i(t), \quad \forall i = 1, \dots, l, \quad (6.9)$$

that minimize the H_2 norm of the transfer function from the disturbance d to the output

$$z = \begin{bmatrix} Q^{\frac{1}{2}} \\ 0 \end{bmatrix} x + \begin{bmatrix} 0 \\ R^{\frac{1}{2}} \end{bmatrix} u.$$

Here, $Q := \text{diag}(Q_1, \dots, Q_l)$ and $R := \text{diag}(R_1, \dots, R_l)$ are given block-diagonal matrices. The decentralized constraint (6.9) makes the control problem challenging to solve (Furieri, Zheng, Papachristodoulou, & Kamgarpour, 2019; Geromel et al., 1994). One simple strategy is to enforce that the closed-loop system admits a separable Lyapunov function in the form (6.4). This allows translating the decentralized constraint on the controller to other auxiliary design variables (Furieri et al., 2019; Furieri, Zheng, Papachristodoulou, & Kamgarpour, 2020). In particular, a suboptimal decentralized controller can be computed using the formula $K_{ii} = Z_i X_i^{-1}$ for each $i = 1, \dots, l$ (Geromel et al., 1994; Zheng et al., 2020, Section II.B), where the matrices Z_1, \dots, Z_l and X_1, \dots, X_l solve the SDP

$$\min_{X_i, Y_i, Z_i} \sum_{i=1}^l \langle Q_i, X_i \rangle + \langle R_i, Y_i \rangle$$

$$\text{subject to } (AX - BZ) + (AX - BZ)^T + MM^T \leq 0, \quad (6.10a)$$

$$\begin{bmatrix} Y_i & Z_i \\ Z_i^T & X_i \end{bmatrix} \geq 0, \quad X_i > 0 \quad \forall i = 1, \dots, l \quad (6.10b)$$

and $X = \text{diag}(X_1, \dots, X_l)$ and $Z = \text{diag}(Z_1, \dots, Z_l)$ are block-diagonal concatenations of the matrix variables.

The cost function of this SDP and the constraints in (6.10b) are fully separable, as they depend only on variables corresponding to a single subsystem. The coupling constraint (6.10a), instead, has a block sparsity pattern induced by the system graph by virtue of the block-diagonal structure of B , M , X and Z . As in Section 6.1.1, therefore, the chordal decomposition techniques of Section 3 allow for a fast numerical solution when the underlying system graph is sparse, which enables control synthesis for large-scale but sparse networks. In addition, customized distributed design methods that combine chordal decomposition with ADMM can solve (6.10) in a privacy-safe way, without requiring subsystems to share information about their local dynamics (Zheng et al., 2020).

6.2. Relaxation of nonconvex QCQPs

A (nonconvex) quadratically constrained quadratic program (QCQP) is an optimization problem in the form

$$\min_x x^T P_0 x + 2q_0^T x + r_0 \quad (6.11)$$

$$\text{subject to } x^T P_i x + 2q_i^T x + r_i \leq 0, \quad i = 1, \dots, m,$$

where $x \in \mathbb{R}^n$ is the optimization variable and $P_i \in \mathbb{S}^n$, $q_i \in \mathbb{R}^n$, $r_i \in \mathbb{R}$, $i = 0, 1, \dots, m$ are given problem data. QCQPs have very powerful modeling capabilities; for instance, many hard combinatorial and discrete optimization problems can be written in the form (6.11) (Nesterov, Wolkowicz, & Ye, 2000). This also means that QCQPs are hard to solve in general, so many different relaxation strategies have been proposed to find approximate bounds and feasible values for the optimization variable x (Nesterov et al., 2000; Park & Boyd, 2017). One approach that provides good bounds, both empirically and theoretically (Nesterov et al., 2000), is to introduce the positive semidefinite matrix $X = xx^T$ and rewrite (6.11) as

$$\min_{x, X} \langle P_0, X \rangle + 2q_0^T x + r_0$$

$$\text{subject to } \langle P_i, X \rangle + 2q_i^T x + r_i \leq 0, \quad i = 1, \dots, m,$$

$$X = xx^T.$$

Upon relaxing the intractable constraint $X = xx^T$ into the inequality $X \geq xx^T$ and applying Schur's complement to rewrite the latter as an

LMI, we arrive at the semidefinite relaxation

$$\min_{x, X} \langle P_0, X \rangle + 2q_0^T x + r_0$$

$$\text{subject to } \langle P_i, X \rangle + 2q_i^T x + r_i \leq 0, \quad i = 1, \dots, m,$$

$$\begin{bmatrix} 1 & x^T \\ x & X \end{bmatrix} \geq 0,$$

which is equivalent to the following primal-form SDP with nonnegative variables

$$\begin{aligned} \min_{Z \in \mathbb{S}^{n+1}, w} & \left\langle \begin{pmatrix} r_0 & q_0^T \\ q_0 & P_0 \end{pmatrix}, Z \right\rangle \\ \text{subject to} & \left\langle \begin{pmatrix} r_i & q_i^T \\ q_i & P_i \end{pmatrix}, Z \right\rangle + w_i = 0, \quad i = 1, \dots, m, \\ & Z_{11} = 1, \\ & Z \geq 0, \quad w \geq 0. \end{aligned} \quad (6.12)$$

It is not difficult to see that the optimal value of problem (6.12) bounds that of the QCQP (6.11) from below and that, if an optimal solution Z_* has rank one, then the relaxation is exact and $Z_* = \begin{pmatrix} 1 & x_*^T \\ x_* & x_* x_*^T \end{pmatrix}$ where x_* solves (6.11).

If the data matrices P_0, \dots, P_m are sparse, then the aggregate sparsity pattern \mathcal{E} of the SDP (6.12) is also sparse, and the positive semidefinite constraint on Z can be replaced with the conic constraint $Z \in \mathbb{S}_+^{n+1}(\mathcal{E}, ?)$. The chordal decomposition techniques described in Section 3 can therefore be applied to solve (6.12) efficiently. The following subsections briefly discuss two types of problem for which sparsity can be exploited effectively: Max-Cut problems (Goemans & Williamson, 1995) and sensor network location problems (Jing et al., 2019; Kim et al., 2009; Nie, 2009; So & Ye, 2007).

6.2.1. Max-Cut problem

The maximum cut (Max-Cut) problem is a classic problem in graph theory (Goemans & Williamson, 1995). Consider an undirected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with n vertices such that each edge $(i, j) \in \mathcal{E}$ is assigned a nonzero weight W_{ij} , and set $W_{ij} = 0$ if $(i, j) \notin \mathcal{E}$. The Max-Cut problem aims to partition the graph's vertices into two complementary sets \mathcal{V}_1 and \mathcal{V}_2 such that the total weight of all edges linking \mathcal{V}_1 and \mathcal{V}_2 is maximized. Given a binary variable $x \in \{-1, +1\}^n$ assigning nodes to one of the two partitions, one seeks to maximize

$$\frac{1}{2} \sum_{i,j: x_i x_j = -1} W_{ij} = \frac{1}{4} \sum_{i,j} W_{ij} (1 - x_i x_j).$$

This is equivalent to solving

$$\min_x x^T W x \quad (6.13)$$

$$\text{subject to } x_i^2 = 1, \quad i = 1, \dots, n,$$

where W is the given matrix of weights.

This problem is a particular QCQP, and can easily be rewritten in the generic form (6.11) using data matrices P_0, P_1, \dots, P_n whose aggregate sparsity graph coincides with the original graph \mathcal{G} . If \mathcal{G} is sparse with small maximal cliques, therefore, SDP relaxations of (6.13) can be solved efficiently using the sparsity-exploiting techniques in Section 3. Indeed, numerical experiments by Andersen, Dahl et al. (2010) and Zheng et al. (2020) demonstrated that the sparsity-exploiting solvers SMCP and CDCS can solve benchmark Max-Cut problems from the SDPLIB problem library (Borchers, 1999) order of magnitude faster than standard conic solvers.

6.2.2. Sensor network location

The sensor network location problem, also known as *Graph Realization* (So & Ye, 2007), has important applications such as inventory management and environment monitoring. At a basic level, the problem is to find unknown *sensor points* $x_1, \dots, x_n \in \mathbb{R}^d$ ($d = 2$ or 3) satisfying some specified distance constraints, as well as distance

constraints with respect to m known *anchor points* $a_1, \dots, a_m \in \mathbb{R}^d$. Precisely, given pairing sets

$$\mathcal{E}_x \subseteq \{1, \dots, n\} \times \{1, \dots, n\},$$

$$\mathcal{E}_a \subseteq \{1, \dots, m\} \times \{1, \dots, n\},$$

we seek to find sensor locations $x_1, \dots, x_n \in \mathbb{R}^d$ such that

$$\begin{aligned} \|x_i - x_j\|^2 &= d_{ij}^2, & (i, j) \in \mathcal{E}_x, \\ \|a_i - x_j\|^2 &= f_{ij}^2, & (i, j) \in \mathcal{E}_a, \end{aligned} \quad (6.14)$$

where the numbers d_{ij} and f_{ij} are specified distances.

One way to relax the sensor location problem into an SDP is to consider (6.14) as a set of quadratic constraints for x_1, \dots, x_n , and apply the generic SDP relaxation strategy to the QCQP (Kim et al., 2009)

$$\begin{aligned} \min_{x_1, \dots, x_n \in \mathbb{R}^d} \quad & 0 \\ \text{subject to} \quad & (6.14). \end{aligned} \quad (6.15)$$

It is clear that the data matrices and vectors of this QCQP are very sparse, and that the aggregate sparsity pattern of the corresponding SDP relaxation is determined only by the edge sets \mathcal{E}_a and \mathcal{E}_x . Then, the techniques in Section 3 can be applied to solve the relaxed problem quickly; we refer the interested reader to Kim et al. (2009) for a more detailed discussion and experiment results. Similar ideas can be used to analyze sensor location problems where the distance measurements d_{ij} and f_{ij} are affected by noise (Kim et al., 2009).

