Lagrangian Duality for Mixed-Integer Semidefinite Programming: Theory and Algorithms

Frank de Meijer [*]   Renata Sotirov [†]

January 20, 2025

Abstract

This paper presents the Lagrangian duality theory for mixed-integer semidefinite programming (MISDP). We derive the Lagrangian dual problem and prove that the resulting Lagrangian dual bound dominates the bound obtained from the continuous relaxation of the MISDP problem. We present a hierarchy of Lagrangian dual bounds by exploiting the theory of integer positive semidefinite matrices and propose three algorithms for obtaining those bounds. Our algorithms are variants of well-known algorithms for minimizing non-differentiable convex functions. The numerical results on the max-k-cut problem show that the Lagrangian dual bounds are substantially stronger than the semidefinite programming bound obtained by relaxing integrality, already for lower levels in the hierarchy. Computational costs for computing our bounds are small.

Keywords: mixed-integer semidefinite programming, discrete positive semidefinite matrices, Lagrangian duality theory, max-k-cut problem

1 Introduction

Mixed-integer semidefinite programming (MISDP) may be seen as a generalization of mixed-integer linear programming (MILP), where the vector of variables is replaced by a positive semidefinite matrix variable in which some of the entries are integer and others real. In integer semidefinite programming (ISDP) all entries of the positive semidefinite matrix variable are required to be integer.

The combination of positive semidefiniteness and integrality allows to formulate various optimization problems as mixed-integer semidefinite programs (MISDPs). Since the 1990s, a few papers appeared that present ISDP formulations for classical discrete optimization problems e.g., [8, 11, 42], and several MISDP formulations of applied optimization problems, see e.g., [20, 58, 59]. In [40], the authors introduce a generic approach for deriving MISDP formulations of binary quadratically constrained quadratic programs and binary quadratic matrix programs. They also show that several optimization problems allow for novel MISDP formulations, which provide new perspectives on solution approaches for those problems.

The MISDPs can be used to compute exact solutions for the optimization problems using generic MISDP solvers, e.g., branch-and-bound algorithms [17, 36] or branch-and-cut algorithms [39, 30]. Another solver that supports solving MISDPs while exploiting sparsity is GravitySDP [23]. Numerical tests in those and related papers show that solving an MISDP formulation of an optimization problem might be more beneficial than solving a mixed-integer programming (MIP) formulation of the same problem via state-of-the-art MIP solvers.

A crucial feature of an efficient solver is being able to compute strong bounds fast. Most of the current solving approaches for MISDPs solve SDP relaxations in each node of a branching tree. Since a generic way of solving an SDP relaxation is an interior point method, computational effort in each node of a branching tree may be (very) high. It is well-known that interior point-based SDP solvers exhibit problems in terms of both time and memory for solving even medium-size semidefinite programs (SDPs). Moreover, interior point methods have difficulties with handling additional cutting planes, including nonnegativity constraints that significantly strengthen SDP bounds. There exist alternative approaches for solving SDP relaxations based on Alternating Direction Augmented Lagrangian (ADAL) methods, see e.g., [25, 56, 46]. Those algorithms are first-order methods that have low memory requirement, but

---

[*]Delft Institute of Applied Mathematics, Delft University of Technology, The Netherlands, `f.j.j.demeijer@tudelft.nl`

[†]CentER, Department of Econometrics and OR, Tilburg University, The Netherlands, `r.sotirov@uvt.nl`

might suffer a tailing-off effect when there is need for a high precision solution. However, ADAL methods can handle a large number of cutting planes together with semidefiniteness efficiently [37, 38, 41]. The bounds resulting from relaxing a MISDP to an SDP do not exploit integrality of the variables. The goal of this work is to provide alternative bounding approaches that do exploit both positive semidefiniteness and integrality.

