

Faster, but Weaker, Relaxations for Quadratically Constrained Quadratic Programs

Samuel Burer* Sunyoung Kim† Masakazu Kojima‡

April 21, 2013

Abstract

We introduce a new relaxation framework for nonconvex quadratically constrained quadratic programs (QCQPs). In contrast to existing relaxations based on semidefinite programming (SDP), our relaxations incorporate features of both SDP and second order cone programming (SOCP) and, as a result, solve more quickly than SDP. A downside is that the calculated bounds are weaker than those gotten by SDP. The framework allows one to choose a block-diagonal structure for the mixed SOCP-SDP, which in turn allows one to control the speed and bound quality. For a fixed block-diagonal structure, we also introduce a procedure to improve the bound quality without increasing computation time significantly. The effectiveness of our framework is illustrated on a large sample of QCQPs from various sources.

Keywords: nonconvex quadratic programming, semidefinite programming, second-order cone programming, difference of convex.

Mathematics Subject Classification: 90C26, 90C22, 90C25.

1 Introduction

Nonconvex quadratically constrained quadratic programs (QCQPs) model many different types of optimization problems and are generally NP-hard. Semidefinite programming (SDP) relaxations can provide tight bounds [6, 11, 12], but they can also be expensive to solve by

*Department of Management Sciences, University of Iowa, Iowa City, IA, 52242-1994, USA. Email: samuel-burer@uiowa.edu. The research of this author was supported in part by NSF Grant CCF-0545514.

†Department of Mathematics, Ewha W. University, 11-1 DaHyun-Dong, SuhDaeMoon-Gu, Seoul, Korea 120-750. Email: skim@ewha.ac.kr. The research of this author was supported by NRF 2012-R1A1A2-038982 and NRF 2010-000-8784.

‡Research and Development Initiative & JST CREST, Chuo University, 1-13-27, Kasuga, Bunkyo-ku, Tokyo 112-8551 Japan. This research was partially supported by the Japan Science and Technology Agency (JST), the Core Research of Evolutionary Science and Technology (CREST) Research Project. Email: kojimamasakazu@mac.com.

classical interior-point methods [14]. Many researchers have thus suggested various alternatives to interior-point methods for improving solution times [1, 2, 4, 8, 15, 19]. These efforts have been quite successful on many classes of SDP relaxations.

Others have studied different types of relaxations, for example, ones based on linear programming [10, 13, 18] or second-order cone programming (SOCP) [9, 17, 20]. Generally speaking, one would expect such relaxations to provide weaker bounds in less time compared to SDP relaxations. In fact, SOCP-based relaxations are often constructed as further relaxations of SDP relaxations. So, in a certain sense, one can see (as discussed in Section 2) that SOCP relaxations are never tighter than their SDP counterparts.

In this paper, we describe a “middle way” between SOCP and SDP relaxations. Our contribution is a framework for constructing mixed SOCP-SDP relaxations of QCQPs that allows one to balance the trade-off between solution time and bound quality. The key idea is a particular d.c. (difference-of-convex) strategy for further relaxing the linear constraints of an SDP relaxation of a QCQP. While related d.c. approaches have been studied previously, ours uniquely ensures a favorable block-diagonal structure on the resultant SOCP-SDP while simultaneously improving the bound quality. We illustrate the effectiveness of our approach on a large sample of QCQPs from various sources.

Our framework is not without drawbacks. Notably, there are several different choices one must make before applying the framework to a specific QCQP instance, and we do not know how to predict effectively the impact of these choices on the final solution time and bound quality. Indeed, it may be impossible to do so, but this is an interesting avenue for future research. For this paper, we simply illustrate the general behavior of our framework under various options for these choices.

1.1 Notation

In this paper, bold capital letters, such as \mathbf{M} , indicate matrices, and bold lowercase letters, such as \mathbf{v} , indicate vectors. The notation $\mathbf{A} \succeq \mathbf{B}$, $\mathbf{B} \preceq \mathbf{A}$ means that $\mathbf{A} - \mathbf{B}$ is positive semidefinite, and $\mathbf{A} \succ \mathbf{B}$, $\mathbf{B} \prec \mathbf{A}$ means $\mathbf{A} - \mathbf{B}$ is positive definite. The special matrices \mathbf{I} and \mathbf{O} are the identity matrix and all-zeros matrix, respectively. For matrices \mathbf{M} and \mathbf{N} of the same size, the notation $\mathbf{M} \bullet \mathbf{N} := \text{trace}(\mathbf{M}^T \mathbf{N})$ is the matrix inner product, and for vectors \mathbf{v} and \mathbf{w} of the same size, $\mathbf{v} \circ \mathbf{w}$ is the Hadamard product, i.e., component-wise product. For positive integer n , we define $[n] := \{1, \dots, n\}$.

2 The Problem and Existing Techniques

We study the QCQP

$$\begin{aligned} & \text{minimize} && \mathbf{x}^T \mathbf{A}_0 \mathbf{x} + \mathbf{a}_0^T \mathbf{x} \\ & \text{subject to} && \mathbf{x}^T \mathbf{A}_i \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0 \quad (i = 1, \dots, m) \end{aligned} \quad (1)$$

where $\mathbf{A}_i \in \mathbb{S}^n$, $\mathbf{a}_i \in \mathbb{R}^n$, and $\alpha_i \in \mathbb{R}$ for all $i = 0, 1, \dots, m$. Let F denote the feasible set of (1). The basic SDP relaxation of (1) is

$$\begin{aligned} & \text{minimize} && \mathbf{A}_0 \bullet \mathbf{X} + \mathbf{a}_0^T \mathbf{x} \\ & \text{subject to} && \mathbf{A}_i \bullet \mathbf{X} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0 \quad (i = 1, \dots, m) \\ & && \mathbf{X} \succeq \mathbf{x} \mathbf{x}^T. \end{aligned} \quad (2)$$

In some cases, it may be possible to strengthen the SDP relaxation — say, by first adding redundant quadratic constraints to (1) before deriving the SDP relaxation — but here we assume that (1) already contains all constraints of interest.

It may also be possible to construct SOCP relaxations of (1) in the original variable space. For example, such a relaxation could be represented as

$$\begin{aligned} & \text{minimize} && \mathbf{x}^T \mathbf{B}_0 \mathbf{x} + \mathbf{b}_0^T \mathbf{x} \\ & \text{subject to} && \mathbf{x}^T \mathbf{B}_l \mathbf{x} + \mathbf{b}_l^T \mathbf{x} + \beta_l \leq 0 \quad (l \in L), \end{aligned} \quad (3)$$

where L is an arbitrary index set and all $\mathbf{B}_0, \mathbf{B}_l \succeq \mathbf{O}$. More precisely, we say that (3) is an *SOCP relaxation of (1)* if $\mathbf{x} \in F$ implies that \mathbf{x} is feasible for (3) and the inequality $\mathbf{x}^T \mathbf{B}_0 \mathbf{x} + \mathbf{b}_0^T \mathbf{x} \leq \mathbf{x}^T \mathbf{A}_0 \mathbf{x} + \mathbf{a}_0^T \mathbf{x}$ holds. In fact, it is well known that SOCPs can also be solved as SDPs. Specifically, Fujie and Kojima [3] prove that (3) is equivalent to the SDP

$$\begin{aligned} & \text{minimize} && \mathbf{B}_0 \bullet \mathbf{X} + \mathbf{b}_0^T \mathbf{x} \\ & \text{subject to} && \mathbf{B}_l \bullet \mathbf{X} + \mathbf{b}_l^T \mathbf{x} + \beta_l \leq 0 \quad (l \in L) \\ & && \mathbf{X} \succeq \mathbf{x} \mathbf{x}^T. \end{aligned} \quad (4)$$

The following proposition states the equivalence of (3) and (4) in terms of our current setting.

Proposition 1. *Suppose the SOCP (3) is a valid relaxation of the QCQP (1). Then the SDP (4) is also a valid relaxation, and its optimal value equals that of (3).*

The upshot of Proposition 1 is that SOCP relaxations in \mathbf{x} can never provide better bounds than SDP relaxations in (\mathbf{x}, \mathbf{X}) .

Kim and Kojima [9] provided the first SOCP relaxation of (1) not relying on the variable

X. First, the authors assume without loss of generality that the objective of (1) is linear. This can be achieved, for example, by introducing a new variable $t \in \mathbb{R}$ and a new quadratic constraint $\mathbf{x}^T \mathbf{A}_0 \mathbf{x} + \mathbf{a}_0^T \mathbf{x} \leq t$ and then minimizing t . Next, each \mathbf{A}_i ($i = 1, \dots, m$) is written as the difference of two carefully chosen positive semidefinite $\mathbf{A}_i^+, \mathbf{A}_i^- \succeq \mathbf{O}$, i.e., $\mathbf{A}_i = \mathbf{A}_i^+ - \mathbf{A}_i^-$, so that i -th constraint may be expressed equivalently as

$$\mathbf{x}^T \mathbf{A}_i^+ \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq \mathbf{x}^T \mathbf{A}_i^- \mathbf{x}.$$

Then, an auxiliary variable $z_i \in \mathbb{R}$ is introduced to represent $\mathbf{x}^T \mathbf{A}_i^- \mathbf{x}$ but also immediately relaxed as $\mathbf{x}^T \mathbf{A}_i^- \mathbf{x} \leq z_i$ resulting in the convex system

$$\begin{aligned} \mathbf{x}^T \mathbf{A}_i^+ \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i &\leq z_i \\ \mathbf{x}^T \mathbf{A}_i^- \mathbf{x} &\leq z_i. \end{aligned}$$

Finally, z_i must be bounded in some fashion, say as $z_i \leq \mu_i \in \mathbb{R}$, or else the relaxation will in fact be useless. Bounding z_i depends very much on the problem and the choice of $\mathbf{A}_i^+, \mathbf{A}_i^-$. In particular, [9] provides strategies for doing so.