Remark 6.2. There are other ways to formulate an SDP relaxation for (6.15). One (So & Ye, 2007) is to introduce a matrix variable $Y = XX^T$ with $X = [x_1, x_2, \dots, x_n] \in \mathbb{R}^{d \times n}$, rewrite all the constraints in (6.15) as linear equalities in X and Y , relax the nonconvex relation between these variables into the inequality $Y \succeq XX^T$ and apply Schur's complement to obtain an SDP. A sparsity-exploiting version of this approach is described by Kim et al. (2009, Section 3.3). Another option (Nie, 2009) is to formulate the search for the sensor locations as an unconstrained polynomial optimization problem,

$$\min_{x_1, \dots, x_n} \sum_{(i,j) \in \mathcal{E}_a} (\|a_i - x_j\|^2 - f_{ij}^2)^2 + \sum_{(i,j) \in \mathcal{E}_x} (\|x_i - x_j\|^2 - d_{ij}^2)^2.$$

The polynomial objective is term-sparse when the coupling set \mathcal{E}_x contains only a small subset of all pairs (i, j) (precisely, it is correlatively sparse; see Section 4.2 for definitions of these concepts). Therefore, the sparse SOS techniques outlined in Section 4.2 can be applied to solve the problem efficiently. The interested reader is referred to Nie (2009) for experiment results. ■

6.3. Machine learning: Verification of neural networks

Neural networks are one of the fundamental building blocks of modern machine-learning methods. For safety-critical applications, it is essential to ensure that they are provably robust to input perturbations. Given a neural network $f(x_0) : \mathbb{R}^d \rightarrow \mathbb{R}^m$, a nominal input $\bar{x} \in \mathbb{R}^d$, a linear function $\phi : \mathbb{R}^m \rightarrow \mathbb{R}$ on the network's output, and a perturbation radius $\epsilon \in \mathbb{R}$, the network verification problem (Raghunathan et al., 2018; Salman, Yang, Zhang, Hsieh, & Zhang, 2019; Tjandraatmadja, Anderson, Huchette, Ma, Patel, & Vielma, 2020) asks to either verify that

$$\phi(f(x_0)) > 0 \quad \forall x_0 : \|x_0 - \bar{x}\|_\infty \leq \epsilon, \quad (6.16)$$

or to identify at least one counterexample to this relation.

Consider an L -layer feedforward neural network where

$$f(x_0) = W_L x_L + b_L,$$

$$x_{i+1} = \text{ReLU}(W_i x_i + b_i), \quad i = 0, \dots, L-1,$$

where $W_i \in \mathbb{R}^{n_{i+1} \times n_i}$ and $b_i \in \mathbb{R}^{n_{i+1}}$ are the network weights and biases, respectively, and the so-called Rectified Linear Unit (ReLU) activation




Fig. 6.3. Abstraction of (a) a 4-layer neural network into (b) a chordal chain graph with four vertices and maximal cliques $\{i-1, i\}$, $\{i, i+1\}$ and $\{i+1, i+2\}$.




Fig. 6.4. CPU time (seconds) required to solve SDP relaxations of the neural network verification problem (6.17) for image verification, with and without sparsity exploitation. The SDP solver was MOSEK, and sparsity was exploited using SparseCoLO. The neural network, with $L = 2$ layers and $n_i = 64$ neurons per layer, was trained for image classification on the MNIST dataset.

function $\text{ReLU} : \mathbb{R}^k \rightarrow \mathbb{R}^k$ is the element-wise positive part of its argument, $\text{ReLU}(z) = [\max(z_i, 0)]_{i=1}^k$. Condition (6.16) can be decided by solving the optimization problem

$$\gamma^* := \min_{x_0, \dots, x_L} c^T x_L + c_0$$

$$\text{subject to} \quad x_{i+1} = \text{ReLU}(W_i x_i + b_i), \quad i \in [L], \quad (6.17a)$$

$$\|x_0 - \bar{x}\|_\infty \leq \epsilon, \quad (6.17b)$$

where $[L] := \{0, 1, \dots, L-1\}$ and c, c_0 are problem data related to the linear function $\phi(\cdot)$. If $\gamma^* > 0$, then (6.16) holds, otherwise counterexamples can be found.

Since the action of the ReLU function can be described by quadratic constraints,

$$y = \text{ReLU}(z) \iff y \geq z, y \geq 0, y(y-z) = 0,$$

problem (6.17) can be reformulated into a QCQP with variable $x = [x_0^T, x_1^T, \dots, x_L^T]^T$ (Raghunathan et al., 2018), and subsequently relaxed into an SDP as described in Section 6.2 above. If the optimal value of this SDP is positive, the network is verified; otherwise, nothing can be said.

Since the constraints (6.17a) have a very natural cascading structure, the interaction among variables x_0, \dots, x_L can be modeled by a line graph with maximal cliques $C_i = \{i, i+1\}$ for $i = 0, \dots, L-1$ (see Fig. 6.3 for illustration with $L = 4$). The SDP relaxation of (6.17) inherits this cascading structure, in addition to any sparsity coming from the structure of the weight matrices W_i . The chordal decomposition techniques described in Section 3 can therefore be applied to solve it efficiently. This idea has been recently validated by Batten et al. (2021), who considered robustness verification in the context

of image classifiers. For instance, the results reproduced in Fig. 6.4 for a neural network with $L = 2$ layers and $n_i = 64$ neurons per layer show that exploiting sparsity reduced by two orders of magnitude the CPU time required to verify the robustness of an image classifier on the MNIST dataset. Similar results were obtained by Newton and Papachristodoulou (2021), and interested readers are invited to consult Table 2 for references to more machine learning applications where sparsity exploitation can dramatically reduce computational complexity.

7. Conclusion and outlook

In this paper, we reviewed theory and applications of decomposition methods for large-scale semidefinite and polynomial optimization. Specifically, we presented classical chordal decomposition results for sparse positive semidefinite matrices (cf. Theorems 2.1–2.4) and we discussed how they can be exploited to implement efficient first- and second-order algorithms for SDPs (Section 3). We then showed how matrix decomposition (primarily, but not necessarily, chordal) can be leveraged to exploit *term sparsity* and *structural sparsity* in large-scale polynomial optimization (Section 4). In particular, we demonstrated that many sparsity-exploiting techniques for polynomial inequalities—including the well-known correlatively sparse SOS representations and the recent TSSOS, CS-TSSOS and chordal-TSSOS hierarchies—are based on the general matrix decomposition strategy outlined in Section 4.2.1. We also discussed how the classical chordal decomposition theorem (Theorem 2.1) can be generalized in different ways to obtain SOS chordal decomposition theorems for sparse polynomial matrices (cf. Theorems 4.4–4.6 and further results by Zheng and Fantuzzi (2020)). In Section 5, we reviewed factor-width decompositions for SDPs with dense semidefinite constraints, to which chordal decomposition cannot be applied. Finally, in Section 6 we demonstrated how some of these techniques can be used to reduce the computational complexity of SDPs and polynomial optimization problems encountered in a number of control and machine learning applications. References to these and other applications are summarized in Table 2.

Despite the considerable progress made in recent years, numerical methods for semidefinite and polynomial optimization are still far from being mature. The most pressing open challenge, in our opinion, lies in bridging the gap between the size of SDPs that can currently be solved with tractable computational resources, and the size of the SDPs that arise from complex control applications. Indeed, the state-of-the-art decomposition techniques reviewed in this article are often still not enough to enable the use of semidefinite programming to analyze and control large-scale nonlinear systems. The same is true for control problems with systems of smaller size, but which require real-time computations.

Achieving significant progress is likely to require theoretical extensions of the decomposition approaches we have discussed, as well as the development of efficient software that can effectively exploit modern multi-core and distributed-memory computer architectures. We conclude this article by outlining some possible research directions that may bear fruit in the near future.

Combining matrix decomposition with other structures

SDPs encountered in applications often have structural properties beyond sparsity, which can also be leveraged to reduce computational complexity; examples are symmetries, the existence of low-rank solutions, and low-rank data matrices (De Klerk, 2010; Gatermann & Parrilo, 2004; Majumdar et al., 2020). It is natural to try and combine the exploitation of such additional structure with matrix decomposition, but, to the best of our knowledge, a unified and theoretically robust framework to do so is yet to be developed. Particular questions to be answered in this context include whether there exist symmetry reduction techniques that preserve (or even promote) sparsity in SDPs,

and whether low-rank positive semidefinite completions (Dancis, 1992, Theorem 1.5) can be exploited in SDPs with aggregate sparsity and low-rank optimal solutions (see Jiang, 2017 and Miller, Zheng, Roig-Solvas et al., 2019 for some results in this direction).

In addition, although we have presented chordal and factor-width decompositions separately, they can be combined if either one, applied in isolation, does not reduce the complexity of a large-scale SDP enough. A relatively straightforward approach (Miller, Zheng, Sznajder et al., 2019) is to first apply the standard chordal decomposition, and then enforce the positive semidefinite constraints associated to large maximal cliques using factor-width approximations. This idea can be taken forward in various directions; for instance, one could use block-chordal and block-factor-width decompositions, or extend ideas by Garstka, Cannon, and Goulart (2020) to formulate adaptive strategies wherein cliques are either combined or factor-width decomposed, depending on their relative sizes and on the available computational resources. Both ideas remain largely unexplored, and further work is required to determine if they can be brought to bear on real-life control problems.

Tailored hierarchies for sparse polynomial optimization

Almost all existing methods for exploiting term sparsity in polynomial optimization rely on the general matrix decomposition approach presented in Section 4.2.1, where the Gram matrix associated with SOS certificates of nonnegativity is decomposed according to the maximal cliques of a sparsity graph to be prescribed *a priori*. While the correlatively sparse, TSSOS, and related hierarchies described in Section 4.2 give useful general strategies to select this sparsity graph, there is ample scope for tailoring the graph structure in particular control applications. It is not unreasonable to expect that problem-specific choices, motivated for example by physical intuition on the dynamical system one is trying to analyze or control, may bring significant further gains. However, it remains to be seen whether this expectation can be met in practice. Better integration between the development of optimization tools and application-related modeling, discussed further below, seems key to achieving progress in this direction.

Decomposition and completion of polynomial matrices

The exploitation of sparsity for polynomial matrix inequalities can be improved in various directions, reducing computational complexity beyond what can be achieved using only the SOS chordal decomposition results summarized in Section 4.3. For instance, those results can be combined in a natural way with techniques to leverage term-sparsity in scalar polynomial inequalities. Indeed, when a polynomial matrix inequality $P(x) \geq 0$ is “scalarized” into a nonnegativity condition for the polynomial $p(x, y) = y^T P(x) y$, the structural sparsity of P translates into correlative sparsity of p with respect to y . The matrix decomposition results of Section 4.3 have equivalent statement at the scalar level (Zheng & Fantuzzi, 2020, Section 4) that can be used to refine or extend term-sparse SOS decomposition hierarchies for polynomials. The latter, in turn, can be used to efficiently handle (scalarized) polynomial matrix inequalities.

It would also be interesting to establish SOS completion results for sparse polynomial matrices, in the spirit of Theorem 2.2. Preliminary results in this direction exist (Zheng, Fantuzzi, & Papachristodoulou, 2018a), but are far from complete. Extension of the results in this reference will contribute to building a comprehensive theory for SOS chordal decomposition and completion of polynomial matrices, which can be used to build tractable SDP approximations of large-scale optimization problems with sparse polynomial matrix inequalities.

To chordality and beyond

Exploiting sparsity in semidefinite and polynomial optimization without modifying the problem usually requires chordality (cf. [Theorems 2.1–2.4](#) for SDPs, and [Theorems 4.3, 4.5 and 4.6](#) for polynomial optimization). Enforcing chordality with traditional chordal extension strategies, even if approximately minimal, may lead to graphs with unacceptably large maximal cliques. The largest maximal clique size plays a major role in determining the computational complexity of a decomposed SDP (or SDP relaxation of a polynomial optimization problem). Therefore, systematic techniques to produce chordal extensions that approximately minimize the largest maximal cliques size would be very valuable.