In this paper we derive the Lagrangian duality theory for MISDP that enables us to derive a hiearchy of Lagrangian dual bounds, and show how to compute those bounds using first-order methods. There exist results on superadditive duality theory for mixed-integer conic optimization [31]. However, the subadditive dual for conic mixed-integer programs does not yield straightforward solution procedures, while the Lagrangian duality theory for MISDP does.

Main results and outline

We first introduce the Lagrangian duality theory for MISDP. We derive a Lagrangian relaxation of the MISDP problem by dualizing the positive semidefinite (PSD) constraint and a part of the linear constraints. The aim is to dualize constraints that are intractable in combination with the integrality constraints. From the Lagrangian relaxation we derive the Lagrangian dual problem. We relate the optimal value of the continuous relaxation of the MISDP problem, the Lagrangian dual bound, and the optimal value of the MISDP problem within a sandwich theorem, showing that the Lagrangian dual bound is always at least as strong as the bound obtained from the continuous SDP relaxation. We also derive conditions under which equalities throughout the sequence of the sandwich theorem follow.

For the case of purely integer SDPs, we introduce a hierarchy of Lagrangian dual bounds by exploiting the theory of integer PSD matrices. Namely, we propose partitioning a matrix variable into submatrices of prescribed sizes and solving the Lagrangian dual problem by exploiting finite generating sets of integer PSD matrices. The larger the size of the matrices in the partitions is, the stronger the Lagrangian dual bounds become. The resulting bounds are related to the so-called exact subgraph approach considered in [1, 16]. In fact, exact subgraph bounds are Lagrangian dual bounds resulting from appropriate ISDP models.

We propose three algorithms for computing the Lagrangian dual bounds. We consider a variant of the deflected subgradient algorithm with the update scheme from [52] and Polyak's stepsize [47]. We also consider a version of the accelerated gradient method introduced by Nesterov [45], and a version of the proximal bundle algorithm, see e.g., [29]. All three variants include projections onto the PSD cone to preserve dual feasibility in each iteration of the algorithm, which distinguishes them from their classical versions. The subgradient methods were not used, up to date, for computing SDP bounds. Variants of the proximal bundle method have been used for obtaining SDP bounds of several optimization problems, see e.g., [13, 16, 49]. The Lagrangian relaxations in these existing algorithms require solving an SDP relaxation in each bundle iteration. Instead, we compute the Lagrangian relaxations over a discrete set.

We evaluate our algorithms on the max-k-cut problem for $k \in \{2, 3, 4\}$. The Lagrangian dual bounds turn out to be stronger than the SDP bounds obtained from the continuous relaxations of MISDPs, already for lower levels in the hierarchy, however the bounds keep on improving for higher levels in the hierarchy. Computational costs for computing our strong bounds are small.

This paper is structured as follows. Lagrangian duality theory for MISDPs is introduced in Section 2. A hierarchy of Lagrangian dual bounds for ISDPs is derived in Section 3. In Section 4 we present three approaches for solving the Lagrangian dual problem. In particular, in Section 4.1 we present a projected-deflected subgradient method, in Section 4.2 a projected-accelerated subgradient method, and in Section 4.3 a projected bundle method. In Section 5 we present an ISDP formulation of the max-k-cut problem, after which numerical results on this problem are provided in Section 6. Section 7 concludes our work.

Notation

For $n \in \mathbb{Z}_+$, we define the set $[n] := \{1, \ldots, n\}$. We denote by $\mathbf{0}_n \in \mathbb{R}^n$ the vector of all zeros, and by $\mathbf{1}_n \in \mathbb{R}^n$ the vector of all ones. The matrix of all-ones of size $n \times n$ is denoted by $\mathbf{J}_n$. We omit the subscripts of these matrices when there is no confusion about the size. The Euclidean norm of vector x is denoted by $||x||$. For any matrix $X \in \mathbb{R}^{n \times n}$ and indicator set $S \subseteq [n]$, we let $X[S]$ denote the principal submatrix induced by S.