Closely related approaches have recently been developed by Saxena et al. [17] and Zheng et al. [20]. In [17], the authors study the relaxation obtained by the following spectral splitting of \mathbf{A}_i :

$$\mathbf{A}_i = \left(\sum_{\lambda_{ij} > 0} \lambda_{ij} \mathbf{v}_{ij} \mathbf{v}_{ij}^T \right) - \left(\sum_{\lambda_{ij} < 0} |\lambda_{ij}| \mathbf{v}_{ij} \mathbf{v}_{ij}^T \right)$$

where $\{\lambda_{i1}, \dots, \lambda_{in}\}$ and $\{\mathbf{v}_{i1}, \dots, \mathbf{v}_{in}\}$ are the eigenvalues and eigenvectors of \mathbf{A}_i , respectively. The constraint $\mathbf{x}^T \mathbf{A}_i \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0$ can thus be reformulated as $\sum_{\lambda_{ij} > 0} \lambda_{ij} (\mathbf{v}_{ij}^T \mathbf{x})^2 + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq \sum_{\lambda_{ij} < 0} |\lambda_{ij}| (\mathbf{v}_{ij}^T \mathbf{x})^2$. Assuming known lower and upper bounds on the entries of \mathbf{x} , the non-convex terms $(\mathbf{v}_{ij}^T \mathbf{x})^2$ for $\lambda_{ij} < 0$ can be relaxed via a secant approximation to derive a convex relaxation of the above constraint. The paper [20] employs similar ideas but further solves a secondary SDP over different splittings of \mathbf{A}_i to improve the resultant SOCP relaxation quality.

3 A New Relaxation Framework

We are motivated by the idea that there should exist relaxations between the two extremes introduced in the previous section: SOCPs in only \mathbf{x} versus SDPs in both (\mathbf{x}, \mathbf{X}) . By “between,” we mean that these hypothesized relaxations solve faster than SDPs while providing better bounds than SOCPs. In this section, we present a general construction, which

provides these in-between relaxations.

Before detailing the general construction in Sections 3.1–3.4, we first introduce a specific case as a gentle introduction. For all $i = 0, 1, \dots, m$, define $\lambda_i := -\lambda_{\min}[\mathbf{A}_i]$ so that $\lambda_i \geq 0$ and $\mathbf{A}_i + \lambda_i \mathbf{I} \succeq \mathbf{O}$. Then (1) is equivalent to

$$\begin{aligned} & \text{minimize} && -\lambda_0 \mathbf{x}^T \mathbf{x} + \mathbf{x}^T (\mathbf{A}_0 + \lambda_0 \mathbf{I}) \mathbf{x} + \mathbf{a}_0^T \mathbf{x} \\ & \text{subject to} && -\lambda_i \mathbf{x}^T \mathbf{x} + \mathbf{x}^T (\mathbf{A}_i + \lambda_i \mathbf{I}) \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0 \quad (i = 1, \dots, m) \end{aligned}$$

which has the following mixed SOCP-SDP relaxation:

$$\begin{aligned} & \text{minimize} && -\lambda_0 \text{trace}(\mathbf{X}) + \mathbf{x}^T (\mathbf{A}_0 + \lambda_0 \mathbf{I}) \mathbf{x} + \mathbf{a}_0^T \mathbf{x} \\ & \text{subject to} && -\lambda_i \text{trace}(\mathbf{X}) + \mathbf{x}^T (\mathbf{A}_i + \lambda_i \mathbf{I}) \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0 \quad (i = 1, \dots, m) \quad (5) \\ & && \mathbf{X} \succeq \mathbf{x} \mathbf{x}^T. \end{aligned}$$

Notice that, other than the constraint $\mathbf{X} \succeq \mathbf{x} \mathbf{x}^T$, the only variables in \mathbf{X} to appear in the objective and remaining constraints are the variables \mathbf{X}_{jj} . Said differently, the entries \mathbf{X}_{jk} for $j \neq k$ are only relevant for the semidefinite constraint. In addition, one can see that, with \mathbf{x} fixed, the diagonal entries of \mathbf{X} can be made arbitrarily large to satisfy all constraints with $\lambda_i > 0$ as well as drive the objective to $-\infty$ if $\lambda_0 > 0$. So, in general, one should bound \mathbf{X}_{jj} to form a sensible relaxation. For the sake of presentation, let us assume that each \mathbf{x}_j is bounded in $[0, 1]$ in (1) so that the constraints $\mathbf{X}_{jj} \leq \mathbf{x}_j$ are valid for the SDP relaxation.

We remark that, instead of defining $\lambda_i := -\lambda_{\min}[\mathbf{A}_i]$ as above, another possibility would be to set λ_i to $\max\{0, -\lambda_{\min}[\mathbf{A}_i]\}$, still ensuring $\mathbf{A}_i + \lambda_i \mathbf{I}$. In words, $\mathbf{A}_i + \lambda_i \mathbf{I}$ would equal \mathbf{A}_i whenever \mathbf{A}_i is already positive semidefinite. However, we have chosen the stated definition of λ_i because it may lead to a stronger relaxation (5). For example, consider the inequality $\mathbf{x}^T \mathbf{x} - 1 \leq 0$. Our definition of λ_i yields the constraint $\text{trace}(\mathbf{X}) - 1 \leq 0$, whereas the alternative would keep $\mathbf{x}^T \mathbf{x} - 1 \leq 0$. Since $\mathbf{X} \succeq \mathbf{x} \mathbf{x}^T$ implies $\mathbf{x}^T \mathbf{x} \leq \text{trace}(\mathbf{X})$, the former inequality is stronger than the latter.

The definition of λ_i has another benefit. If the original problem (1) contains a strictly convex inequality constraint with $\mathbf{A}_i \succ \mathbf{O}$, the feasible region of (5) will be bounded without the assumption $\mathbf{x}_j \in [0, 1]$ and constraint $\mathbf{X}_{jj} \leq \mathbf{x}_j$ as described above. For example, the spherical constraint $\mathbf{x}^T \mathbf{x} - 1 \leq 0$ would be transformed to $\text{trace}(\mathbf{X}) - 1 \leq 0$, thus bounding the feasible region of (5) in conjunction with $\mathbf{X} \succeq \mathbf{x} \mathbf{x}^T$. Still, in this paper, our preference is to use $\mathbf{x}_j \in [0, 1]$ and $\mathbf{X}_{jj} \leq \mathbf{x}_j$ to establish boundedness.

Consider now the following proposition:

Proposition 2 (Grone et. al [7]). *Given a vector \mathbf{x} and scalars $\mathbf{X}_{11}, \dots, \mathbf{X}_{nn}$, there exists a symmetric-matrix completion \mathbf{X} of $\mathbf{X}_{11}, \dots, \mathbf{X}_{nn}$ satisfying $\mathbf{X} \succeq \mathbf{x} \mathbf{x}^T$ if and only if*

$\mathbf{X}_{jj} \geq \mathbf{x}_j^2$ for all $j = 1, \dots, n$.

Proof. This follows from the chordal arrow structure of the matrix

$$\begin{pmatrix} 1 & \mathbf{x}^T \\ \mathbf{x} & \mathbf{Diag}(\mathbf{X}_{11}, \dots, \mathbf{X}_{nn}) \end{pmatrix},$$

where $\mathbf{Diag}(\cdot)$ places its arguments in a diagonal matrix. We refer the reader to Grone et. al [7] and Fukuda et. al [4] for further details. \square

In light of Proposition 2, problem (5) with additional bounding constraints $\mathbf{X}_{jj} \leq \mathbf{x}_j$ is equivalent to

$$\begin{aligned} \text{minimize} \quad & -\lambda_0 \sum_{j=1}^n \mathbf{X}_{jj} + \mathbf{x}^T (\mathbf{A}_0 + \lambda_0 \mathbf{I}) \mathbf{x} + \mathbf{a}_0^T \mathbf{x} \\ \text{subject to} \quad & -\lambda_i \sum_{j=1}^n \mathbf{X}_{jj} + \mathbf{x}^T (\mathbf{A}_i + \lambda_i \mathbf{I}) \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0 \quad (i = 1, \dots, m) \\ & \mathbf{x}_j^2 \leq \mathbf{X}_{jj} \leq \mathbf{x}_j \quad (j = 1, \dots, n). \end{aligned} \quad (6)$$

Compared to the SDP relaxation (2), which has $O(n^2)$ variables, problem (6) has only $O(n)$ variables and hence is much faster to solve. Of course, its bound on (1) should generally be weaker than the SDP bound.

It may actually be possible to improve the bound quality of (6) without changing its basic structure and computational complexity. Instead of splitting \mathbf{A}_i into $\mathbf{A}_i + \lambda_i \mathbf{I} \succeq \mathbf{O}$ and $-\lambda_i \mathbf{I} \preceq 0$, we could more generally split it into $\mathbf{A}_i + \mathbf{D}_i \succeq \mathbf{O}$ and $-\mathbf{D}_i$, where \mathbf{D}_i is a diagonal (not necessarily positive semidefinite) matrix. The resultant relaxation

$$\begin{aligned} \text{minimize} \quad & \sum_{j=1}^n [\mathbf{D}_0]_{jj} \mathbf{X}_{jj} + \mathbf{x}^T (\mathbf{A}_0 + \mathbf{D}_0) \mathbf{x} + \mathbf{a}_0^T \mathbf{x} \\ \text{subject to} \quad & \sum_{j=1}^n [\mathbf{D}_i]_{jj} \mathbf{X}_{jj} + \mathbf{x}^T (\mathbf{A}_i + \mathbf{D}_i) \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0 \quad (i = 1, \dots, m) \\ & \mathbf{x}_j^2 \leq \mathbf{X}_{jj} \leq \mathbf{x}_j \quad (j = 1, \dots, n), \end{aligned}$$

could provide a better bound than (6) if the \mathbf{D}_i are chosen intelligently, but it should still require roughly the same amount of time to solve.