If good chordal extensions prove hard to find, a compelling alternative is to sacrifice chordality and use nonchordal graphs with small cliques that can be determined analytically. This was done, for instance, by [Nie and Demmel \(2009\)](#) and [Kočvara \(2020\)](#). While clique decompositions of matrix inequalities based on nonchordal graphs are conservative in general, it may still be possible to identify classes of matrices for which the equivalence between the original and decomposed inequalities can be guaranteed. For example, sparse (scaled)-diagonally dominant matrices always admit a clique decomposition, even when their sparsity graph is not chordal ([Miller, Zheng, Sznajder et al., 2019](#), Proposition 1). The same is true for certain positive semidefinite matrices whose sparsity pattern can be extended to be of a “block-arrow” type ([Kočvara, 2020](#)). Necessary and sufficient cycle conditions for positive semidefinite completion problem with nonchordal sparsity graphs were investigated by [Barrett, Johnson, and Loewy \(1996\)](#). Extensions of these results, even if limited to particular application domains, are likely to enable considerable progress in the solution of large-scale SDPs with nonchordal sparsity.

Efficient software for modern computers

Reliable and user-friendly implementations of the cutting-edge decomposition techniques for SDPs and polynomial optimization problems reviewed in this paper are, in our opinion, just as important as further theoretical advances. Most of the available open-source packages mentioned in [Sections 3.5 and 4.5](#) have not yet reached the level of maturity required to solve robustly a wide range of SDPs or polynomial optimization problems arising from real-life applications. Moreover, many of the commonly-used optimization modeling environments on which these packages rely are by now over a decade old, and often cannot handle extremely large problems of industrial relevance efficiently.

The lack of very-high-performance software currently limits the scale of problems that can be solved without ad-hoc implementations. Since such implementations require considerable expertise in large-scale optimization, the deployment of SDP-based frameworks for system analysis and control to real-world problems is currently hindered. We expect that improvements in software reliability, efficiency, user-friendliness, and the ability to leverage modern multi-processor and/or distributed computing platforms will considerably increase the practical impact of decomposition methods for SDPs, bringing great benefit to the community of application-oriented researchers.

Blending application-driven modeling with optimization

The decomposition techniques reviewed in [Sections 3–5](#) apply to generic standard-form SDPs and polynomial optimization problems, irrespective of the context in which they arise. In control-related application, however, SDPs and polynomial optimization problems often come from modeling or relaxation frameworks for the study of dynamical systems, the details of which strongly affect the structure of the eventual optimization problem. Bridging the existing gaps between

application-driven modeling and the development of large-scale optimization algorithm promises to enable significant progress in the study of linear and nonlinear systems. On the one hand, it may be possible to implement tailored SDP solvers that target special structures arising in particular applications. On the other hand, given a particular control or analysis task, one should attempt to formulate modeling approaches that lead to optimization problems with a “computationally friendly” structure. For example, when studying fluid flows using semidefinite programming (see, e.g., [Fantuzzi et al., 2018](#) and [Arslan et al., 2021](#)), a smart discretization of the flow field leads to SDPs with chordal aggregate sparsity that can be solved in minutes even though their linear matrix inequalities have more than 10 000 rows/columns. Similarly, using structured Lyapunov (or Lyapunov-like) functions as explained in [Section 6.1](#) can lead to structured SDPs, enabling the analysis of increasingly large systems in fields such as robotics, smart energy grid, and autonomous transportation.

Of course, the design of analysis and control frameworks that combine system-level modeling with algorithmic considerations will present a number of challenges. Resolving these challenges, however, promises to remove long-standing barriers to the study of complex systems, especially nonlinear ones. Success seems likely to require a collaborative effort between researchers working in different areas and an increasing awareness of outstanding problems in particular application domains, as well as of state-of-the-art tools for large-scale optimization. We hope that the present review of decomposition methods for semidefinite and polynomial optimization takes a step in the right direction and can inspire new discoveries in the near future.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Y.Z. was supported in part by Clarendon Scholarship. G.F. gratefully acknowledges funding from an Imperial College Research Fellowship. A.P. was supported in part by the Engineering and Physical Sciences Research Council (EPSRC) under project EP/M002454/1.

Appendix A. Cholesky factorization with no fill-in

The no fill-in property of the Cholesky factorization for positive definite matrices with chordal sparsity is one of the most important results for sparsity exploitation in matrix calculations. For instance, it enables a simple proof of [Theorem 2.1](#) and efficient computations involving barrier functions for sparse matrix cones (cf. [Section 3.4.2](#)). To formally introduce this no fill-in property, we first define the notions of simplicial vertices and perfect elimination ordering for graphs.

Definition A.1. A vertex v in a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is called *simplicial* if all its neighbors are connected to each other.

Definition A.2. An ordering $\sigma = \{v_1, \dots, v_n\}$ of the vertices in a graph \mathcal{G} is a *perfect elimination ordering* if each v_i , $i = 1, \dots, n$, is a simplicial vertex in the subgraph induced by the vertices $\{v_i, v_{i+1}, \dots, v_n\}$.

For example, vertices 2, 4, 6 are simplicial for the graph in [Fig. A.1\(a\)](#), and the ordering $\sigma = \{2, 4, 6, 1, 3, 5\}$ is a perfect elimination ordering. A graph \mathcal{G} is chordal if and only if it has at least one perfect elimination ordering ([Vandenberghe & Andersen, 2015](#), Theorem 4.1). The maximal cardinality search ([Algorithm 1](#)) either returns one of the perfect elimination orderings or certifies that none exists in $\mathcal{O}(|\mathcal{V}| + |\mathcal{E}|)$ time ([Tarjan & Yannakakis, 1984](#)).

Algorithm 1 Maximal cardinality search

Input: A graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$
Output: An elimination ordering α of \mathcal{G}

```

for all vertices  $v$  in  $\mathcal{G}$  do
     $w(v) = 0$ .
end for
for  $i = n$  to 1 do
    pick an unnumbered vertex  $v$  with maximum weight in  $w$ ;
    set  $\alpha(v) = i$ ;
    for all unnumbered vertices  $u$  adjacent to  $v$  do
         $w(u) \leftarrow w(u) + 1$ ;
    end for
end for

```




Fig. A.1. Chordal graph decomposition: (a) a chordal graph with six nodes; (b) maximal cliques; (c) a clique tree that satisfies the clique intersection property.




Fig. A.2. (a) A symbolic 6×6 sparse positive definite matrix Z with chordal sparsity graph shown in Fig. A.1(a). (b) Sparsity pattern of $P_\sigma Z P_\sigma^T$ for the perfect elimination ordering $\sigma = \{2, 4, 6, 1, 3, 5\}$. (c) Cholesky factor of Z ; the entries marked by \times denote nonzero fill-ins. (d) Cholesky factor of $P_\sigma Z P_\sigma^T$ with no fill-in.

Now, given a positive definite matrix $Z \in \mathbb{S}_+^n(\mathcal{E}, 0)$ with a chordal sparsity pattern \mathcal{E} , we have a sparse Cholesky factorization with zero fill-in (Rose, 1970), (Vandenberghe & Andersen, 2015, Theorem 9.1):

$$P_\sigma Z P_\sigma^T = LL^T, \quad P_\sigma^T(L + L^T)P_\sigma \in \mathbb{S}^n(\mathcal{E}, 0), \quad (\text{A.1})$$

where P_σ is a permutation matrix corresponding to the perfect elimination ordering σ and L is a lower-triangular matrix. This can be proven using an elimination process according to the perfect elimination ordering σ ; see Vandenberghe and Andersen (2015, Chapter 9.1) and Kakimura (2010) for details. Fig. A.2 illustrates the process of sparse Cholesky factorization for a 6×6 positive definite matrix with chordal sparsity graph shown in Fig. A.1(a).

Appendix B. A proof of Theorem 2.1

The sparse Cholesky factorization (A.1) with zero fill-in allows for a simple proof of Theorem 2.1. For simplicity, but without loss of generality, assume that the matrix Z has already been permuted in such a way that $\sigma = \{1, 2, \dots, n\}$ is a perfect elimination ordering, so $P_\sigma = I$ in (A.1). We denote the columns of L by l_1, l_2, \dots, l_n , and write

$$Z = LL^T = \sum_{i=1}^n l_i l_i^T.$$

Since $L + L^T$ has the same sparsity pattern \mathcal{E} as Z , the non-zero elements of each column vector l_i must be indexed by a maximal clique C_{h_i} for some $h_i \in \{1, \dots, t\}$. Thus, the non-zero elements of l_i can be extracted through multiplication by the matrix E_{C_i} , and we have

$$l_i = E_{C_{h_i}}^T E_{C_{h_i}} l_i \Rightarrow l_i l_i^T = E_{C_{h_i}}^T \underbrace{(E_{C_{h_i}} l_i l_i^T E_{C_{h_i}}^T)}_{Q_i} E_{C_{h_i}}.$$

Now, let $J_k = \{i : h_i = k\}$ be the set of column indices i such that column i is indexed by clique C_k . These index sets are disjoint and $\cup_k J_k = \{1, \dots, n\}$, so we obtain

$$\begin{aligned} Z = LL^T &= \sum_{i=1}^n E_{C_{h_i}}^T Q_i E_{C_{h_i}} \\ &= \sum_{k=1}^t \sum_{i \in J_k} E_{C_k}^T Q_i E_{C_k} \\ &= \sum_{k=1}^t E_{C_k}^T \left(\sum_{i \in J_k} Q_i \right) E_{C_k}. \end{aligned}$$

Introducing the matrices $Z_k = \sum_{i \in J_k} Q_i$, each of which is in $\mathbb{S}_+^{|C_k|}$ by construction, we obtain the decomposition (2.3) in Theorem 2.1.

Appendix C. Some properties of maximal cliques

A connected chordal graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with n vertices has at most $n - 1$ maximal cliques that can be identified in linear time—more precisely, with a complexity of $\mathcal{O}(|\mathcal{V}| + |\mathcal{E}|)$ (Berry et al., 2004; Tarjan & Yannakakis, 1984). Algorithm 2 is a simple strategy with such a complexity, and it finds all maximal cliques based on a perfect elimination ordering. For example, the chordal graph in Fig. A.1(a) has the perfect elimination ordering $\sigma = \{2, 4, 6, 1, 3, 5\}$, and Algorithm 2 constructs the sets

$$\begin{aligned} C_1 &= \{2, 1, 3\}, & C_2 &= \{4, 3, 5\}, & C_3 &= \{6, 5, 1\}, \\ C_4 &= \{1, 3, 5\}, & C_5 &= \{3, 5\}, & C_6 &= \{5\}. \end{aligned}$$

The sets C_1, \dots, C_4 are maximal cliques, while C_5, C_6 are not because they are subsets of C_4 .

The maximal cliques of a chordal graph can be arranged in a so-called *clique tree*, that is, a graph $\mathcal{T}(\Gamma, \Xi)$ with the maximal cliques $\Gamma = \{C_1, \dots, C_t\}$ as its vertices and an edge set $\Xi \subseteq \Gamma \times \Gamma$. In particular, the clique tree can be chosen to satisfy the *clique intersection property*, meaning that $C_i \cap C_j \subseteq C_k$ if clique C_k lies on the path between cliques C_i and C_j in the tree and the intersection $C_i \cap C_j$ is nonempty (Blair & Peyton, 1993). For example, the clique tree in Fig. A.1(c) satisfies the clique intersection property.