We denote the set of all $n \times n$ real symmetric matrices by $\mathcal{S}^n$. The cone of symmetric positive semidefinite matrices is defined as $\mathcal{S}^n_+ := \{X \in \mathcal{S}^n : X \succeq \mathbf{0}\}$. Thus, $X \succeq \mathbf{0}$ means that X is a positive semidefinite matrix. The trace of a square matrix $X = (X_{ij})$ is given by $\mathrm{tr}(X) = \sum_i X_{ii}$. For any $X, Y \in \mathcal{S}^n$ the trace inner product is defined as $\langle X, Y \rangle := \mathrm{tr}(XY) = \sum_{i,j=1}^n X_{ij} Y_{ij}$. The associated norm is the Frobenius norm, denoted by $||X||_F := \sqrt{\mathrm{tr}(X^\top X)}$.

The operator $\mathrm{diag} : \mathbb{R}^{n \times n} \to \mathbb{R}^n$ maps a square matrix to a vector consisting of its diagonal elements. We denote by $\mathrm{Diag} : \mathbb{R}^n \to \mathbb{R}^{n \times n}$ its adjoint operator.

2 Lagrangian duality theory for MISDP

In this section we extend the Lagrangian duality theory from mixed-integer linear programming to the case of mixed-integer semidefinite programming.

Let $C \in \mathcal{S}^n$, $b \in \mathbb{R}^m$ and let $\mathcal{A} : \mathcal{S}^n \to \mathbb{R}^m$ be a linear operator that is defined by $\mathcal{A}(X)_i := \langle A_i, X \rangle$ with $A_i \in \mathcal{S}^n$ for all $i \in [m]$. We define $\mathcal{A}^* : \mathbb{R}^m \to \mathcal{S}^n$ to be its adjoint. Moreover, we let $\mathcal{J} \subseteq [n] \times [n]$ be an index set of the integer variables in the program. For each $(i,j) \in \mathcal{J}$, the set $B_{ij} \subseteq \mathbb{Z}$ denotes the integer solution space of the variable indexed by (i,j). Finally, for all $(i,j) \in \mathcal{J}$, let $l_{ij} \in \mathbb{Z}$ and $u_{ij} \in \mathbb{Z}$ denote the lower and upper bound, respectively, with respect to the set B_{ij}. We assume these lower and upper bounds to be finite, implying that the sets B_{ij} are finite. We consider a MISDP problem in the following general form:

$$z_{MISDP} := \quad \min \quad \langle C, X \rangle$$
$$\text{s.t.} \quad \mathcal{A}(X) = b, \ X \succeq \mathbf{0} \qquad\qquad (MISDP)$$
$$X_{ij} \in B_{ij} \qquad \text{for all } (i,j) \in \mathcal{J}.$$

By relaxing the integrality constraints, we obtain the continuous SDP relaxation of $(MISDP)$:

$$z_{SDP} := \quad \min \quad \langle C, X \rangle$$
$$\text{s.t.} \quad \mathcal{A}(X) = b, \ X \succeq \mathbf{0} \qquad\qquad (SDP)$$
$$l_{ij} \leq X_{ij} \leq u_{ij} \qquad \text{for all } (i,j) \in \mathcal{J}.$$

The inequality $z_{SDP} \leq z_{MISDP}$ clearly holds. Throughout this section, we make the following assumption, which is natural for many MISDPs originating from discrete optimization [40].

Assumption 1. *The feasible set of $(MISDP)$ is bounded.*

Since all integer variables are bounded by l_{ij} and u_{ij}, respectively, Assumption 1 implies that bounds on the continuous variables are enforced by the constraints of $(MISDP)$. Observe that the boundedness assumption on $(MISDP)$ also implies that a solution to the continuous SDP is attained. Namely, if the feasible set of (SDP) would be unbounded, then there exists a ray $R \in \mathcal{S}^n$ with $\mathcal{A}(R) = \mathbf{0}$ and $R \succeq \mathbf{0}$. Since all bounds on variables in $\mathcal{J}$ are finite, $R_{ij} = 0$ for all $(i,j) \in \mathcal{J}$. Therefore, R would also be a ray of the feasible set of $(MISDP)$, contradicting Assumption 1.