3.1 The mixed SOCP-SDP relaxation

Given a positive integer $r \leq n$, let $\mathcal{C} := \{C_1, \dots, C_r\}$ be a partition of the indices $[n]$. We assume without loss of generality that \mathcal{C} satisfies $\max(C_k) \leq \min(C_{k+1})$ for all $k = 1, \dots, r-1$. In words, C_1 consists of the first few indices in $[n]$, C_2 consists of the next few, and so on. We say that a symmetric matrix \mathbf{D} is \mathcal{C} -block diagonal if $\mathbf{D}_{jk} = 0$ whenever j and k are members of different sets in \mathcal{C} . That is, \mathbf{D} is block diagonal with blocks $\mathbf{D}_{C_1 C_1}, \dots, \mathbf{D}_{C_r C_r}$.

Based on the partition \mathcal{C} , we construct a mixed SOCP-SDP relaxation of (1) as follows. For each $i = 0, 1, \dots, m$, let \mathbf{D}_i be a \mathcal{C} -block diagonal matrix satisfying $\mathbf{A}_i + \mathbf{D}_i \succeq \mathbf{O}$. This yields the following problem equivalent to (1), where we assume $\mathbf{x} \in [0, 1]^n$ as above:

$$\begin{aligned} & \text{minimize} && -\mathbf{x}^T \mathbf{D}_0 \mathbf{x} + \mathbf{x}^T (\mathbf{A}_0 + \mathbf{D}_0) \mathbf{x} + \mathbf{a}_0^T \mathbf{x} \\ & \text{subject to} && -\mathbf{x}^T \mathbf{D}_i \mathbf{x} + \mathbf{x}^T (\mathbf{A}_i + \mathbf{D}_i) \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0 \quad (i = 1, \dots, m) \\ & && \mathbf{x} \circ \mathbf{x} \leq \mathbf{x}. \end{aligned}$$

Its mixed SOCP-SDP relaxation is

$$\begin{aligned} & \text{minimize} && -\mathbf{D}_0 \bullet \mathbf{X} + \mathbf{x}^T (\mathbf{A}_0 + \mathbf{D}_0) \mathbf{x} + \mathbf{a}_0^T \mathbf{x} \\ & \text{subject to} && -\mathbf{D}_i \bullet \mathbf{X} + \mathbf{x}^T (\mathbf{A}_i + \mathbf{D}_i) \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0 \quad (i = 1, \dots, m) \\ & && \mathbf{diag}(\mathbf{X}) \leq \mathbf{x} \\ & && \mathbf{X} \succeq \mathbf{x} \mathbf{x}^T. \end{aligned} \tag{7}$$

Relaxation (7) has the property that, due to the block structure of \mathbf{D}_i , the only portions of \mathbf{X} appearing in the objective and linear constraints are the blocks $\mathbf{X}_{C_1 C_1}, \dots, \mathbf{X}_{C_r C_r}$. The other entries of \mathbf{X} are only relevant for the semidefiniteness constraint $\mathbf{X} \succeq \mathbf{x} \mathbf{x}^T$. The following proposition allows us to eliminate all entries of \mathbf{X} except for the blocks $\mathbf{X}_{C_k C_k}$:

Proposition 3 (Grone et. al [7]). *Given a vector \mathbf{x} and symmetric blocks $\mathbf{X}_{C_1 C_1}, \dots, \mathbf{X}_{C_r C_r}$, there exists a symmetric-matrix completion \mathbf{X} of $\mathbf{X}_{C_1 C_1}, \dots, \mathbf{X}_{C_r C_r}$ satisfying $\mathbf{X} \succeq \mathbf{x} \mathbf{x}^T$ if and only $\mathbf{X}_{C_k C_k} \succeq \mathbf{x}_{C_k} \mathbf{x}_{C_k}^T$ for all $k = 1, \dots, r$.*

Proof. This follows from the chordal arrow structure of the matrix

$$\begin{pmatrix} 1 & & \mathbf{x}^T \\ & & \\ \mathbf{x} & \mathbf{Diag}(\mathbf{X}_{C_1 C_1}, \dots, \mathbf{X}_{C_r C_r}) & \end{pmatrix},$$

where $\mathbf{Diag}(\cdot)$ places its arguments in a diagonal matrix. We refer the reader to Grone et. al [7] and Fukuda et. al [4] for further details. \square

In light of the Proposition 3, problem (7) is equivalent to

$$\begin{aligned} & \text{minimize} && -\sum_{k=1}^r [\mathbf{D}_0]_{C_k C_k} \bullet \mathbf{X}_{C_k C_k} + \mathbf{x}^T (\mathbf{A}_0 + \mathbf{D}_0) \mathbf{x} + \mathbf{a}_0^T \mathbf{x} \\ & \text{subject to} && -\sum_{k=1}^r [\mathbf{D}_i]_{C_k C_k} \bullet \mathbf{X}_{C_k C_k} + \mathbf{x}^T (\mathbf{A}_i + \mathbf{D}_i) \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0 \quad (i = 1, \dots, m) \\ & && \mathbf{X}_{jj} \leq \mathbf{x}_j \quad (j = 1, \dots, n) \\ & && \mathbf{X}_{C_k C_k} \succeq \mathbf{x}_{C_k} \mathbf{x}_{C_k}^T \quad (k = 1, \dots, r). \end{aligned} \tag{8}$$

It is important to keep in mind that only the portions $\mathbf{X}_{C_k C_k}$ of the original \mathbf{X} actually remain in the problem. We say that (8) has a \mathcal{C} -block structure.

Problem (8) is the “in-between” relaxation that we propose to solve. However, it remains to choose the splitting of \mathbf{A}_i intelligently. To make the notation of the next subsections easier, we apply the linear change of variables $\mathbf{B}_i := \mathbf{A}_i + \mathbf{D}_i$ so that $\mathbf{B}_i \succeq \mathbf{O}$ and $\mathbf{B}_i - \mathbf{A}_i$ is \mathcal{C} -block diagonal. We also rewrite (8) as

$$\begin{aligned}
& \text{minimize} && (\mathbf{A}_0 - \mathbf{B}_0) \bullet \mathbf{X} + \mathbf{x}^T \mathbf{B}_0 \mathbf{x} + \mathbf{a}_0^T \mathbf{x} \\
& \text{subject to} && (\mathbf{A}_i - \mathbf{B}_i) \bullet \mathbf{X} + \mathbf{x}^T \mathbf{B}_i \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + \alpha_i \leq 0 \quad (i = 1, \dots, m) \\
& && \text{diag}(\mathbf{X}) \leq \mathbf{x} \\
& && \mathbf{X} \succeq \mathbf{x} \mathbf{x}^T
\end{aligned} \tag{9}$$

in terms of \mathbf{B}_i and the full matrix \mathbf{X} . We stress that the use of \mathbf{X} is only to simplify the presentation. In computation, the alternative yet equivalent representation in terms of $\mathbf{X}_{C_k C_k}$ would be employed to exploit the \mathcal{C} -block structure of (9).

3.2 How to choose the matrices \mathbf{B}_i

We now discuss how to choose the matrices \mathbf{B}_i intelligently. Consistent with the discussion in the previous subsection, \mathbf{B}_i is a feasible choice as long as it is an element of

$$\mathcal{F}(\mathbf{A}_i) := \{\mathbf{M} \succeq \mathbf{O} : \mathbf{A}_i - \mathbf{M} \text{ is } \mathcal{C}\text{-block diagonal}\}.$$

One idea would be to choose the collection $\{\mathbf{B}_i\}$ that results in the tightest bound from (9). In other words, $\{\mathbf{B}_i\}$ could be chosen as an optimal solution of

$$\begin{aligned}
& \text{maximize} && f(\mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_m) \\
& \text{subject to} && \mathbf{B}_i \in \mathcal{F}(\mathbf{A}_i) \quad (i = 0, \dots, m)
\end{aligned}$$

where $f(\mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_m)$ is the optimal value of (9). However, the objective of this problem is non-convex, and so the problem is likely to require as much time to solve as the SDP relaxation (2). It is then not worth the effort since it cannot even produce a better bound than (2) by Proposition 1. So the choice of the \mathbf{B}_i 's must balance the cost to compute them with the resultant bound quality.