The maximal cliques of a chordal graph play a central role in the sparse matrix decomposition results stated in Theorems 2.1–2.4. It is important to remember that these theorems require one to use *all* maximal cliques in the (chordal) sparsity graph of a matrix X , even when a subset of cliques already covers all nonzero entries of X . For

Algorithm 2 Maximal clique search**Input:** A chordal graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, and a perfect elimination ordering

$$\alpha = \{v_1, \dots, v_n\}$$

Output: All maximal cliques C_1, \dots, C_t Initialize $C_0 = \emptyset$;**for** $i = 1$ to n **do**

$$C_i = \{v_i\} \cup \{u \text{ adjacent to } v_i \text{ and behind } v_i \text{ in } \alpha\};$$

if C_i is not a subset of C_0 **then** C_i is a maximal clique;

$$C_0 = C_i;$$

end if**end for**

example, consider the indefinite matrix

$$X = \begin{bmatrix} 2 & 2 & 2 & 0 & 1 & 1 \\ 2 & 2 & 2 & 0 & 0 & 0 \\ 2 & 2 & 2 & 2 & 2 & 0 \\ 0 & 0 & 2 & 2 & 2 & 0 \\ 1 & 0 & 2 & 2 & 2 & 1 \\ 1 & 0 & 0 & 0 & 1 & 2 \end{bmatrix},$$

whose chordal sparsity graph is shown in Fig. A.1 and has the four maximal cliques identified above. Even though the maximal cliques C_1 , C_2 , and C_3 already cover all nonzero entries of the matrix, the maximal clique C_4 is necessary when applying Theorem 2.2 to check whether X admits a positive semidefinite completion. Indeed, observing that

$$E_{C_1} X E_{C_1}^T = E_{C_2} X E_{C_2}^T = \begin{bmatrix} 2 & 2 & 2 \\ 2 & 2 & 2 \\ 2 & 2 & 2 \end{bmatrix} \in \mathbb{S}_+^3$$

and

$$E_{C_3} X E_{C_3}^T = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix} \in \mathbb{S}_+^3,$$

is not sufficient to conclude $X \in \mathbb{S}_+^6(\mathcal{E}, ?)$ because the submatrix

$$E_{C_4} X E_{C_4}^T = \begin{bmatrix} 2 & 2 & 1 \\ 2 & 2 & 2 \\ 1 & 2 & 2 \end{bmatrix}$$

indexed by clique C_4 has one negative eigenvalue. Similarly, the matrix

$$Z = \begin{bmatrix} 4 & 2 & 2 & 0 & 1 & 1 \\ 2 & 4 & 2 & 0 & 0 & 0 \\ 2 & 2 & 3 & 2 & 2 & 0 \\ 0 & 0 & 2 & 3 & 2 & 0 \\ 1 & 0 & 2 & 2 & 3 & 1 \\ 1 & 0 & 0 & 0 & 1 & 3 \end{bmatrix}$$

is positive semidefinite and has the same sparsity graph as above, but it does not admit a decomposition

$$Z = \sum_{k=1}^3 E_{C_k}^T Z_k E_{C_k}, \quad Z_k \geq 0$$

that uses only cliques C_1 , C_2 , and C_3 ; the last maximal clique C_4 is necessary for Theorem 2.1 to apply. Indeed, any decomposition using only the first three maximal cliques requires

$$Z_1 = \begin{pmatrix} \alpha & 2 & 2 \\ 2 & 4 & 2 \\ 2 & 2 & \beta \end{pmatrix},$$

$$Z_2 = \begin{pmatrix} 3-\beta & 2 & 2 \\ 2 & 3 & 2 \\ 2 & 2 & \gamma \end{pmatrix},$$

$$Z_3 = \begin{pmatrix} 4-\alpha & 1 & 1 \\ 1 & 3-\gamma & 1 \\ 1 & 1 & 3 \end{pmatrix},$$

where α , β and γ must be selected to make these three matrices positive semidefinite. For this, it is necessary that the diagonal elements and all 2×2 principal minors of Z_1 , Z_2 and Z_3 be nonnegative; in particular,

$$\alpha, \beta, \gamma \geq 0, \quad 3-\beta \geq 0, \quad 4-\alpha \geq 0, \quad (\text{C.1a})$$

$$4\alpha-4 \geq 0, \quad \alpha\beta-4 \geq 0, \quad (3-\beta)\gamma-4 \geq 0, \quad (\text{C.1b})$$

$$(4-\alpha)(3-\gamma)-1 \geq 0. \quad (\text{C.1c})$$

However, this set of inequalities is infeasible. Specifically, inequality (C.1c) can be rearranged to show that

$$\gamma \leq \frac{11-3\alpha}{4-\alpha}.$$

We must also have $3-4/\alpha \geq 3-\beta \geq 0$, so $\alpha \geq 4/3$, and therefore

$$\frac{4\alpha}{3\alpha-4} \leq \frac{4}{3-\beta} \leq \gamma \leq \frac{11-3\alpha}{4-\alpha}.$$

But this cannot be true because $4/3 \leq \alpha \leq 4$, so $\frac{4\alpha}{3\alpha-4} > \frac{11-3\alpha}{4-\alpha}$.

References

- Agler, J., Helton, W., McCullough, S., & Rodman, L. (1988). Positive semidefinite matrices with a given sparsity pattern. *Linear Algebra and its Applications*, 107, 101–149.
- Ahmadi, A. A., Dash, S., & Hall, G. (2017). Optimization over structured subsets of positive semidefinite matrices via column generation. *Discrete Optimization*, 24, 129–151.
- Ahmadi, A. A., & Gunluk, O. (2018). Robust-to-dynamics optimization. *arXiv:1805.03682*, [math.OC].
- Ahmadi, A. A., & Hall, G. (2017). Sum of squares basis pursuit with linear and second order cone programming. *Contemporary Mathematics*, 25–54.
- Ahmadi, A. A., Hall, G., Papachristodoulou, A., Saunderson, J., & Zheng, Y. (2017). Improving efficiency and scalability of sum of squares optimization: Recent advances and limitations. In *2017 IEEE 56th annual conference on decision and control (CDC)* (pp. 453–462). IEEE.
- Ahmadi, S. P., Hansson, A., & Pakazad, S. K. (2019). Efficient robust model predictive control using chordality. In *2019 18th European control conference (ECC)* (pp. 4270–4275). IEEE.
- Ahmadi, A. A., & Majumdar, A. (2019). DSOS And SDSOS optimization: more tractable alternatives to sum of squares and semidefinite optimization. *SIAM Journal on Applied Algebra and Geometry*, 3(2), 193–230.
- Ahmadi, M., Valmorbida, G., Gayme, D., & Papachristodoulou, A. (2019). A framework for input-output analysis of wall-bounded shear flows. *Journal of Fluid Mechanics*, 873, 742–785.
- Alizadeh, F., & Goldfarb, D. (2003). Second-order cone programming. *Mathematical Programming*, 95(1), 3–51.
- Andersen, M. S. (2011). *Chordal sparsity in interior-point methods for conic optimization*. (Ph.D. thesis), University of California, Los Angeles.
- Andersen, M. S., Dahl, J., & Vandenberghe, L. (2010). Implementation of nonsymmetric interior-point methods for linear optimization over sparse matrix cones. *Mathematical Programming Comput.*, 2(3–4), 167–201.
- Andersen, M. S., Dahl, J., & Vandenberghe, L. (2013). Logarithmic barriers for sparse matrix cones. *Optimization Methods & Software*, 28(3), 396–423.
- Andersen, M. S., Hansson, A., & Vandenberghe, L. (2014). Reduced-complexity semidefinite relaxations of optimal power flow problems. *IEEE Transactions on Power Systems*, 29(4), 1855–1863.
- Andersen, M. S., Pakazad, S. K., Hansson, A., & Rantzer, A. (2014). Robust stability analysis of sparsely interconnected uncertain systems. *IEEE Transactions on Automatic Control*, 59(8), 2151–2156.
- Andersen, M. S., & Vandenberghe, L. (2010). Support vector machine training using matrix completion techniques. *Technical report*, Los Angeles: University of California, 2010. 6.2.
- Andersen, M., & Vandenberghe, L. (2014). SMCP: Python extension for sparse matrix cone programs, version 0.4. <https://github.com/cvxopt/smcpl>.