The Lagrangian dual of $(MISDP)$ is obtained by dualizing the constraints that are intractable in combination with the integrality constraints. In our setting, it is natural to dualize the constraint $X \succeq \mathbf{0}$. Moreover, we can distinguish between the equalities in $\mathcal{A}(X) = b$ that are tractable with the integrality constraints and the equalities that are not. Consequently, we split $\mathcal{A}(X) = b$ into $\mathcal{A}_1(X) = b_1$ and $\mathcal{A}_2(X) = b_2$ where $b_1 \in \mathbb{R}^{m_1}$ and $b_2 \in \mathbb{R}^{m_2}$ with $m_1 + m_2 = m$. Here we assume that the equalities $\mathcal{A}_1(X) = b_1$ are not tractable in combination with the integrality constraints. We define the Lagrangian $\mathcal{L}(\cdot)$ after dualizing the constraints $X \succeq \mathbf{0}$ and $\mathcal{A}_1(X) = b_1$ as

$$\mathcal{L}(X, S, \lambda) := \langle C, X \rangle - \langle S, X \rangle + \lambda^\top (\mathcal{A}_1(X) - b_1),$$

where $S \succeq \mathbf{0}$ and $\lambda \in \mathbb{R}^{m_1}$ are the corresponding Lagrange multipliers. Moreover, we let P denote the set of mixed-integer symmetric matrices induced by the remaining constraints, i.e.,

$$P := \{X \in \mathcal{S}^n : \mathcal{A}_2(X) = b_2, \ X_{ij} \in B_{ij} \text{ for all } (i,j) \in \mathcal{J}\}. \qquad (1)$$

Without loss of generality, we may assume that P is bounded. Namely, if not, we can add to $\mathcal{A}_2(X) = b_2$ the variable bounds on the continuous variables (which exist due to Assumption 1). Now, we define the Lagrangian dual function $g : \mathcal{S}^n_+ \times \mathbb{R}^{m_1} \to \mathbb{R}$ as follows:

$$g(S, \lambda) := \min \{\mathcal{L}(X, S, \lambda) : X \in P\}. \qquad (LR(S, \lambda))$$

Obviously, for all $S \in \mathcal{S}_+^n$, $\lambda \in \mathbb{R}^{m_1}$ we have $g(S, \lambda) \leq \mathcal{L}(X^*, S, \lambda) \leq \langle C, X^* \rangle = z_{MISDP}$, where X^* is an optimal solution to $(MISDP)$. We call the minimization problem $(LR(S, \lambda))$ the *Lagrangian relaxation of $(MISDP)$ with parameters* (S, λ). To obtain the best lower bound for z_{MISDP}, we take the supremum of $g(S, \lambda)$ with respect to the dual variables S and λ. This leads to the *Lagrangian dual* of $(MISDP)$:

$$z_{LD} := \quad \sup \quad g(S, \lambda)$$
$$\text{s.t.} \quad S \succeq \mathbf{0}, \ \lambda \in \mathbb{R}^{m_1}. \tag{LD}$$

The following result follows by construction.

Proposition 1 (Weak duality). $z_{LD} \leq z_{MISDP}$

An optimal solution of the Lagrangian relaxation $(LR(S, \lambda))$ may be an optimal solution to $(MISDP)$, as stated in the following proposition.