As a compromise, consider the following observation:

Proposition 4. *For each $i = 0, 1, \dots, m$, suppose $\mathbf{B}_i, \widehat{\mathbf{B}}_i \in \mathcal{F}(\mathbf{A}_i)$ satisfy $\mathbf{B}_i \succeq \widehat{\mathbf{B}}_i$. Then*

$$f(\mathbf{B}_0, \mathbf{B}_1, \dots, \mathbf{B}_m) \leq f(\widehat{\mathbf{B}}_0, \widehat{\mathbf{B}}_1, \dots, \widehat{\mathbf{B}}_m).$$

Proof. We must show that the optimal value of (9) cannot deteriorate when all \mathbf{B}_i are replaced by $\widehat{\mathbf{B}}_i$. To see this, let (\mathbf{x}, \mathbf{X}) be feasible for the problem with $\widehat{\mathbf{B}}_i$. Note that $\mathbf{B}_i \succeq \widehat{\mathbf{B}}_i$ and $\mathbf{X} \succeq \mathbf{x}\mathbf{x}^T$ imply $(\mathbf{B}_i - \widehat{\mathbf{B}}_i) \bullet (\mathbf{X} - \mathbf{x}\mathbf{x}^T) \geq 0$, and so, for all i , the inequalities

$$\begin{aligned} & (\mathbf{A}_i - \mathbf{B}_i) \bullet \mathbf{X} + \mathbf{x}^T \mathbf{B}_i \mathbf{x} + \mathbf{a}_i^T \mathbf{x} \\ & \leq (\mathbf{A}_i - \mathbf{B}_i) \bullet \mathbf{X} + \mathbf{x}^T \mathbf{B}_i \mathbf{x} + \mathbf{a}_i^T \mathbf{x} + (\mathbf{B}_i - \widehat{\mathbf{B}}_i) \bullet (\mathbf{X} - \mathbf{x}\mathbf{x}^T) \\ & = (\mathbf{A}_i - \widehat{\mathbf{B}}_i) \bullet \mathbf{X} + \mathbf{x}^T \widehat{\mathbf{B}}_i \mathbf{x} + \mathbf{a}_i^T \mathbf{x} \end{aligned}$$

imply that (\mathbf{x}, \mathbf{X}) is feasible for (9) with no higher objective value. \square

Proposition 4 suggests that a reasonable choice of $\{\mathbf{B}_i\}$ should at least have the property that each \mathbf{B}_i is a “minimally positive semidefinite” member of $\mathcal{F}(\mathbf{A}_i)$. In the next subsection, we deal with the issue of finding such a minimal member of $\mathcal{F}(\mathbf{A}_i)$ for each i .

3.3 Minimal and minimum elements

The previous subsection has discussed choosing good matrices $\{\mathbf{B}_i\}$ to form (9). Our idea is to choose a “minimal” element in the set $\mathcal{F}(\mathbf{A}_i)$. Here, we discuss this issue formally and computationally. We also discuss the related idea of “minimum” elements. The results in this subsection are general and will be applied specifically to choose $\{\mathbf{B}_i\}$ in Section 3.4.

Given a nonempty, closed, convex subset \mathcal{T} of the positive semidefinite matrices, a member $\mathbf{T} \in \mathcal{T}$ is called *minimal* if $\{\mathbf{M} \in \mathcal{T} : \mathbf{T} \succeq \mathbf{M}\} = \{\mathbf{T}\}$. In other words, \mathbf{T} is minimal if there exists no \mathbf{M} in \mathcal{T} distinct from \mathbf{T} such that $\mathbf{M} \preceq \mathbf{T}$. Furthermore, $\mathbf{T} \in \mathcal{T}$ is called *minimum* if $\mathbf{T} \preceq \mathbf{M}$ for all $\mathbf{M} \in \mathcal{T}$. It is not difficult to see that every \mathcal{T} has at least one minimal element and that \mathcal{T} has a minimum element \mathbf{T} if and only if \mathbf{T} is the unique minimal element. We let $\text{Minimal}(\mathcal{T})$ denote the set of minimal elements in \mathcal{T} and $\text{minimum}(\mathcal{T})$ denote the minimum element if it exists.

Example 3.1. Let

$$\mathcal{T} := \left\{ \begin{pmatrix} t_{11} & t_{12} \\ t_{21} & t_{22} \end{pmatrix} \in \mathbb{S}_+^2 : t_{12} = 1 \right\} = \left\{ \begin{pmatrix} t_{11} & 1 \\ 1 & t_{22} \end{pmatrix} \in \mathbb{S}^2 : \begin{array}{l} t_{11} > 0, t_{22} > 0 \\ t_{11}t_{22} \geq 1 \end{array} \right\}.$$

One can show

$$\text{Minimal}(\mathcal{T}) = \left\{ \begin{pmatrix} t_{11} & 1 \\ 1 & t_{22} \end{pmatrix} \in \mathbb{S}^2 : \begin{array}{l} t_{11} > 0, t_{22} > 0 \\ t_{11}t_{22} = 1 \end{array} \right\} = \left\{ \begin{pmatrix} t_{11} & 1 \\ 1 & t_{11}^{-1} \end{pmatrix} \in \mathbb{S}^2 : t_{11} > 0 \right\}.$$

In this case, there are multiple minimal elements and hence no minimum element. We also

note that the rank of all minimal elements is small, namely 1.

By solving an SDP, it is always possible to calculate some $\mathbf{T} \in \text{Minimal}(\mathcal{T})$. Choose any positive definite matrix \mathbf{C} , and let \mathbf{T} be an optimal solution of $\min\{\mathbf{C} \bullet \mathbf{X} : \mathbf{X} \in \mathcal{T}\}$. To see that \mathbf{T} is minimal, assume on the contrary that $\mathbf{X} \in \mathcal{T}$ satisfies $\mathbf{T} \neq \mathbf{X} \preceq \mathbf{T}$. Then $\mathbf{C} \bullet (\mathbf{X} - \mathbf{T}) < 0$, contradicting the optimality of \mathbf{T} .

However, in this paper we do not suggest solving an SDP to get a minimal element in \mathcal{T} because—at least for our application—doing so would require as much time as solving (2). So we suggest a different technique that will work quickly in our specialized setting. Accordingly, let us focus on specific sets \mathcal{T} relevant to this paper. Given $\mathbf{A} \in \mathbb{S}^n$, define

$$\mathcal{F}(\mathbf{A}) := \{\mathbf{M} \succeq \mathbf{O} : \mathbf{A} - \mathbf{M} \text{ is } \mathcal{C}\text{-block diagonal}\}.$$

One can think of \mathbf{A} as generically representing a matrix \mathbf{A}_i in (1), although \mathbf{A} is not necessarily limited in this way. We would like to compute some $\mathbf{B} \in \text{Minimal}(\mathcal{F}(\mathbf{A}))$ without solving an SDP. Our approach is proposed as Algorithm 2 below.

We start with an observation that $\mathcal{F}(\mathbf{A})$ does not change even if \mathbf{A} is shifted by a \mathcal{C} -block diagonal matrix:

Proposition 5. *Given $\mathbf{A} \in \mathbb{S}^n$, let Δ be any \mathcal{C} -block diagonal matrix. Then $\mathcal{F}(\mathbf{A}) = \mathcal{F}(\mathbf{A} + \Delta)$.*

Proof. This follows from the fact that $\mathbf{A} - \mathbf{M}$ is \mathcal{C} -block diagonal if and only if $(\mathbf{A} + \Delta) - \mathbf{M}$ is. □

This proposition will give us some useful flexibility. To calculate $\mathbf{B} \in \text{Minimal}(\mathcal{F}(\mathbf{A}))$, Algorithm 2 will first shift \mathbf{A} by a \mathcal{C} -block diagonal matrix Δ such that $\mathbf{A} + \Delta \succeq \mathbf{O}$ and then calculate $\mathbf{B} \in \text{Minimal}(\mathcal{F}(\mathbf{A} + \Delta))$. The choice of the shift Δ may have an effect on the final calculated \mathbf{B} , and we will discuss different choices of Δ in Section 3.4.

Algorithm 2 will make use of a specialized subroutine (Algorithm 1). Abusing notation, let an input $\mathbf{A} \succeq \mathbf{O}$ to this subroutine be fixed, and for each $k = 1, \dots, r$, define the following restriction of $\mathcal{F}(\mathbf{A})$, which depends on the subset C_k of $[n]$:

$$\mathcal{F}_k(\mathbf{A}) := \{\mathbf{O} \preceq \mathbf{M} \preceq \mathbf{A} : \mathbf{A} - \mathbf{M} \text{ is } C_k\text{-block diagonal}\}$$

where a matrix is defined to be C_k -block diagonal if all entries outside of its $C_k C_k$ block are zero. We present Algorithm 1 for calculating $\mathbf{B} = \text{minimum}(\mathcal{F}_k(\cdot))$. Lemma 1 and Proposition 6 establish the correctness of Algorithm 1.

Algorithm 1 Calculate the minimum element of $\mathcal{F}_k(\mathbf{A})$ for $\mathbf{A} \succeq \mathbf{O}$

Input: $C_k \subseteq [n]$, $\mathbf{A} \succeq \mathbf{O}$, $p := \text{rank}(\mathbf{A})$, $\mathbf{L} \in \mathbb{R}^{n \times p}$ s.t. $\mathbf{A} = \mathbf{L}\mathbf{L}^T$ if $p \geq 1$.

Output: $\mathbf{B} = \text{minimum}(\mathcal{F}_k(\mathbf{A}))$ $q := \text{rank}(\mathbf{B}) \leq p$, $\widehat{\mathbf{L}} \in \mathbb{R}^{n \times q}$ s.t. $\mathbf{B} = \widehat{\mathbf{L}}\widehat{\mathbf{L}}^T$ if $q \geq 1$.

- 1: If $p = 0$ or $C_k = [n]$, then return $\mathbf{B} = \mathbf{O}$.
 - 2: Define $S := \text{span}\{\mathbf{L}_{i\cdot}^T : i \notin C_k\} \subseteq \mathbb{R}^p$. Let $q := \dim(S)$.
 - 3: Construct $\mathbf{V} \in \mathbb{R}^{p \times q}$ with columns forming an orthonormal basis of S .
 - 4: Return $\widehat{\mathbf{L}} := \mathbf{L}\mathbf{V}$ and $\mathbf{B} := \widehat{\mathbf{L}}\widehat{\mathbf{L}}^T$.
-

Lemma 1. Given \mathbf{M} such that $\mathbf{M}\mathbf{M}^T \preceq \mathbf{I}$, suppose \mathbf{N} has orthonormal columns with $\text{range}(\mathbf{M}) \subseteq \text{range}(\mathbf{N})^\perp$. Then $\mathbf{M}\mathbf{M}^T + \mathbf{N}\mathbf{N}^T \preceq \mathbf{I}$.