- Andersen, M., & Vandenbergh, L. (2015). *Chompack: a python package for chordal matrix computations*. Version.
- Andersen, M. S., Vandenbergh, L., & Dahl, J. (2010). Linear matrix inequalities with chordal sparsity patterns and applications to robust quadratic optimization. In *2010 IEEE international symposium on computer-aided control system design* (pp. 7–12). IEEE.
- Anderson, J., & Papachristodoulou, A. (2011). A decomposition technique for non-linear dynamical system analysis. *IEEE Transactions on Automatic Control*, 57(6), 1516–1521.
- Anderson, J., & Papachristodoulou, A. (2015). Advances in computational Lyapunov analysis using sum-of-squares programming. *Discrete and Continuous Dynamical Systems. Series B*, 20, 2361.
- Arsalan, A., Fantuzzi, G., Craske, J., & Wynn, A. (2021). Bounds on heat transport for convection driven by internal heating. *Journal of Fluid Mechanics*, 919, A15.
- Astrom, K. J., & Kumar, P. (2014). Control: A perspective. *Automatica*, 50(1), 3–43.
- Aylward, E. M., Itani, S. M., & Parrilo, P. A. (2007). Explicit SOS decompositions of univariate polynomial matrices and the Kalman-yakubovich-popov lemma. In *Proceedings of the 46 th IEEE conference on decision and control* (pp. 5660–5665). IEEE.
- Bai, X., Wei, H., Fujisawa, K., & Wang, Y. (2008). Semidefinite programming for optimal power flow problems. *International Journal of Electrical Power & Energy Systems*, 30(6–7), 383–392.
- Balante-Lugoan, R., Bonami, P., Misener, R., & Tramontani, A. (2019). Scoring positive semidefinite cutting planes for quadratic optimization via trained neural networks. http://www.optimization-online.org/DB_HTML/2018/11/6943.html.
- Banjac, G., Goulart, P., Stellato, B., & Boyd, S. (2019). Infeasibility detection in the alternating direction method of multipliers for convex optimization. *Journal of Optimization Theory and Applications*, 183(2), 490–519.
- Barrett, W. W., Johnson, C. R., & Loewy, R. (1996). vol. 584, *The Real positive definite completion problem: Cycle completability*. American Mathematical Soc..
- Barvinok, A. I. (1995). Problems of distance geometry and convex properties of quadratic maps. *Discrete & Computational Geometry*, 13(2), 189–202.
- Batten, B., Kouvaros, P., Lomuscio, A., & Zheng, Y. (2021). Efficient neural network verification via layer-based semidefinite relaxations and linear cuts. In Z.-H. Zhou (Ed.), *Proceedings of the thirtieth international joint conference on artificial intelligence, IJCAI-21* (pp. 2184–2190). International Joint Conferences on Artificial Intelligence Organization.
- Beck, A. (2017). *First-order methods in optimization*. SIAM.
- Ben-Tal, A., & Nemirovski, A. (1998). Robust convex optimization. *Mathematics of Operations Research*, 23(4), 769–805.
- Ben-Tal, A., & Nemirovski, A. (2001). *Lectures on modern convex optimization: analysis, algorithms, and engineering applications*. SIAM.
- Benson, S. J., Ye, Y., & Zhang, X. (2000). Solving large-scale sparse semidefinite programs for combinatorial optimization. *SIAM Journal on Optimization*, 10(2), 443–461.
- Berry, A., Blair, J. R., Heggernes, P., & Peyton, B. W. (2004). Maximum cardinality search for computing minimal triangulations of graphs. *Algorithmica*, 39(4), 287–298.
- Blair, J. R., & Peyton, B. (1993). An introduction to chordal graphs and clique trees. In *Graph theory and sparse matrix computation* (pp. 1–29). Springer.
- Blekherman, G., Dey, S. S., Molinaro, M., & Sun, S. (2020). Sparse PSD approximation of the PSD cone. *arXiv preprint arXiv:2002.02988*.
- Blekherman, G., Parrilo, P. A., & Thomas, R. R. (2012). *Semidefinite optimization and convex algebraic geometry*. SIAM.
- Boman, E. G., Chen, D., Parekh, O., & Toledo, S. (2005). On factor width and symmetric H-matrices. *Linear Algebra and its Applications*, 405, 239–248.
- Borchers, B. (1999). SDPLIB 1.2, a library of semidefinite programming test problems. *Optimization Methods & Software*, 11(1–4), 683–690.
- Boumal, N., Voroninski, V., & Bandeira, A. S. (2020). Deterministic guarantees for Burer-Monteiro factorizations of smooth semidefinite programs. *Communications on Pure and Applied Mathematics*, 73(3), 581–608, URL <https://doi.org/10.1002/cpa.21830>.
- Boyd, S., Diaconis, P., & Xiao, L. (2004). Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4), 667–689.
- Boyd, S., El Ghaoui, L., Feron, E., & Balakrishnan, V. (1994). *Linear matrix inequalities in system and control theory*. Society for Industrial and Applied Mathematics.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1), 1–122.
- Boyd, S., & Yang, Q. (1989). Structured and simultaneous Lyapunov functions for system stability problems. *International Journal of Control*, 49(6), 2215–2240.
- Burer, S. (2003). Semidefinite programming in the space of partial positive semidefinite matrices. *SIAM Journal on Optimization*, 14(1), 139–172.
- Burer, S., & Choi, C. (2006). Computational enhancements in low-rank semidefinite programming. *Optimization Methods & Software*, 21(3), 493–512.
- Burer, S., & Monteiro, R. D. (2003). A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2), 329–357.
- Burer, S., & Monteiro, R. D. (2005). Local minima and convergence in low-rank semidefinite programming. *Mathematical Programming*, 103(3), 427–444.
- Burer, S., Monteiro, R. D. C., & Zhang, Y. (2002). Solving a class of semidefinite programs via nonlinear programming. *Mathematical Programming, Series A*, 93(1), 97–122.
- Carlson, D., Hershkowitz, D., & Shasha, D. (1992). Block diagonal semistability factors and Lyapunov semistability of block triangular matrices. *Linear Algebra and its Applications*, 172, 1–25.
- Chen, T., Lasserre, J.-B., Magron, V., & Pauwels, E. (2020). Semialgebraic optimization for Lipschitz constants of ReLU networks. *arXiv:2002.03657*.
- Chen, H., Liu, H.-T. D., Jacobson, A., & Levin, D. I. (2020). Chordal decomposition for spectral coarsening. *arXiv preprint arXiv:2009.02294*.
- Chesi, G. (2011). vol. 415, *Domain of attraction: Analysis and control via SOS programming*. Springer Science & Business Media.
- Cifuentes, D., & Parrilo, P. A. (2016). Exploiting chordal structure in polynomial ideals: A grobner bases approach. *SIAM Journal of Discrete Mathematics*, 30(3), 1534–1570.
- Cifuentes, D., & Parrilo, P. A. (2017). Chordal networks of polynomial ideals. *SIAM Journal on Applied Algebra and Geometry*, 1(1), 73–110.
- Coey, C., Kapelevich, L., & Vielma, J. P. (2020). Towards practical generic conic optimization. *arXiv preprint arXiv:2005.01136*.
- Dahl, J., Vandenbergh, L., & Roychowdhury, V. (2008). Covariance selection for nonchordal graphs via chordal embedding. *Optimization Methods & Software*, 23(4), 501–520.
- Dall'Anese, E., Zhu, H., & Giannakis, G. B. (2013). Distributed optimal power flow for smart microgrids. *IEEE Transactions on Smart Grid*, 4(3), 1464–1475.
- Dancis, J. (1992). Positive semidefinite completions of partial hermitian matrices. *Linear Algebra and its Applications*, 175, 97–114.
- De Klerk, E. (2010). Exploiting special structure in semidefinite programming: A survey of theory and applications. *European Journal of Operational Research*, 201(1), 1–10.
- Deroo, F., Meinel, M., Ulbrich, M., & Hirche, S. (2014). Distributed control design with local model information and guaranteed stability. *IFAC Proceedings Volumes*, 47(3), 4010–4017.
- Deroo, F., Meinel, M., Ulbrich, M., & Hirche, S. (2015). Distributed stability tests for large-scale systems with limited model information. *IEEE Transactions on Control of Network Systems*, 2(3), 298–309.
- Dvijotham, K. D., Stanforth, R., Goyal, S., Qin, C., De, S., & Kohli, P. (2020). Efficient neural network verification with exactness characterization. In *Uncertainty in artificial intelligence* (pp. 497–507). PMLR.
- Eltved, A., Dahl, J., & Andersen, M. S. (2020). On the robustness and scalability of semidefinite relaxation for optimal power flow problems. *Optimization and Engineering*, 21(2), 375–392.
- Fantuzzi, G. (2020). Aeroimperial-YALMIP. <https://github.com/aeroimperial-optimization/aeroimperial-yalmip>.
- Fantuzzi, G., & Goluskin, D. (2020). Bounding extreme events in nonlinear dynamics using convex optimization. *SIAM Journal on Applied Dynamical Systems*, 19(3), 1823–1864.
- Fantuzzi, G., Goluskin, D., Huang, D., & Chernyshenko, S. I. (2016). Bounds for deterministic and stochastic dynamical systems using sum-of-squares optimization. *SIAM Journal on Applied Dynamical Systems*, 15(4), 1962–1988.
- Fantuzzi, G., Pershin, A., & Wynn, A. (2018). Bounds on heat transfer for Bénard-Marangoni convection at infinite Prandtl number. *Journal of Fluid Mechanics*, 837, 562–596.
- Fosson, S. M., & Abuabiah, M. (2019). Recovery of binary sparse signals from compressed linear measurements via polynomial optimization. *IEEE Signal Processing Letters*, 26(7), 1070–1074.
- Fujisawa, K., Fukuda, M., Kojima, M., Nakata, K., & Yamashita, M. (2004). *SDPA-C (semidefinite programming algorithm-completion method). user's manual-version 6-10*. Inst. of Technology.
- Fujisawa, K., Kim, S., Kojima, M., Okamoto, Y., & Yamashita, M. (2009). *User's manual for SparseCoLO: Conversion methods for sparse conic-form linear optimization problems, Research Report B-453: Tech. Rep.*, (pp. 152–8552). Dept. of Math. and Comp. Sci. Japan.
- Fukuda, M., Kojima, M., Murota, K., & Nakata, K. (2001). Exploiting sparsity in semidefinite programming via matrix completion I: General framework. *SIAM Journal on Optimization*, 11(3), 647–674.
- Furieri, L., Zheng, Y., Papachristodoulou, A., & Kamgarpour, M. (2019). On separable quadratic Lyapunov functions for convex design of distributed controllers. In *2019 18th European control conference (ECC)* (pp. 42–49). IEEE.
- Furieri, L., Zheng, Y., Papachristodoulou, A., & Kamgarpour, M. (2020). Sparsity invariance for convex design of distributed controllers. *IEEE Transactions on Control of Network Systems*.
- Gabay, D., & Mercier, B. (1976). A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & mathematics with applications*, 2(1), 17–40.
- Garstka, M., Cannon, M., & Goulart, P. (2020). A clique graph based merging strategy for decomposable SDPs. *IFAC-PapersOnLine*, 53(2), 7355–7361.
- Garstka, M., Cannon, M., & Goulart, P. (2021). COSMO: A conic operator splitting method for convex conic problems. *Journal of Optimization Theory and Applications*.