Proposition 2. *Let $S \succeq \mathbf{0}$ and $\lambda \in \mathbb{R}^{m_1}$ be given. If $X^* \in \mathcal{S}^n$ is an optimal solution for $(LR(S, \lambda))$ that is feasible for $(MISDP)$ and satisfies the complementarity slackness condition $SX^* = 0$, then X^* is optimal for $(MISDP)$.*

In the sequel, we show that it is possible to obtain z_{LD} as the solution of a continuous semidefinite programming problem. We consider

$$\min \quad \langle C, X \rangle$$
$$\text{s.t.} \quad \mathcal{A}_1(X) = b_1, \ X \succeq \mathbf{0} \tag{2}$$
$$X \in \text{conv}(P).$$

Observe that P is defined as the set of mixed-integer points contained in a polyhedron. Since the feasible set of (2) is contained in the feasible set of (SDP) and the latter one is compact, it follows that an optimal solution to (2) is attained. We now show that the optimization problem (2) is equivalent to the Lagrangian dual of $(MISDP)$ based on a similar result for MILP by Geoffrion [18].

Theorem 1. *Let $\hat{z}$ denote the optimal objective value to (2), then $z_{LD} = \hat{z}$.*

Proof. Since $\mathcal{L}(\cdot, S, \lambda)$ is linear on $\mathcal{S}^n$ for all fixed $S \in \mathcal{S}_+^n$ and $\lambda \in \mathbb{R}^{m_1}$, we have

$$z_{LD} = \sup_{S \succeq \mathbf{0}, \lambda} \min_{X \in \mathcal{S}^n} \left\{ \langle C, X \rangle - \langle S, X \rangle + \lambda^\top (\mathcal{A}_1(X) - b_1) \ : \ X \in P \right\}$$
$$= \sup_{S \succeq \mathbf{0}, \lambda} \min_{X \in \mathcal{S}^n} \left\{ \langle C, X \rangle - \langle S, X \rangle + \lambda^\top (\mathcal{A}_1(X) - b_1) \ : \ X \in \text{conv}(P) \right\}. \tag{3}$$

Since $\mathcal{L}(X, \cdot, \cdot)$ is also linear on $\mathcal{S}_+^n \times \mathbb{R}^{m_1}$ for all fixed $X \in \text{conv}(P)$, and $\text{conv}(P)$ is compact, Sion's minimax theorem [55] implies that we may interchange the order of taking the minimum and the supremum, yielding

$$z_{LD} = \min_{X \in \mathcal{S}^n} \sup_{S \succeq \mathbf{0}, \lambda} \left\{ \langle C, X \rangle - \langle S, X \rangle + \lambda^\top (\mathcal{A}_1(X) - b_1) \ : \ X \in \text{conv}(P) \right\}$$
$$= \min_{X \in \mathcal{S}^n} \left\{ \langle C, X \rangle \ : \ X \in \text{conv}(P), \ \mathcal{A}_1(X) = b_1, \ X \succeq \mathbf{0} \right\},$$

where the final equality follows from the fact that if $\mathcal{A}_1(X) \neq b_1$ or $X \not\succeq \mathbf{0}$, then the inner supremum is unbounded. $\qquad \square$

The combination of the results from Proposition 1 and Theorem 1 leads to the following sandwich relation.

Corollary 1 (Sandwich theorem). *Under Assumption 1, $z_{SDP} \leq z_{LD} \leq z_{MISDP}$.*

Remark 1. Instead of Assumption 1, we can also establish the sandwich theorem of Corollary 1 by assuming that the problem (2) has a Slater feasible point and all data matrices are rational (while allowing P to be unbounded). In that case, P is the intersection of a mixed-integer linear set and a rational polyhedron, hence by Meyer's theorem [43] $\text{conv}(P)$ is a polyhedron. Let $\mathcal{A}_3 : \mathcal{S}^n \to \mathbb{R}^{m_3}$ and $b_3 \in \mathbb{R}^{m_3}$ be the linear operator and vector such that $\text{conv}(P) = \{ X \in \mathcal{S}^n \ : \ \mathcal{A}_3(X) \leq b_3 \}$, respectively. Then, (3) still holds and can be rewritten to

$$z_{LD} = \sup_{S \succeq \mathbf{0}, \lambda} -b_1^\top \lambda + \min_{X \in \mathcal{S}^n} \left\{ \langle C - S + \mathcal{A}_1^*(\lambda), X \rangle \ : \ \mathcal{A}_3(X) \leq b_3 \right\}.$$