Proof. Let \mathbf{W} be a matrix of orthonormal columns spanning $\text{range}(\mathbf{N})^\perp$. Then $\text{range}(\mathbf{M}) \subseteq \text{range}(\mathbf{W})$ and there exists \mathbf{H} such that $\mathbf{M} = \mathbf{W}\mathbf{H}$.

We claim $\mathbf{H}\mathbf{H}^T \preceq \mathbf{I}$, where \mathbf{I} is of the appropriate size. If not, then there exists \mathbf{x} such that $\mathbf{x}^T \mathbf{H}\mathbf{H}^T \mathbf{x} > \mathbf{x}^T \mathbf{x}$. Defining $\mathbf{v} := \mathbf{W}\mathbf{x}$, we see

$$\mathbf{v}^T \mathbf{M}\mathbf{M}^T \mathbf{v} = (\mathbf{W}\mathbf{x})^T (\mathbf{W}\mathbf{H}\mathbf{H}^T \mathbf{W}^T) (\mathbf{W}\mathbf{x}) = \mathbf{x}^T \mathbf{H}\mathbf{H}^T \mathbf{x} > \mathbf{x}^T \mathbf{x} = \mathbf{x}^T \mathbf{W}^T \mathbf{W}\mathbf{x} = \mathbf{v}^T \mathbf{v},$$

which contradicts $\mathbf{M}\mathbf{M}^T \preceq \mathbf{I}$.

Next, for arbitrary \mathbf{v} , write $\mathbf{v} = \mathbf{W}\mathbf{x} + \mathbf{N}\mathbf{y}$; note that $\mathbf{v}^T \mathbf{v} = \mathbf{x}^T \mathbf{x} + \mathbf{y}^T \mathbf{y}$. Then

$$\begin{aligned} \mathbf{v}^T (\mathbf{M}\mathbf{M}^T + \mathbf{N}\mathbf{N}^T) \mathbf{v} &= (\mathbf{W}\mathbf{x} + \mathbf{N}\mathbf{y})^T (\mathbf{W}\mathbf{H}\mathbf{H}^T \mathbf{W}^T + \mathbf{N}\mathbf{N}^T) (\mathbf{W}\mathbf{x} + \mathbf{N}\mathbf{y}) \\ &= \mathbf{x}^T \mathbf{H}\mathbf{H}^T \mathbf{x} + \mathbf{y}^T \mathbf{y} \\ &= \mathbf{x}^T (\mathbf{H}\mathbf{H}^T - \mathbf{I}) \mathbf{x} + \mathbf{v}^T \mathbf{v} \\ &\leq \mathbf{v}^T \mathbf{v}. \end{aligned}$$

This proves the result. □

Proposition 6. Algorithm 1 correctly calculates $\mathbf{B} = \text{minimum}(\mathcal{F}_k(\mathbf{A}))$ when $\mathbf{A} \succeq \mathbf{O}$.

Proof. Clearly, $\mathbf{B} \succeq \mathbf{O}$. We next show $\mathbf{B} \preceq \mathbf{A}$. If $p = 0$ or $C_k = [n]$ so that $\mathbf{B} = \mathbf{O}$, then $\mathbf{B} \preceq \mathbf{A}$ is obvious. Otherwise, if $1 \leq p \leq r$, then $\mathbf{A} - \mathbf{B} = \mathbf{L}\mathbf{L} - \widehat{\mathbf{L}}\widehat{\mathbf{L}}^T = \mathbf{L}(\mathbf{I}_r - \mathbf{V}\mathbf{V}^T)\mathbf{L}$. This is positive semidefinite because $\mathbf{V}\mathbf{V}^T \preceq \mathbf{I}_r$ from Lemma 1 with $(\mathbf{M}, \mathbf{N}) = (\mathbf{0}, \mathbf{V})$.

We now argue that $\mathbf{A} - \mathbf{B}$ is C_k -block diagonal. If $p = 0$ or $C_k = [n]$ so that $\mathbf{B} = \mathbf{O}$, then this is clear. If $1 \leq p \leq r$, then note that $\mathbf{V}\mathbf{V}^T \in \mathbb{R}^{r \times r}$ serves as the orthogonal projection matrix from \mathbb{R}^r onto the subspace S . Since $\mathbf{L}_{i\cdot}^T \in S$ for each $i \notin C_k$, we see that $\mathbf{V}\mathbf{V}^T \mathbf{L}_{i\cdot}^T = \mathbf{L}_{i\cdot}^T$ for $i \notin C_k$. Therefore, $(i, j) \notin C_k \times C_k$ ensures

$$\mathbf{A}_{ij} - \mathbf{B}_{ij} = \mathbf{A}_{ij} - \mathbf{L}_{i\cdot} \mathbf{V}\mathbf{V}^T \mathbf{L}_{j\cdot}^T = \mathbf{A}_{ij} - \mathbf{L}_{i\cdot} \mathbf{L}_{j\cdot}^T = 0,$$

i.e., $\mathbf{A} - \mathbf{B}$ is C_k -block diagonal.

The preceding two paragraphs establish $\mathbf{B} \in \mathcal{F}_k(\mathbf{A})$. We now prove that \mathbf{B} is in fact minimum in $\mathcal{F}_k(\mathbf{A})$. That is, we prove $\mathbf{B} \preceq \mathbf{Z}$ for all $\mathbf{Z} \in \mathcal{F}_k(\mathbf{A})$. So let $\mathbf{Z} \in \mathcal{F}_k(\mathbf{A})$ be arbitrary. If $\mathbf{Z} = \mathbf{A}$, we already know $\mathbf{B} \preceq \mathbf{Z}$. Likewise, if $\mathbf{B} = \mathbf{O}$, the result is clear. So assume $p \geq 1$ and $q := \text{rank}(\mathbf{A} - \mathbf{Z}) \geq 1$.

Since $\mathbf{O} \preceq \mathbf{A} - \mathbf{Z}$, there exists a rank- q $\mathbf{F} \in \mathbb{R}^{n \times q}$ such that $\mathbf{A} - \mathbf{Z} = \mathbf{F}\mathbf{F}^T$. Hence, $\mathbf{Z} = \mathbf{A} - \mathbf{F}\mathbf{F}^T = \mathbf{L}\mathbf{L}^T - \mathbf{F}\mathbf{F}^T \succeq \mathbf{O}$. It follows that $\text{null}(\mathbf{L}^T) \subseteq \text{null}(\mathbf{F}^T)$, which implies $\text{range}(\mathbf{F}) \subseteq \text{range}(\mathbf{L})$. So there exists rank- q $\mathbf{G} \in \mathbb{R}^{r \times q}$ such that $\mathbf{F} = \mathbf{L}\mathbf{G}$. We also see that, for all $i \notin C_k$, the equation $0 = [\mathbf{A} - \mathbf{Z}]_{ii} = \|\mathbf{F}_i\|^2 = \|\mathbf{L}_i\mathbf{G}\|^2$ implies that each column of \mathbf{G} lies in S^\perp , the orthogonal complement of the linear subspace S . Thus, we have $\mathbf{Z} - \mathbf{B} = \mathbf{L}\mathbf{L}^T - \mathbf{F}\mathbf{F}^T - \widehat{\mathbf{L}}\widehat{\mathbf{L}}^T = \mathbf{L}(\mathbf{I}_r - \mathbf{G}\mathbf{G}^T - \mathbf{V}\mathbf{V}^T)\mathbf{L}^T$, where $\mathbf{I}_r - \mathbf{G}\mathbf{G}^T - \mathbf{V}\mathbf{V}^T \succeq \mathbf{O}$ by Lemma 1 with $(\mathbf{M}, \mathbf{N}) = (\mathbf{G}, \mathbf{V})$. \square

We are now ready to present Algorithm 2, based on the subroutine Algorithm 1, for computing a minimal element of $\mathcal{F}(\mathbf{A})$ for arbitrary $\mathbf{A} \in \mathbb{S}^n$. Note that Algorithm 1 requires a factorization of its input and provides a factorization of its output, which is useful for repeatedly calling Algorithm 1 within Algorithm 2. Theorem 1 establishes the correctness of Algorithm 2 via Lemma 2.

Algorithm 2 Calculate a minimal element of $\mathcal{F}(\mathbf{A})$

Input: $\mathbf{A} \in \mathbb{S}^n$.

Output: $\mathbf{B} \in \text{Minimal}(\mathcal{F}(\mathbf{A}))$.

- 1: Choose \mathcal{C} -block diagonal Δ such that $\mathbf{A} + \Delta \succeq \mathbf{O}$.
 - 2: Initialize $\mathbf{B}^0 := \mathbf{A} + \Delta$.
 - 3: **for** $k = 1, 2, \dots, r$ **do**
 - 4: Calculate $\mathbf{B}^k := \text{minimum}(\mathcal{F}_k(\mathbf{B}^{k-1}))$ via Algorithm 1.
 - 5: **end for**
 - 6: Set $\mathbf{B} := \mathbf{B}^r$.
-

Lemma 2. Let $\mathbf{A} \in \mathbb{S}^n$ and $k \in \{1, \dots, r\}$. If $\mathbf{B} \in \mathcal{F}(\mathbf{A})$, then $\mathcal{F}_k(\mathbf{B}) \subseteq \mathcal{F}(\mathbf{A})$.

Proof. Let $\mathbf{M} \in \mathcal{F}_k(\mathbf{B})$, i.e., $\mathbf{O} \preceq \mathbf{M} \preceq \mathbf{B}$ and $\mathbf{B} - \mathbf{M}$ is C_k -block diagonal. Since $\mathbf{B} \in \mathcal{F}(\mathbf{A})$, we also know $\mathbf{A} - \mathbf{B}$ is \mathcal{C} -block diagonal. Hence $\mathbf{B} \succeq \mathbf{O}$ and $\mathbf{A} - \mathbf{M} = (\mathbf{A} - \mathbf{B}) + (\mathbf{B} - \mathbf{M})$ is \mathcal{C} -block diagonal. So $\mathbf{M} \in \mathcal{F}(\mathbf{A})$. \square

Theorem 1. Algorithm 2 correctly calculates $\mathbf{B} \in \text{Minimal}(\mathcal{F}(\mathbf{A}))$.