- Gatermann, K., & Parrilo, P. A. (2004). Symmetry groups, semidefinite programs, and sums of squares. *Journal of Pure and Applied Algebra*, 192(1–3), 95–128.
- Geromel, J. C., Bernussou, J., & Peres, P. L. D. (1994). Decentralized control through parameter space optimization. *Automatica*, 30(10), 1565–1578.
- Giulietti, F., Pollini, L., & Innocenti, M. (2000). Autonomous formation flight. *IEEE Control Systems Magazine*, 20(6), 34–44.
- Glowinski, R., & Marroco, A. (1975). Sur l'approximation, par éléments finis d'ordre un, et la résolution, par pénalisation-dualité d'une classe de problèmes de Dirichlet non linéaires. *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique Et Analyse Numérique*, 9(R2), 41–76.
- Goemans, M. X., & Williamson, D. P. (1995). Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6), 1115–1145.
- Goldfarb, D., & Iyengar, G. (2003). Robust convex quadratically constrained programs. *Mathematical Programming*, 97(3), 495–515.
- Golumbic, M. C. (2004). *Algorithmic Graph Theory and Perfect Graphs*. Elsevier.
- Goluskin, D. (2020). Bounding extrema over global attractors using polynomial optimisation. *Nonlinearity*, 33(9), 4878–4899.
- Griewank, A., & Toint, P. L. (1984). On the existence of convex decompositions of partially separable functions. *Mathematical Programming*, 28(1), 25–49.
- Grimm, D., Netzer, T., & Schweighofer, M. (2007). A note on the representation of positive polynomials with structured sparsity. *Arch. Math. (Basel)*, 89(5), 399–403.
- Grone, R., Johnson, C. R., Sá, E. M., & Wolkowicz, H. (1984). Positive definite completions of partial Hermitian matrices. *Linear Algebra and its Applications*, 58, 109–124.
- Han, W., & Tedrake, R. (2018). Convex optimization of nonlinear state feedback controllers for discrete-time polynomial systems via occupation measures. (pp. 1–8). arXiv:1803.09022 [math.OC].
- Hansson, D., & Pakazad, S. K. (2018). Exploiting chordality in optimization algorithms for model predictive control. In *Large-scale and distributed optimization* (pp. 11–32). Springer.
- Heinke, S., Schug, A.-K., & Werner, H. (2020). Distributed controller design for systems interconnected over chordal graphs. In *2020 American control conference (ACC)* (pp. 1569–1574). IEEE.
- Henrion, D., & Garulli, A. (2005). *Lecture notes in control and information sciences: vol. 312, Positive polynomials in control*. Springer-Verlag Berlin Heidelberg.
- Henrion, D., & Korda, M. (2014). Convex computation of the region of attraction of polynomial control systems. *IEEE Transactions on Automatic Control*, 59(2), 297–312.
- Henrion, D., & Lasserre, J.-B. (2006). Convergent relaxations of polynomial matrix inequalities and static output feedback. *IEEE Transactions on Automatic Control*, 51(2), 192–202.
- Henrion, D., & Lasserre, J.-B. (2011). Inner approximations for polynomial matrix inequalities and robust stability regions. *IEEE Transactions on Automatic Control*, 57(6), 1456–1467.
- Henrion, D., Lasserre, J.-B., & Löfberg, J. (2009). Gloptipoly 3: moments, optimization and semidefinite programming. *Optimization Methods & Software*, 24(4–5), 761–779.
- Hilbert, D. (1888). Ueber die darstellung definiter formen als summe von formenquadraten. *Mathematische Annalen*, 32(3), 342–350.
- Jabr, R. A. (2011). Exploiting sparsity in SDP relaxations of the OPF problem. *IEEE Transactions on Power Systems*, 27(2), 1138–1139.
- Jeyakumar, V., Kim, S., Lee, G., & Li, G. (2016). Semidefinite programming relaxation methods for global optimization problems with sparse polynomials and unbounded semialgebraic feasible sets. *Journal of Global Optimization*, 65(2), 175–190.
- Jiang, X. (2017). *Minimum rank positive semidefinite matrix completion with chordal sparsity pattern* (Master's thesis), UCLA.
- Jiang, X., & Vandenbergh, L. (2021). Bregman primal-dual first-order method and application to sparse semidefinite programming. http://www.optimization-online.org/DB_HTML/2020/03/7702.html.
- Jing, G., Wan, C., & Dai, R. (2019). Angle-based sensor network localization. arXiv preprint arXiv:1912.01665.
- Jones, M., & Peet, M. M. (2019). Using SOS and sublevel set volume minimization for estimation of forward reachable sets. *IFAC-PapersOnLine*, 52(16), 484–489.
- Kailath, T. (1980). *vol. 156, Linear systems*. Prentice-Hall Englewood Cliffs, NJ.
- Kakimura, N. (2010). A direct proof for the matrix decomposition of chordal-structured positive semidefinite matrices. *Linear Algebra and its Applications*, 433(4), 819–823.
- Kalbat, A., & Lavaei, J. (2015). A fast distributed algorithm for decomposable semidefinite programs. In *2015 54th IEEE conference on decision and control (CDC)* (pp. 1742–1749). IEEE.
- Karisch, S. E., & Rendl, F. (1998). Semidefinite programming and graph equipartition. *Topics in Semidefinite and Interior-Point Methods*, 18(77–95), 25.
- Kim, S., Kojima, M., Mevissen, M., & Yamashita, M. (2011). Exploiting sparsity in linear and nonlinear matrix inequalities via positive semidefinite matrix completion. *Mathematical Programming*, 129(1), 33–68.
- Kim, S., Kojima, M., & Waki, H. (2009). Exploiting sparsity in SDP relaxation for sensor network localization. *SIAM Journal on Optimization*, 20(1), 192–215.
- Klep, I., Magron, V., & Povh, J. (2019). Sparse noncommutative polynomial optimization. arXiv:1909.00569 [math.OC].
- Kočvara, M. (2020). Decomposition of arrow type positive semidefinite matrices with application to topology optimization. *Mathematical Programming*, 1–30.
- Kojima, M. (2003). *Sums of squares relaxations of polynomial semidefinite programs: Research reports on mathematical and computing sciences series B : Operations research B-397*, Tokyo Institute of Technology.
- Korda, M., Henrion, D., & Jones, C. N. (2013). Inner approximations of the region of attraction for polynomial dynamical systems. *IFAC Proceedings Volumes (IFAC-PapersOnline)*, 43(23), 534–539.
- Korda, M., Henrion, D., & Mezić, I. (2021). Convex computation of extremal invariant measures of nonlinear dynamical systems and Markov processes. *Journal of Nonlinear Science*, 31, 14(1–26).
- Kuntz, J., Ottobre, M., Stan, G.-B., & Barahona, M. (2016). Bounding stationary averages of polynomial diffusions via semidefinite programming. *SIAM Journal on Scientific Computing*, 38(6), A3891–A3920.
- Lam, A. Y., Zhang, B., & David, N. T. (2012). Distributed algorithms for optimal power flow problem. In *2012 IEEE 51st IEEE conference on decision and control (CDC)* (pp. 430–437). IEEE.
- Lasagna, D., Huang, D., Tutty, O. R., & Chernyshenko, S. I. (2016). Sum-of-squares approach to feedback control of laminar wake flows. *Journal of Fluid Mechanics*, 809, 628–663.
- Lasserre, J.-B. (2006). Convergent SDP-relaxations in polynomial optimization with sparsity. *SIAM Journal on Optimization*, 17(3), 822–843.
- Lasserre, J.-B. (2010). *Moments, positive polynomials and their applications*. Imperial College Press.
- Lasserre, J. B., Henrion, D., Prieur, C., & Trélat, E. (2008). Nonlinear optimal control via occupation measures and LMI-relaxations. *SIAM Journal on Control and Optimization*, 47(4), 1643–1666.
- Lasserre, J. B., Toh, K.-C., & Yang, S. (2017). A bounded degree SOS hierarchy for polynomial optimization. *EURO Journal on Computational Optimization*, 5(1–2), 87–117.
- Latorre, F., Rolland, P., & Cevher, V. (2020). Lipschitz constant estimation of neural networks via sparse polynomial optimization. arXiv preprint arXiv:2004.08688.
- Legat, B., Coey, C., Deits, R., Huchette, J., & Perry, A. (2017). Sum-of-squares optimization in julia. In *The first annual jump-dev workshop*.
- Li, H., Xia, B., Zhang, H., & Zheng, T. (2021). Choosing the variable ordering for cylindrical algebraic decomposition via exploiting chordal structure. arXiv preprint arXiv:2102.00823.
- Li, S. E., Zheng, Y., Li, K., Wu, Y., Hedrick, J. K., Gao, F., et al. (2017). Dynamical modeling and distributed control of connected and automated vehicles: Challenges and opportunities. *IEEE Intelligent Transportation Systems Magazine*, 9(3), 46–58.
- Liu, H.-T. D., Jacobson, A., & Ovsjanikov, M. (2019). Spectral coarsening of geometric operators. *ACM Transactions on Graphics*, 38(4), 1–13.
- Liu, Y., Ryu, E. K., & Yin, W. (2017). A new use of douglas-rachford splitting and ADMM for identifying infeasible, unbounded, and pathological conic programs. arXiv preprint arXiv:1706.02374.
- Löfberg, J. (2004). YALMIP: A toolbox for modeling and optimization in MATLAB. In *Proceedings of the IEEE International Symposium on Computer-Aided Control System Design* (pp. 284–289). IEEE.
- Löfberg, J. (2009a). Dualize it: software for automatic primal and dual conversions of conic programs. *Optimization Methods & Software*, 24(3), 313–325.
- Löfberg, J. (2009b). Pre-and post-processing sum-of-squares programs in practice. *IEEE Transactions on Automatic Control*, 54(5), 1007–1011.
- Lu, Z., Nemirovski, A., & Monteiro, R. D. (2007). Large-scale semidefinite programming via a saddle point mirror-prox algorithm. *Mathematical Programming*, 109(2), 211–237.
- Madani, R., Kalbat, A., & Lavaei, J. (2017). A low-complexity parallelizable numerical algorithm for sparse semidefinite programming. *IEEE Transactions on Control of Network Systems*, 5(4), 1898–1909.
- Madani, R., Sojoudi, S., Fazelnia, G., & Lavaei, J. (2017). Finding low-rank solutions of sparse linear matrix inequalities using convex optimization. *SIAM Journal on Optimization*, 27(2), 725–758.
- Magron, V., Garoche, P.-L., Henrion, D., & Thiriaux, X. (2019). Semidefinite approximations of reachable sets for discrete-time polynomial systems. *SIAM Journal on Control and Optimization*, 57(4), 2799–2820.
- Magron, V., & Wang, J. (2021). TSSOS: A Julia library to exploit sparsity for large-scale polynomial optimization. arXiv:2103.00915 [math.OC].
- Mai, N. H. A., Magron, V., & Lasserre, J.-B. (2020). A sparse version of Reznick's positivstellensatz. arXiv:2002.05101 [math.OC].
- Majumdar, A., Hall, G., & Ahmadi, A. A. (2020). Recent scalability improvements for semidefinite programming with applications in machine learning, control, and robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, 3, 331–360.
- Majumdar, A., Vasudevan, R., Tobenkin, M. M., & Tedrake, R. (2014). Convex optimization of nonlinear feedback controllers via occupation measures. *International Journal of Robotics Research*, 33(9), 1209–1230.
- Mason, R. (2015). *A chordal sparsity approach to scalable linear and nonlinear systems analysis* (Ph.D. thesis), University of Oxford.
- Mason, R. P., & Papachristodoulou, A. (2014). Chordal sparsity, decomposing SDPs and the Lyapunov equation. In *2014 American control conference* (pp. 531–537). IEEE.