Observe that the inner minimization problem is linear, so we replace it by its dual. Let $z \geq \mathbf{0}$ denote the dual variable corresponding to $\mathcal{A}_3(X) \leq b_3$, then:

$$
\begin{aligned}
&= \sup_{S \succeq \mathbf{0}, \lambda} -b_1^\top \lambda + \sup_{z \geq \mathbf{0}} \left\{ -b_3^\top z \ : \ S = C + \mathcal{A}_1^*(\lambda) + \mathcal{A}_3^*(z) \right\} \\
&= \sup_{S \succeq \mathbf{0}, z \geq \mathbf{0}, \lambda} \left\{ -b_1^\top \lambda - b_3^\top z \ : \ S = C + \mathcal{A}_1^*(\lambda) + \mathcal{A}_3^*(z) \right\} \\
&= \min_X \left\{ \langle C, X \rangle \ : \ X \in \mathrm{conv}(P), \ \mathcal{A}_1(X) = b_1, \ X \succeq \mathbf{0} \right\},
\end{aligned}
$$

where the last equality follows from the latter problem having a Slater feasible point. The sequence $z_{SDP} \leq z_{LD} \leq z_{MISDP}$ immediately follows. As boundedness is a more natural property in the MISDPs that are considered in later sections, we stick to Assumption 1 in the sequel. $\qquad \square$

Remark 2. In the light of Remark 1, we cannot relax both Slater feasibility of (2) and boundedness of P in order to maintain the result of Corollary 1. For example, consider

$$
z_{MISDP} = \min \quad \left\langle \begin{pmatrix} -1 & -1 \\ -1 & 0 \end{pmatrix}, \begin{pmatrix} x_1 & x_2 \\ x_2 & x_3 \end{pmatrix} \right\rangle
$$

$$
\text{s.t.} \quad \begin{pmatrix} x_1 & x_2 \\ x_2 & x_3 \end{pmatrix} \succeq 0, \quad x_1 = 0, \quad x_2, x_3 \in \mathbb{Z}.
$$

It is clear that $z_{MISDP} = z_{SDP} = 0$. Now, let the operator $\mathcal{A}_2(\cdot)$ capture the constraint $x_1 = 0$. Then the set P is unbounded and $\mathrm{conv}(P)$ does not have a Slater feasible point. The Lagrangian dual function $g(S)$ becomes

$$
g(S) = \min_{x_2, x_3 \in \mathbb{Z}} \left\langle \begin{pmatrix} -1 - s_1 & -1 - s_2 \\ -1 - s_2 & -s_3 \end{pmatrix}, \begin{pmatrix} 0 & x_2 \\ x_2 & x_3 \end{pmatrix} \right\rangle.
$$

Since the system $S \succeq \mathbf{0}$, $s_3 = 0$ and $s_2 = -1$ is infeasible, we have $g(S) = -\infty$ for all $S \succeq \mathbf{0}$. Thus, $z_{LD} = -\infty$. $\qquad \square$

In the sequel, we briefly describe conditions under which equality throughout the sequence $z_{SDP} \leq z_{LD} \leq z_{MISDP}$ holds. Let $\mathcal{F}_{MISDP}, \mathcal{F}_{SDP}$ and $\mathcal{F}_{LD}$ denote the feasible sets of $(MISDP)$, (SDP) and (2), respectively. Moreover, for any convex set $K \subseteq \mathcal{S}^n$, its normal cone at $X \in K$ is defined as

$$
\mathcal{N}_K(X) := \left\{ Z \in \mathcal{S}^n \ : \ \langle Z, X \rangle \geq \langle Z, Y \rangle \text{ for all } Y \in K \right\}.
$$