Proof. We first argue that $\mathbf{B} \in \mathcal{F}(\mathbf{A} + \Delta)$. Let $\mathbf{A} + \Delta =: \mathbf{B}^0, \mathbf{B}^1, \dots, \mathbf{B}^r =: \mathbf{B}$ be the sequence generated by Algorithm 2. Clearly $\mathbf{B}^0 \in \mathcal{F}(\mathbf{A} + \Delta)$. For induction, assume $\mathbf{B}^{k-1} \in \mathcal{F}(\mathbf{A} + \Delta)$. Then $\mathbf{B}^k \in \mathcal{F}_k(\mathbf{B}^{k-1}) \subseteq \mathcal{F}(\mathbf{A} + \Delta)$ by Lemma 2. So $\mathbf{B}^r \in \mathcal{F}(\mathbf{A} + \Delta)$.

Next we prove that $\mathcal{F}_k(\mathbf{B}) = \{\mathbf{B}\}$ for all $k = 1, \dots, r$. Fix k . If $\mathbf{B} = \mathbf{O}$ then the assertion holds easily. For $\mathbf{B}^r =: \mathbf{B} \neq \mathbf{O}$, we assume on the contrary the existence of $\mathbf{Z} \in \mathcal{F}_k(\mathbf{B}^r)$ with $\mathbf{Z} \neq \mathbf{B}^r$. Define $\tilde{\mathbf{B}} := \mathbf{B}^k - (\mathbf{B}^r - \mathbf{Z})$. Then

$$\mathbf{B}^{k-1} \succeq \mathbf{B}^k \succeq \tilde{\mathbf{B}} = (\mathbf{B}^k - \mathbf{B}^{k+1}) + (\mathbf{B}^{k+1} - \mathbf{B}^{k+2}) + \dots + (\mathbf{B}^{r-1} - \mathbf{B}^r) + \mathbf{Z} \succeq \mathbf{O}$$

and the fact that $\mathbf{B}^r - \mathbf{Z} = \mathbf{B}^k - \tilde{\mathbf{B}}$ is C_k -block diagonal implies also that $\mathbf{B}^{k-1} - \tilde{\mathbf{B}} = (\mathbf{B}^{k-1} - \mathbf{B}^k) + (\mathbf{B}^k - \tilde{\mathbf{B}})$ is C_k -block diagonal. Hence, $\tilde{\mathbf{B}} \in \mathcal{F}_k(\mathbf{B}^{k-1})$ and $\mathbf{B}^k \neq \tilde{\mathbf{B}} \preceq \mathbf{B}^k$, but this contradicts the fact that $\mathbf{B}^k = \text{minimum}(\mathcal{F}_k(\mathbf{B}^{k-1}))$.

We are now ready to prove $\mathbf{B} \in \text{Minimal}(\mathcal{F}(\mathbf{A} + \Delta))$. Assume on the contrary that there exists $\tilde{\mathbf{B}} \in \mathcal{F}(\mathbf{A} + \Delta)$ such that $\mathbf{B} \neq \tilde{\mathbf{B}} \preceq \mathbf{B}$. Let $\mathbf{D} := \mathbf{B} - \tilde{\mathbf{B}}$. Then \mathbf{D} is nonzero, positive semidefinite, and \mathcal{C} -block diagonal. So $\mathbf{D}_{C_k C_k} \succeq \mathbf{O}$ is nonzero for some k . Let \mathbf{D}_k be the C_k -block diagonal matrix with block $\mathbf{D}_{C_k C_k}$, and define $\mathbf{Z} := \mathbf{B} - \mathbf{D}_k$. Then we see that $\mathbf{Z} \in \mathcal{F}_k(\mathbf{B})$, and so $\mathbf{Z} = \mathbf{B}$ by the previous paragraph. However, this is a contradiction.

Finally, we know $\mathbf{B} \in \text{Minimal}(\mathcal{F}(\mathbf{A} + \Delta)) = \text{Minimal}(\mathcal{F}(\mathbf{A}))$ by Proposition 5. \square

3.4 Our practical choice of matrices \mathbf{B}_i

Section 3.2 argued that each \mathbf{B}_i should be a minimal member of $\mathcal{F}(\mathbf{A}_i)$ in order to improve the bound gotten from the mixed SOCP-SDP relaxation (9), and Section 3.3 presented Algorithm 2 to compute a member of $\text{Minimal}(\mathcal{F}(\mathbf{A}_i))$. We will see in Section 4 that, as intended, Algorithm 2 computes \mathbf{B}_i efficiently.

Given input \mathbf{A}_i , Algorithm 2 relies in step 1 on the choice of a \mathcal{C} -block diagonal shift Δ_i such that $\mathbf{A}_i + \Delta_i \succeq \mathbf{O}$, so we now discuss choices for Δ_i . Of course, there are many possible choices, but it is important that the choice be made quickly so that the overall time to construct and solve (9) is less than directly solving the SDP (2). We suggest and study two choices of Δ_i , both of which are based on the spectral decomposition, which is reasonably quick to compute for the sizes of problems that we consider in Section 4.

Before stating our choices of Δ_i , we first need some definitions. Given any symmetric matrix \mathbf{M} , define $\rho(\mathbf{M}) := -\lambda_{\min}[\mathbf{M}]$ so that $\mathbf{M} + \rho(\mathbf{M})\mathbf{I} \succeq \mathbf{O}$. Also decompose \mathbf{M} into the sum of two matrices $\mathbf{M}(\mathcal{C})$ and $\overline{\mathbf{M}}(\mathcal{C}) := \mathbf{M} - \mathbf{M}(\mathcal{C})$ such that $\mathbf{M}(\mathcal{C})$ is \mathcal{C} -block diagonal and $\overline{\mathbf{M}}(\mathcal{C})$ has nonzeros only in the complementary positions. In other words, the blocks of $\mathbf{M}(\mathcal{C})$ are precisely $\mathbf{M}_{C_k C_k}$ ($k = 1, \dots, r$).

Our first choice for Δ_i takes diagonal $\Delta_i := \rho(\mathbf{A}_i)\mathbf{I}$ and applies Algorithm 2 to $\mathbf{A}_i + \Delta_i \succeq \mathbf{O}$. This choice was also discussed at the beginning of Section 3. We call it the *first shift*.

Our second choice is motivated by the observation that, in a certain sense, the choice of $\mathbf{B}_i \in \text{Minimal}(\mathcal{F}(\mathbf{A}_i))$ should not depend on the \mathcal{C} -block diagonal portion $\mathbf{A}_i(\mathcal{C})$ of \mathbf{A}_i ,

where $\mathbf{A}_i = \mathbf{A}_i(\mathcal{C}) + \overline{\mathbf{A}}_i(\mathcal{C})$ as defined above. This is because the set $\mathcal{F}(\mathbf{A}_i)$ may be defined equivalently as $\{\mathbf{M} \succeq \mathbf{O} : \overline{\mathbf{M}}(\mathcal{C}) = \overline{\mathbf{A}}_i(\mathcal{C})\}$, i.e., it may be defined in a way that does not depend on $\mathbf{A}_i(\mathcal{C})$. Proposition 5 further supports this observation because $\mathcal{F}(\mathbf{A}_i) = \mathcal{F}(\overline{\mathbf{A}}_i(\mathcal{C}))$ as \mathbf{A}_i and $\overline{\mathbf{A}}_i(\mathcal{C})$ differ by $\mathbf{A}_i(\mathcal{C})$. So, loosely speaking, our second choice of shift is to apply the first shift to $\overline{\mathbf{A}}_i(\mathcal{C})$. More precisely, define

$$\Delta_i := \rho(\overline{\mathbf{A}}_i(\mathcal{C})) I - \mathbf{A}_i(\mathcal{C})$$

so that $\mathbf{A}_i + \Delta_i = \overline{\mathbf{A}}_i(\mathcal{C}) + \rho(\overline{\mathbf{A}}_i(\mathcal{C})) I \succeq \mathbf{O}$. We call this the *second shift*.

4 Computational Results

In this section, we describe our computational experience with the framework introduced and described in Section 3. Our first goal is to verify that our technique—including preprocessing time—can speed up the solution of the mixed SOCP-SDP (9) compared to solving the SDP (2). Of course, the resultant bound is weaker, and so our second goal is to quantify this bound loss. We test the framework on a set of interesting instances from the literature, and we also test the two types of shifts (*first* and *second*) introduced in Section 3.4.

One important detail when applying the framework to a particular instance is the choice of partition $\mathcal{C} := \{C_1, \dots, C_r\}$ of $[n]$. Clearly the precise structure of the partition will have a big impact on the solution time of (9) as well as the resultant bound. As we describe in the next subsection, we take a straightforward approach for choosing \mathcal{C} , one which depends only on the value of n of an instance. Although this approach does not take full advantage of the data $(\mathbf{A}_i, \mathbf{a}_i, \alpha_i)$, we feel this approach allows us to study the basic features of our framework.

4.1 The instances and relaxations

We collected a total of 400 instances of (1) from the literature, which consisted of three groups: (i) 199 instances of the maximum cut (MaxCut) problem coming from [8] and [16] (21 instances of the *Gset* library and 178 instances of the *BiqMac* library, respectively); (ii) 64 instances of binary quadratic programming (BinQP) coming from the *BiqMac* library [16]; and (iii) 36 instances from GlobalLib [5] having bounded feasible sets. In particular, all instances had between $n = 16$ and $n = 800$ variables. The instance sizes (number of variables and constraints) are depicted in Figure 1 on a log-log scale. We note that all instances were formulated precisely in the standard form of (1). For example, quadratic equations were

split into two quadratic inequalities. In addition, we made sure to bound the diagonal of X as shown in (9).