- Mevisen, M. (2010). *Sparse semidefinite programming relaxations for large scale polynomial optimization and their applications to differential equations* (Ph.D. thesis), Tokyo Institute of Technology.
- Mevisen, M., Kojima, M., Nie, J., & Takayama, N. (2008). Solving partial differential equations via sparse SDP relaxations. *Pacific Journal of Optimization*, 4(2), 213–241.
- Mevisen, M., Lasserre, J. B., & Henrion, D. (2011). Moment and SDP relaxation techniques for smooth approximations of problems involving nonlinear differential equations. *IFAC Proceedings Volumes*, 44(1), 10887–10892.
- Mevisen, M., Yokoyama, K., & Takayama, N. (2009). Solutions of polynomial systems derived from the steady cavity flow problem. In *Proceedings of the 2009 international symposium on symbolic and algebraic computation* (pp. 255–262). Seoul, Republic of Korea: Association for Computing Machinery.
- Miller, J., Henrion, D., & Sznaier, M. (2021). Peak estimation recovery and safety analysis. *IEEE Control Systems Letters*, 5(6), 1982–1987.
- Miller, J., Zheng, Y., Roig-Solvas, B., Sznaier, M., & Papachristodoulou, A. (2019). Chordal decomposition in rank minimized semidefinite programs with applications to subspace clustering. In *2019 IEEE 58th conference on decision and control (CDC)* (pp. 4916–4921). IEEE.
- Miller, J., Zheng, Y., Sznaier, M., & Papachristodoulou, A. (2019). Decomposed structured subsets for semidefinite and sum-of-squares optimization. arXiv:1911.12859 [math.OC].
- Molzahn, D. K., & Hiskens, I. A. (2014). Sparsity-exploiting moment-based relaxations of the optimal power flow problem. *IEEE Transactions on Power Systems*, 30(6), 3168–3180.
- Molzahn, D. K., Holzer, J. T., Lesieutre, B. C., & DeMarco, C. L. (2013). Implementation of a large-scale optimal power flow solver based on semidefinite programming. *IEEE Transactions on Power Systems*, 28(4), 3987–3998.
- Mosek, A. (2015). *The MOSEK optimization toolbox for MATLAB manual*. Version.
- Motzkin, T. S. (1967). The arithmetic-geometric inequality. In *Inequalities (Proc. sympos. wright-patterson air force base, Ohio, 1965)* (pp. 205–224).
- Mou, C., Bai, Y., & Lai, J. (2021). Chordal graphs in triangular decomposition in top-down style. *Journal of Symbolic Computation*, 102, 108–131.
- Murray, R. M., Astrom, K. J., Boyd, S. P., Brockett, R. W., & Stein, G. (2003). Future directions in control in an information-rich world. *IEEE Control Systems Magazine*, 23(2), 20–33.
- Murty, K. G., & Kabadi, S. N. (1987). Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2), 117–129.
- Nakata, K., Fujisawa, K., Fukuda, M., Kojima, M., & Murota, K. (2003). Exploiting sparsity in semidefinite programming via matrix completion II: implementation and numerical results. *Mathematical Programming B*, 95(2), 303–327.
- Nemirovski, A. (2006). Advances in convex optimization: Conic programming. 1, In *International congress of mathematicians* (pp. 413–444).
- Nesterov, Y. (2003). vol. 87, *Introductory lectures on convex optimization: A basic course*. Springer Science & Business Media.
- Nesterov, Y. (2012). Towards non-symmetric conic optimization. *Optimization Methods & Software*, 27(4–5), 893–917.
- Nesterov, Y., & Nemirovski, A. (1994). *Interior-point polynomial algorithms in convex programming* (pp. IX–396). SIAM.
- Nesterov, Y., Wolkowicz, H., & Ye, Y. (2000). Semidefinite programming relaxations of nonconvex quadratic optimization. In *Handbook of semidefinite programming* (pp. 361–419). Springer.
- Newton, M., & Papachristodoulou, A. (2021). Exploiting sparsity for neural network verification. In *Proceedings of the 3rd Annual Learning for Dynamics and Control Conference* (pp. 715–727). PMLR.
- Nie, J. (2009). Sum of squares method for sensor network localization. *Computational Optimization and Applications*, 43(2), 151–179.
- Nie, J., & Demmel, J. (2009). Sparse SOS relaxations for minimizing functions that are summations of small polynomials. *SIAM Journal on Optimization*, 19(4), 1534–1558.
- O'Donoghue, B., Chu, E., Parikh, N., & Boyd, S. (2016). Conic optimization via operator splitting and homogeneous self-dual embedding. *Journal of Optimization Theory and Applications*, 169(3), 1042–1068.
- O'Donoghue, B., Chu, E., Parikh, N., & Boyd, S. (2019). SCS: Splitting conic solver, version 2.1.2. <https://github.com/cvxgrp/scs>.
- Pakazad, S. K., Hansson, A., Andersen, M. S., & Nielsen, I. (2017). Distributed primal-dual interior-point methods for solving tree-structured coupled convex problems using message-passing. *Optimization Methods & Software*, 32(3), 401–435.
- Pakazad, S. K., Hansson, A., Andersen, M. S., & Rantzer, A. (2017). Distributed semidefinite programming with application to large-scale system analysis. *IEEE Transactions on Automatic Control*, 63(4), 1045–1058.
- Papachristodoulou, A., & Prajna, S. (2005). A tutorial on sum of squares techniques for systems analysis. In *American control conference, 2005. Proceedings of the 2005* (pp. 2686–2700). IEEE.
- Park, J., & Boyd, S. (2017). General heuristics for nonconvex quadratically constrained quadratic programming. arXiv preprint arXiv:1703.07870.
- Parrilo, P. A. (2000). *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization* (Ph.D. thesis), California Institute of Technology.
- Parrilo, P. A. (2003). Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96(2), 293–320.
- Parrilo, P. A. (2013). Polynomial optimization, sums of squares and applications. In G. Blekherman, P. A. Parrilo, & R. R. Thomas (Eds.), *Semidefinite optimization and convex algebraic geometry* (1st ed.). (pp. 47–157). SIAM.
- Parrilo, P. A., & Lall, S. (2003). Semidefinite programming relaxations and algebraic optimization in control. *European Journal of Control*, 9(2–3), 307–321.
- Pataki, G. (1998). On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues. *Mathematics of Operations Research*, 23(2), 339–358.
- Peet, M. M., & Papachristodoulou, A. (2012). A converse sum of squares Lyapunov result with a degree bound. *IEEE Transactions on Automatic Control*, 57(9), 2281–2293.
- Peet, M. M., Papachristodoulou, A., & Lall, S. (2009). Positive forms and stability of linear time-delay systems. *SIAM Journal on Control and Optimization*, 47(6), 3237–3258.
- Permenter, F., Friberg, H. A., & Andersen, E. D. (2017). Solving conic optimization problems via self-dual embedding and facial reduction: a unified approach. *SIAM Journal on Optimization*, 27(3), 1257–1282.
- Permenter, F., & Parrilo, P. A. (2014a). Basis selection for SOS programs via facial reduction and polyhedral approximations. In *Proceedings of the 53rd IEEE conference on decision and control* (pp. 6615–6620). IEEE.
- Permenter, F., & Parrilo, P. A. (2014b). Partial facial reduction: simplified, equivalent SDPs via approximations of the PSD cone. *Mathematical Programming*, 1–54.
- Ploeg, J., Shukla, D. P., van de Wouw, N., & Nijmeijer, H. (2013). Controller synthesis for string stability of vehicle platoons. *IEEE Transactions on Intelligent Transportation Systems*, 15(2), 854–865.
- Prajna, S., Jadbabaie, A., & Pappas, G. J. (2007). A framework for worst-case and stochastic safety verification using barrier certificates. *IEEE Transactions on Automatic Control*, 52(8), 1415–1428.
- Prajna, S., Papachristodoulou, A., & Parrilo, P. A. (2002). Introducing SOSTOOLS: A general purpose sum of squares programming solver. vol. 1, In *Proceedings of the 41st IEEE conference on decision and control, 2002* (pp. 741–746). IEEE.
- Prajna, S., Papachristodoulou, A., & Wu, F. (2004). Nonlinear control synthesis by sum of squares optimization: A Lyapunov-based approach. vol. 1, In *2004 5th Asian control conference (IEEE Cat. No. 04EX904)* (pp. 157–165). IEEE.
- Putinar, M. (1993). Positive polynomials on compact semi-algebraic sets. *Indiana University Mathematics Journal*, 42(3), 969–984.
- Raghuathan, A. U., & Knyazev, A. V. (2016). Degeneracy in maximal clique decomposition for semidefinite programs. In *2016 American control conference (ACC)* (pp. 5605–5611). IEEE.
- Raghuathan, A., Steinhardt, J., & Liang, P. S. (2018). Semidefinite relaxations for certifying robustness to adversarial examples. In *Advances in neural information processing systems* (pp. 10877–10887).
- Rajamani, R. (2011). *Vehicle dynamics and control*. Springer Science & Business Media.
- Reznick, B. (1978). Extremal PSD forms with few terms. *Duke Mathematical Journal*, 45(2), 363–374.
- Reznick, B. (1995). Uniform denominators in Hilbert's seventeenth problem. *Mathematische Zeitschrift*, 220, 75–97.
- Riener, C., Theobald, T., Andr n, L. J., & Lasserre, J.-B. (2013). Exploiting symmetries in SDP-relaxations for polynomial optimization. *Mathematics of Operations Research*, 38(1), 122–141.
- Rivero, S., Sarzo, F., & Ferrari-Trecate, G. (2014). Plug-and-play voltage and frequency control of islanded microgrids with meshed topology. *IEEE Transactions on Smart Grid*, 6(3), 1176–1184.
- Rose, D. J. (1970). Triangulated graphs and the elimination process. *Journal of Mathematical Analysis and Applications*, 32(3), 597–609.
- Sadabadi, M. S., Shafee, Q., & Karimi, A. (2016). Plug-and-play voltage stabilization in inverter-interfaced microgrids via a robust control strategy. *IEEE Transactions on Control Systems Technology*, 25(3), 781–791.
- Salman, H., Yang, G., Zhang, H., Hsieh, C.-J., & Zhang, P. (2019). A convex relaxation barrier to tight robustness verification of neural networks. arXiv preprint arXiv:1902.08722.
- Scherer, C., & Hol, C. (2006). Matrix sum-of-squares relaxations for robust semi-definite programs. *Mathematical Programming*, 107, 189–211.
- Schlosser, C., & Korda, M. (2020). Sparse moment-sum-of-squares relaxations for nonlinear dynamical systems with guaranteed convergence. (pp. 1–34). arXiv:2012.05572 [math.OC].
- Schm dgen, K. (2009). Noncommutative real algebraic geometry some basic concepts and first ideas. In *Emerging applications of algebraic geometry* (pp. 325–350). Springer.
- Skajaa, A., & Ye, Y. (2015). A homogeneous interior-point algorithm for nonsymmetric convex conic optimization. *Mathematical Programming*, 150(2), 391–422.
- So, A. M.-C., & Ye, Y. (2007). Theory of semidefinite programming for sensor network localization. *Mathematical Programming*, 109(2), 367–384.
- Song, D., & Parrilo, P. A. (2021). On approximations of the PSD cone by a polynomial number of smaller-sized PSD cones. arXiv preprint arXiv:2105.02080.
- Sootla, A., Zheng, Y., & Papachristodoulou, A. (2017). Block-diagonal solutions to Lyapunov inequalities and generalisations of diagonal dominance. In *2017 IEEE 56th annual conference on decision and control (CDC)* (pp. 6561–6566). IEEE.