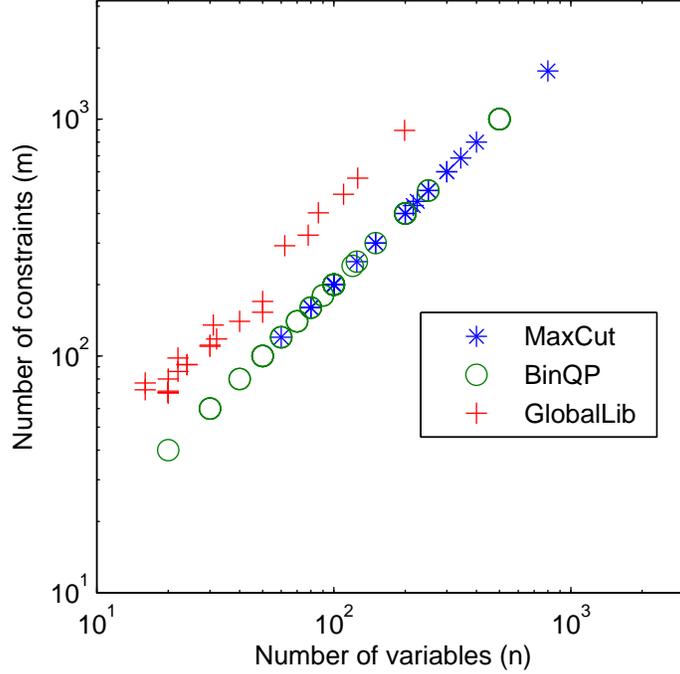


Figure 1: The sizes of the 400 MaxCut, BinQP, and GlobalLib instances. Some single points depict multiple instances of the same size.

To test the framework, we next created partitions \mathcal{C} as utilized in Section 3. Specifically, for each QCQP instance, we created and tested 4 different partitions. The first partition consists simply of all variables in a single set $C_1 = \{1, \dots, n\}$, which yields an instance of (9) that is equivalent to (2); this is our “base case” partition. The second partition is a refinement of the first gotten by (approximately) halving the first partition:

$$C_1 = \{1, \dots, \lceil \frac{n}{2} \rceil\}, C_2 = \{\lceil \frac{n}{2} \rceil + 1, \dots, n\}.$$

Likewise, the third partition approximately halves the second, and the fourth approximately halves the fourth. For example, if $n = 16$ for a QCQP instance, then we would create the 4

partitions

$$\begin{aligned}
C_1 &= \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16\}; \\
C_1 &= \{1, 2, 3, 4, 5, 6, 7, 8\}, C_2 = \{9, 10, 11, 12, 13, 14, 15, 16\}; \\
C_1 &= \{1, 2, 3, 4\}, C_2 = \{5, 6, 7, 8\}, C_3 = \{9, 10, 11, 12\}, C_4 = \{13, 14, 15, 16\}; \\
C_1 &= \{1, 2\}, C_2 = \{3, 4\}, C_3 = \{5, 6\}, C_4 = \{7, 8\}, C_5 = \{9, 10\}, C_6 = \{11, 12\}, \\
C_7 &= \{13, 14\}, C_8 = \{15, 16\}.
\end{aligned}$$

For each QCQP, every partition gives an instance of (9). So our 400 chosen instances of (1) yield a total of 1,600 instances of (9). While our choice of partitions is somewhat arbitrary, it is designed to provide a variety of sizes of partitions that we investigate. It would be interesting to choose the partition using, for example, a clever heuristic to improve the subsequent bound, but we do not do so here.

4.2 Algorithm variants and implementation

The 1,600 instances of (9) described in the previous subsection are solved by several algorithm variants that we describe now. The variants are based on two design choices, each with two options, yielding a total of four algorithm variants.

The first design choice is the type of shift employed, either *first* or *second*, as described in Section 3.4. The second design choice is whether to continue running Algorithm 2 after the shift in step 1 is made. In other words, we can terminate Algorithm 2 immediately after step 1 and simply use $\mathbf{B}_i := \mathbf{A}_i + \Delta_i$ to create the SOCP-SDP relaxation (9). We denote the four variants as *1N* (first shift without continuing Algorithm 2), *1Y* (first shift with the full Algorithm 2), *2N* (second shift without continuing Algorithm 2), and *2Y* (second shift with the full Algorithm 2). The letters *N* and *Y* are meant to indicate “no” and “yes” for the full Algorithm 2.

The purpose of these four variants is to isolate the effects of different aspects of our framework. For example, by comparing 1N with 1Y (or 2N with 2Y), we can determine if calculating a minimal element $\mathbf{B}_i \in \text{Minimal}(\mathcal{F}(\mathbf{A}_i))$ has added benefits over simply constructing the relaxation (9) with a \mathcal{C} -block structure after the shift. Further, comparing 1N with 2N (or 1Y with 2Y), we can determine the different effects of the two types of shifts. Finally, we may also hope to find the best overall choice of the four algorithm variants.

For each of the four variants and any of the 1,600 instances, the mixed SOCP-SDP (9) is converted to standard block-diagonal form derived from the \mathcal{C} -block structure and Proposition 3. This involves adding just a few extra variables and constraints (at most n).

The resultant problems are solved with a default installation of SeDuMi on an Ubuntu Linux computer having an Intel Core 2 Quad CPU running at 2.4 GHz with 4 MB cache and 4 GB RAM. Even though the CPU has multiple cores, we limit Matlab to using at most one core.

4.3 Comparisons

We now compare the four different algorithm variants on the 400 instances of (2), each of which gives rise to four instances of (9) based on different partitions $\mathcal{C} := \{C_1, \dots, C_r\}$. Recall that, for each instance, the four partitions have $r = 1, 2, 4, 8$ blocks, respectively. When $r = 1$, we have the “base case,” which is equivalent to solving the SDP relaxation (2).

Consider a single instance of (2) solved by a single variant (1N, 1Y, 2N, or 2Y) for the four block values $r = 1, 2, 4, 8$. This gives rise to four lower bounds $\{b_r : r = 1, 2, 4, 8\}$ on the optimal value of (2) and four computation times $\{t_r : r = 1, 2, 4, 8\}$. In particular, t_r is the total time to apply our framework including calculating the shift for each i , applying the full Algorithm 2 for each i (as required), and setting up and solving (9). We compare our framework to the base-case SDP by calculating the six ratios

$$\beta_r := \frac{b_r}{b_1} \quad (r = 2, 4, 8) \quad \text{and} \quad \tau_r := \frac{t_r}{t_1} \quad (r = 2, 4, 8).$$

(Note that all instances had $b_1 \leq -0.01$ so that the denominator defining β_r was sufficiently far from 0 and so that each $b_r \leq -0.01$ and each $\beta_r \geq 1$.) By examining these ratios β_r and τ_r , we can compare our framework’s bounds and times to the SDP bounds and times on a standardized, relative scale. For example, if $\beta_r = 1.05$ and $\tau_r = 0.90$, it means that our framework degraded the bound 5% but took 10% less time.

For each combination of an algorithm variant and a block value $r \in \{2, 4, 8\}$, Figure 2 displays a box plot of the bound ratios β_r gotten over all 400 instances of (2). In a similar fashion, Figure 3 shows the time ratios τ_r . The plots were created using Matlab, and we point out the visual features that explain how to read the plots. For each box plot, the blue box spans the 25th and 75th percentiles of the data, and the red line within the blue box indicates the median (50th percentile). The vertical, black, dashed lines that extend from the blue box and terminate with horizontal, black “whiskers” show the rest of the data, which is not considered to be outliers. According to Matlab’s defaults, a data point is considered to be an outlier if it is smaller than $p_{25} - 1.5(p_{75} - p_{25})$ or larger than $p_{75} + 1.5(p_{75} - p_{25})$, where p_{25} and p_{75} are the 25th and 75th percentiles, i.e., edges of the blue box. Outliers are indicted by red plus signs. In addition, Figure 2 also uses Matlab’s “extreme mode,” which collapses the outliers beyond a certain point onto a single, dashed horizontal line to save space.

We first discuss the bound ratios in Figure 2 and make several observations:

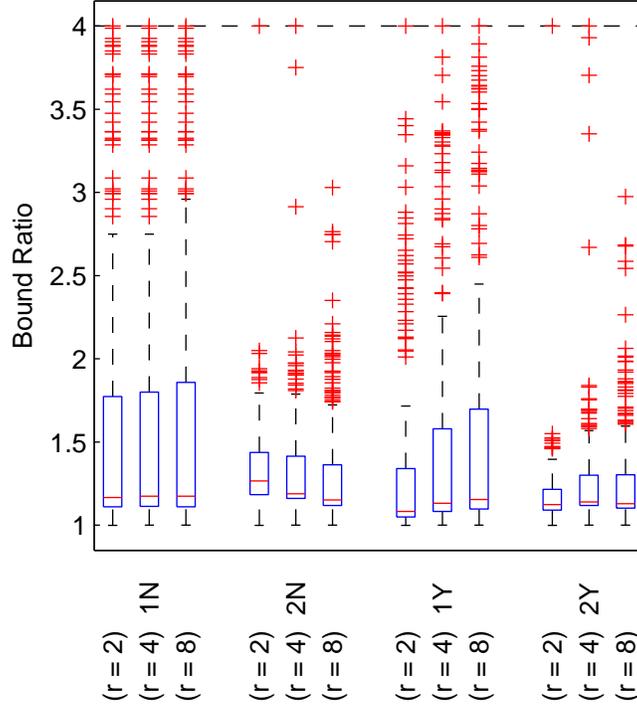


Figure 2: Box plots for the bound ratios β_r over all pairs of algorithm variants (1N, 2N, 1Y, 2Y) and block values ($r = 2, 4, 8$).