- Sootla, A., Zheng, Y., & Papachristodoulou, A. (2019). On the existence of block-diagonal solutions to Lyapunov and H_∞ Riccati inequalities. *IEEE Transactions on Automatic Control*, 65(7), 3170–3175.
- Sturm, J. F. (1999). Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods & Software*, 11(1–4), 625–653.
- Sun, Y. (2015). *Decomposition methods for semidefinite optimization* (Ph.D. thesis), UCLA.
- Sun, Y., Andersen, M. S., & Vandenberghe, L. (2014). Decomposition in conic optimization with partially separable structure. *SIAM Journal on Optimization*, 24(2), 873–897.
- Sun, D., Toh, K.-C., Yuan, Y., & Zhao, X.-Y. (2020). SDPNAL+: A matlab software for semidefinite programming with bound constraints (version 1.0). *Optimization Methods & Software*, 35(1), 87–115.
- Sun, Y., & Vandenberghe, L. (2015). Decomposition methods for sparse matrix nearness problems. *SIAM Journal on Matrix Analysis and Applications*, 36(4), 1691–1717.
- Tacchi, M., Cardozo, C., Henrion, D., & Lasserre, J.-B. (2019). Approximating regions of attraction of a sparse polynomial differential system. (pp. 1–15). arXiv:1911.09500 [math.OC].
- Tacchi, M., Weisser, T., Lasserre, J.-B., & Henrion, D. (2019). Exploiting sparsity for semi-algebraic set volume computation. (pp. 1–15). arXiv:1902.02976 [math.OC].
- Tarjan, R. E., & Yannakakis, M. (1984). Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM Journal on Computing*, 13(3), 566–579.
- Tjandraatmadja, C., Anderson, R., Huchette, J., Ma, W., Patel, K., & Vielma, J. P. (2020). The convex relaxation barrier, revisited: Tightened single-neuron relaxations for neural network verification. arXiv preprint arXiv:2006.14076.
- Tomita, E., Tanaka, A., & Takahashi, H. (2006). The worst-case time complexity for generating all maximal cliques and computational experiments. *Theoretical Computer Science*, 363(1), 28–42.
- Topcu, U., Packard, A. K., Seiler, P., & Balas, G. J. (2009). Robust region-of-attraction estimation. *IEEE Transactions on Automatic Control*, 55(1), 137–142.
- Tütüncü, R. H., Toh, K.-C., & Todd, M. J. (2003). Solving semidefinite-quadratic-linear programs using SDPT3. *Mathematical Programming*, 95(2), 189–217.
- Valmorbida, G., Ahmadi, M., & Papachristodoulou, A. (2015). Stability analysis for a class of partial differential equations via semidefinite programming. *IEEE Transactions on Automatic Control*, 61(6), 1649–1654.
- Valmorbida, G., & Anderson, J. (2017). Region of attraction estimation using invariant sets and rational Lyapunov functions. *Automatica*, 75, 37–45.
- Vandenberghe, L., & Andersen, M. S. (2015). Chordal graphs and semidefinite optimization. *Foundation and Trends in Optimization*, 1(4), 241–433.
- Vandenberghe, L., & Boyd, S. (1996). Semidefinite programming. *SIAM Review*, 38(1), 49–95.
- Waki, H., Kim, S., Kojima, M., & Muramatsu, M. (2006). Sums of squares and semidefinite program relaxations for polynomial optimization problems with structured sparsity. *SIAM Journal on Optimization*, 17(1), 218–242.
- Waki, H., Kim, S., Kojima, M., Muramatsu, M., & Sugimoto, H. (2008). Algorithm 883: Sparsepop—a sparse semidefinite programming relaxation of polynomial optimization problems. *ACM Transactions on Mathematical Software*, 35(2), 1–13.
- Waki, H., & Muramatsu, M. (2010). A facial reduction algorithm for finding sparse SOS representations. *Operations Research Letters*, 38(5), 361–365.
- Waldspurger, I., & Waters, A. (2020). Rank optimality for the Burer–Monteiro factorization. *SIAM Journal on Optimization*, 30(3), 2577–2602, URL <https://doi.org/10.1137/19M1255318>.
- Wang, J., Li, H., & Xia, B. (2019). A new sparse SOS decomposition algorithm based on term sparsity. In *Proceedings of the ACM International Symposium on Symbolic and Algebraic Computation* (pp. 347–354).
- Wang, J., Magron, V., & Lasserre, J.-B. (2021a). Chordal-TSSOS: a moment-SOS hierarchy that exploits term sparsity with chordal extension. *SIAM Journal on Optimization*, 31(1), 114–141.
- Wang, J., Magron, V., & Lasserre, J.-B. (2021b). TSSOS: A moment-SOS hierarchy that exploits term sparsity. *SIAM Journal on Optimization*, 31(1), 30–58.
- Wang, J., Magron, V., Lasserre, J.-B., & Mai, N. H. A. (2020). CS-TSSOS: Correlative and term sparsity for large-scale polynomial optimization. (pp. 1–27). arXiv:2005.02828 [math.OC].
- Wang, Y., Tanaka, A., & Yoshise, A. (2021). Polyhedral approximations of the semidefinite cone and their application. *Computational Optimization and Applications*, 78(3), 893–913.
- Wang, J., Zheng, Y., Chen, C., Xu, Q., & Li, K. (2020). Leading cruise control in mixed traffic flow: System modeling, controllability, and string stability. arXiv preprint arXiv:2012.04313.
- Weisser, T., Lasserre, J.-B., & Toh, K.-C. (2018). Sparse-BSOS: a bounded degree SOS hierarchy for large scale polynomial optimization with sparsity. *Mathematical Programming Computation*, 10(1), 1–32.
- Weisser, T., Legat, B., Coey, C., Kapelevich, L., & Vielma, J. P. (2019). Polynomial and moment optimization in Julia and JuMP. In *JuliaCon*. URL <https://pretalx.com/juliacon2019/talk/QZBKAU/>.
- Wen, Z., Goldfarb, D., & Yin, W. (2010). Alternating direction augmented Lagrangian methods for semidefinite programming. *Mathematical Programming Computation*, 2(3–4), 203–230.
- Weng, Y., Li, Q., Negi, R., & Ilić, M. (2013). Distributed algorithm for SDP state estimation. In *2013 IEEE PES innovative smart grid technologies conference (ISGT)* (pp. 1–6). IEEE.
- Yamashita, M., Fujisawa, K., Fukuda, M., Kobayashi, K., Nakata, K., & Nakata, M. (2012). Latest developments in the SDPA family for solving large-scale SDPs. In *Handbook on semidefinite, conic and polynomial optimization* (pp. 687–713). Springer.
- Yang, H., & Carlone, L. (2020). One ring to rule them all: Certifiably robust geometric perception with outliers. arXiv preprint arXiv:2006.06769.
- Yang, C.-H., & Deng, B. S. (2020). Exploiting sparsity in SDP relaxation for harmonic balance method. *IEEE Access*, 8, 115957–115965.
- Yannakakis, M. (1981). Computing the minimum fill-in is NP-complete. *SIAM Journal on Algebraic Discrete Methods*, 2(1), 77–79.
- Ye, Y. (2011). vol. 44, *Interior point algorithms: Theory and analysis*. John Wiley & Sons.
- Ye, Y., Todd, M. J., & Mizuno, S. (1994). An $O(\sqrt{n}L)$ -iteration homogeneous and self-dual linear programming algorithm. *Mathematics of Operations Research*, 19(1), 53–67.
- Yurtsever, A., Tropp, J. A., Fercoq, O., Udell, M., & Cevher, V. (2021). Scalable semidefinite programming. *SIAM Journal on Mathematics of Data Science*, 3(1), 171–200.
- Zhang, R. Y. (2020). On the tightness of semidefinite relaxations for certifying robustness to adversarial examples. arXiv preprint arXiv:2006.06759.
- Zhang, R., Fattahi, S., & Sojoudi, S. (2018). Large-scale sparse inverse covariance estimation via thresholding and max-det matrix completion. In *International conference on machine learning* (pp. 5766–5775). PMLR.
- Zhang, R. Y., & Lavaei, J. (2020a). Dual-CTC. https://github.com/ryz-codes/dual_ctc.
- Zhang, R. Y., & Lavaei, J. (2020b). Sparse semidefinite programs with guaranteed near-linear time complexity via dualized clique tree conversion. *Mathematical Programming*, 1–43.
- Zhang, Y., Madani, R., & Lavaei, J. (2017). Conic relaxations for power system state estimation with line measurements. *IEEE Transactions on Control of Network Systems*, 5(3), 1193–1205.
- Zhao, X.-Y., Sun, D., & Toh, K.-C. (2010). A Newton-CG augmented Lagrangian method for semidefinite programming. *SIAM Journal on Optimization*, 20(4), 1737–1765.
- Zheng, Y. (2019). *Chordal sparsity in control and optimization of large-scale systems* (Ph.D. thesis, recipient of the 2019 European Systems and Control PhD Thesis Award), University of Oxford.
- Zheng, Y., & Fantuzzi, G. (2020). Sum-of-squares chordal decomposition of polynomial matrix inequalities. arXiv preprint arXiv:2007.11410.
- Zheng, Y., Fantuzzi, G., & Papachristodoulou, A. (2018a). Decomposition and completion of sum-of-squares matrices. In *Proceedings of the 57th IEEE conference on decision and control* (pp. 4026–4031). IEEE.
- Zheng, Y., Fantuzzi, G., & Papachristodoulou, A. (2018b). Decomposition methods for large-scale semidefinite programs with chordal aggregate sparsity and partial orthogonality. In *Large-scale and distributed optimization* (pp. 33–55). Springer.
- Zheng, Y., Fantuzzi, G., & Papachristodoulou, A. (2018c). Fast ADMM for sum-of-squares programs using partial orthogonality. *IEEE Transactions on Automatic Control*, 64(9), 3869–3876.
- Zheng, Y., Fantuzzi, G., & Papachristodoulou, A. (2019). Sparse sum-of-squares (SOS) optimization: A bridge between DSOS/SDSOS and SOS optimization for sparse polynomials. In *Proceedings of the 2019 American control conference* (pp. 5513–5518). IEEE.
- Zheng, Y., Fantuzzi, G., Papachristodoulou, A., Goulart, P., & Wynn, A. (2016). CDCS: Cone decomposition conic solver, version 1.1. <https://github.com/oxfordcontrol/CDCS>.
- Zheng, Y., Fantuzzi, G., Papachristodoulou, A., Goulart, P., & Wynn, A. (2020). Chordal decomposition in operator-splitting methods for sparse semidefinite programs. *Mathematical Programming*, 180, 489–532.
- Zheng, Y., Kamgarpour, M., Sootla, A., & Papachristodoulou, A. (2018). Scalable analysis of linear networked systems via chordal decomposition. In *2018 European control conference (ECC)* (pp. 2260–2265). IEEE.
- Zheng, Y., Kamgarpour, M., Sootla, A., & Papachristodoulou, A. (2020). Distributed design for decentralized control using chordal decomposition and ADMM. *IEEE Transactions on Control of Network Systems*, 7(2), 614–626.
- Zheng, Y., Mason, R. P., & Papachristodoulou, A. (2018). Scalable design of structured controllers using chordal decomposition. *IEEE Transactions on Automatic Control*, 63(3), 752–767.
- Zheng, Y., Sootla, A., & Papachristodoulou, A. (2019). Block factor-width-two matrices and their applications to semidefinite and sum-of-squares optimization. arXiv preprint arXiv:1909.11076.
- Zheng, Y., Wang, J., & Li, K. (2020). Smoothing traffic flow via control of autonomous vehicles. *IEEE Internet of Things Journal*, 7(5), 3882–3896.
- Zhou, K., Doyle, J. C., Glover, K., et al. (1996). vol. 40, *Robust and optimal control*. Prentice hall New Jersey.
- Zhu, H., & Giannakis, G. B. (2014). Power system nonlinear state estimation using distributed semidefinite programming. *IEEE Journal of Selected Topics in Signal Processing*, 8(6), 1039–1050.