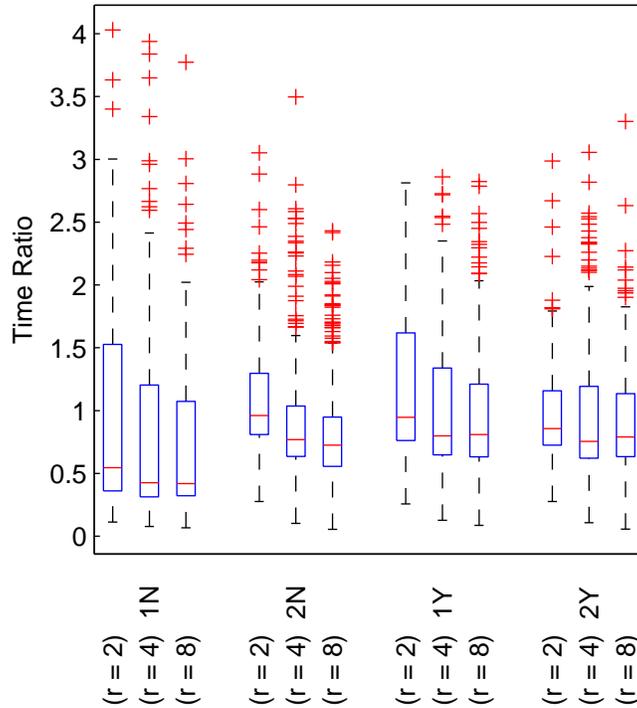


Figure 3: Box plots for the time ratios β_r over all pairs of algorithm variants (1N, 2N, 1Y, 2Y) and block values ($r = 2, 4, 8$).

- For each of the variants 1N and 1Y, we see that the bound ratio generally worsens as r increases, while for 2N and 2Y, the bound ratio generally improves or stays the same. This indicates that, irrespective of the full use of Algorithm 2, the second shift is better for improving or maintaining the bound quality as r increases.

Actually, for 1N, the decreasing bounds are expected because the $(r + 1)$ -st relaxation is itself a relaxation of the r -th relaxation. This is true since the \mathcal{C} for $r + 1$ is a finer partition than the \mathcal{C} for r and since the first shift is not dependent on the partition \mathcal{C} . In contrast, for 2N, the second shift is able to counteract the expected loss of bound due to the finer partition.

- Note, however, that for a fixed r , neither shift dominates the other. For example, when $r = 2$, 1N provides better bounds than 2N, but when $r = 8$, the bounds provided by 1N are worse.
- For each value of r , the bound ratios for 1N are worse than for 1Y. Likewise, for each value of r , the bound ratios for 2N are worse than for 2Y. This indicates that, irrespective of the type of shift, the full use of Algorithm 2 keeps the bound quality closer to that of the SDP relaxation.
- Of the four variants, 2Y (second shift with full Algorithm 2) keeps the bounds consistently low for all values of r . As such, it seems to be the best performing variant overall.

We next discuss the time ratios in Figure 3:

- Generally speaking, for each variant, increasing r results in relaxations that are faster to solve, and the median time ratios are well below the value of 1, which means that our framework is generally faster than solving the SDP directly.
- However, there is considerable variation with many individual ratios above 1. The smallest ratios are near 0.5, meaning that our framework is, roughly speaking, at most twice as fast as the SDP. In particular, we do not see a full order-of-magnitude speed up.
- Variant 1N is clearly the fastest variant, while the other three variants require roughly the same amount of time.

In both figures, we also see an interesting trend regarding the variation of the data as depicted by the heights of the blue boxes. Variants 2N and 2Y clearly exhibit less variation than 1N

and 1Y. The second shift appears to be largely responsible for this decrease in variation, but even in Figure 2 for the bounds, the variation in 2Y is a bit less than the variation in 2N, which shows a slight improvement due to the full Algorithm 2.

While the comparisons just made using Figures 2 and 3 are imperfect—especially due to the overall variation seen in the bound and time ratios due to the outliers—the general trends indicate that our framework can provide faster relaxations but with a corresponding loss in bound quality. It appears that variant 2Y, which uses the second shift and Algorithm 2, achieves the best bound quality relative to the SDP relaxation, while taking about the same amount of time as variants 2N and 1Y. (Variant 1N is noticeably faster, but its bounds are less reliable.)

5 Conclusions

When attempting to solve difficult QCQPs, the bounds provided by SDP relaxations can be tight and hence quite useful, but they can also be time consuming to calculate. This paper has presented a framework for constructing mixed SOCP-SDPs that provide faster, but weaker, bounds. Our framework is unique compared to other related approaches in the literature as it allows one to control the solution speed of the SOCP-SDP (via its block structure) while simultaneously working to improve the bound quality via the idea of minimum and minimal elements, which respect the block structure. We found that the combination of the second shift and the full Algorithm 2 (variant 2Y) performed the best overall.

There are many avenues to improve our framework. In Section 4, we tested simplistic choices for the partition $\mathcal{C} := \{C_1, \dots, C_r\}$ that do not take into account the actual data $\{(\mathbf{A}_i, \mathbf{a}_i, \alpha_i)\}$. We did so with the intent of just testing the basic behavior of our framework, but it would be very interesting to design heuristics that choose \mathcal{C} intelligently to preserve the bound quality even beyond the second shift and the full use of Algorithm 2.

Another way to extend our approach is to allow \mathcal{C} to be a covering of $[n]$ rather than just a partition, that is, to allow the elements C_k to overlap. As long as we can extend Proposition 3 to this case, the framework will extend easily, and extending Proposition 3 relies in turn on the chordal-graph structure of the matrix

$$\begin{pmatrix} 1 & \mathbf{x}^T \\ \mathbf{x} & \mathbf{Y} \end{pmatrix},$$

where $\mathbf{Y}_{C_k C_k} = \mathbf{X}_{C_k C_k}$ for all $k = 1, \dots, r$ and zero otherwise (see again [7]). The flexibility of choosing overlapping C_k should allow further preservation of the bound without dramatic

increases in the computation time.

References

- [1] S. Burer and R. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming (Series B)*, 95:329–357, 2003.
- [2] E. de Klerk and R. Sotirov. Exploiting group symmetry in semidefinite programming relaxations of the quadratic assignment problem. *Math. Programming*, 122(2, Ser. A):225–246, 2010.
- [3] T. Fujie and M. Kojima. Semidefinite programming relaxation for nonconvex quadratic programs. *J. Global Optim.*, 10(4):367–380, 1997.
- [4] M. Fukuda, M. Kojima, K. Murota, and K. Nakata. Exploiting sparsity in semidefinite programming via matrix completion. I. General framework. *SIAM J. Optim.*, 11(3):647–674, 2000/01.
- [5] See the website: www.gamsworld.org/global/globallib/globalstat.htm.
- [6] M. Goemans and D. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of ACM*, 42:1115–1145, 1995.
- [7] R. Grone, C. Johnson, E. Sá, and H. Wolkowicz. Positive definite completions of partial hermitian matrices. *Linear Algebra Applications*, 58:109–124, 1984.
- [8] C. Helmberg and F. Rendl. A spectral bundle method for semidefinite programming. *SIAM Journal on Optimization*, 10:673–696, 2000.
- [9] S. Kim and M. Kojima. Second order cone programming relaxation of nonconvex quadratic optimization problems. *Optim. Methods Softw.*, 15(3-4):201–224, 2001.
- [10] M. Kojima and L. Tunçel. Cones of matrices and successive convex relaxations of nonconvex sets. *SIAM J. Optim.*, 10(3):750–778, 2000.
- [11] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optim.*, 11(3):796–817, 2001.
- [12] M. Laurent and F. Rendl. Semidefinite programming and integer programming. In R. W. Karen Aardal, George Nemhauser, editor, *Handbook on Discrete Optimization*, pages 393–514. Elsevier B.V., 2005.
- [13] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1:166–190, 1991.

- [14] R. D. C. Monteiro. First- and second-order methods for semidefinite programming. *Mathematical Programming*, 97:209–244, 2003.
- [15] K. Nakata, K. Fujisawa, M. Fukuda, M. Kojima, and K. Murota. Exploiting sparsity in semidefinite programming via matrix completion. II. Implementation and numerical results. *Math. Program.*, 95(2, Ser. B):303–327, 2003. Computational semidefinite and second order cone programming: the state of the art.
- [16] F. Rendl, G. Rinaldi, and A. Wiegele. Solving Max-Cut to optimality by intersecting semidefinite and polyhedral relaxations. *Math. Programming*, 121(2):307, 2010.
- [17] A. Saxena, P. Bonami, and J. Lee. Convex relaxations of non-convex mixed integer quadratically constrained programs: projected formulations. *Mathematical Programming*, 2010. DOI 10.1007/s10107-010-0340-3.
- [18] H. D. Sherali and W. P. Adams. *A Reformulation-Linearization Technique for Solving Discrete and Continuous Nonconvex Problems*. Kluwer, 1997.
- [19] X. Zhao, D. Sun, and K. Toh. A Newton-CG augmented Lagrangian method for semidefinite programming. *SIAM J. Optim.*, 20:1737–1765, 2010.
- [20] X. Zheng, X. Sun, and D. Li. Nonconvex quadratically constrained quadratic programming: best d.c. decompositions and their sdp representations. *Journal of Global Optimization*, 50:695–712, 2011.