Theorem 2. *Let χ_{MISDP}^* and χ_{LD}^* denote the set of optimizers to $(MISDP)$ and (2), respectively. Then,*

(i) *$z_{LD} = z_{MISDP}$ if and only if $-C \in \mathcal{N}_{\mathcal{F}_{LD}}(X^*)$ for all $X^* \in \chi_{MISDP}^*$;*

(ii) *$z_{SDP} = z_{LD}$ if and only if $-C \in \mathcal{N}_{\mathcal{F}_{SDP}}(X^*)$ for all $X^* \in \chi_{LD}^*$.*

Proof. (i) Let $X^* \in \chi_{MISDP}^*$. Since $X^* \in \mathcal{F}_{LD}$, $z_{MISDP} = z_{LD}$ if and only if X^* is also an optimizer to (2). The latter holds if and only if X^* is such that $\langle C, X^* \rangle \leq \langle C, Y \rangle$ for all $Y \in \mathcal{F}_{LD}$, which is equivalent to $-C \in \mathcal{N}_{\mathcal{F}_{LD}}(X^*)$.

(ii) The proof of the second statement is very similar, replacing χ_{MISDP}^* by χ_{LD}^* and $\mathcal{F}_{LD}$ by $\mathcal{F}_{SDP}$. $\qquad \square$

It follows from Theorem 2 that a sufficient condition for $z_{LD} = z_{MISDP}$ is that

$$
\mathrm{conv}\left(P \cap \{ X \in \mathcal{S}^n \ : \ \mathcal{A}_1(X) = b_1, \ X \succeq \mathbf{0} \}\right) = \mathrm{conv}(P) \cap \{ X \in \mathcal{S}^n \ : \ \mathcal{A}_1(X) = b_1, \ X \succeq \mathbf{0} \},
$$

whereas a sufficient condition for $z_{SDP} = z_{LD}$ is

$$
\mathrm{conv}(P) = \{ X \in \mathcal{S}^n \ : \ \mathcal{A}_2(X) = b_2, \ l_{ij} \leq X_{ij} \leq u_{ij} \text{ for all } (i,j) \in \mathcal{J} \}.
$$

3 Hierarchy of Lagrangian dual bounds for ISDP

In this section, we extend on the general theory of Section 2 and introduce a hierarchy of Lagrangian dual bounds for purely integer SDPs, i.e., problems of the form $(MISDP)$ with $\mathcal{J} = [n] \times [n]$.

Without loss of generality, we assume that B_{ij} is of the same form $B \subseteq \mathbb{Z}$ for all $(i, j) \in [n] \times [n]$. This assumption is non-restrictive, as tighter bounds for certain entries in X can be enforced by the linear constraints. Let l and u denote the smallest and lowest value of B, respectively. Hence, we consider problems of the following form:

$$
\begin{aligned}
z_{ISDP} := \quad &\min \quad \langle C, X \rangle \\
&\text{s.t.} \quad \mathcal{A}(X) = b, \ X \succeq \mathbf{0} \qquad\qquad (ISDP)\\
&\qquad X \in B^{n \times n}.
\end{aligned}
$$

A main challenge in the construction of the Lagrangian dual problem is to decide on the splitting between $\mathcal{A}_1(X) = b_1$ and $\mathcal{A}_2(x) = b_2$. One choice, which we consider here, is to dualize all of them, i.e., $\mathcal{A}_1(X) = b_1$ captures all linear constraints. To prevent the resulting bound to be equivalent to the continuous SDP relaxation, there is another property of the optimization problem that we can exploit. In many problems that we consider, i.e., the ones resulting from discrete optimization problems, the feasible matrix variables do not only have integer entries, they also have an upper bound on their rank. This upper bound results from the combinatorial nature of the problem setting. To exploit it, we define for a given integer $1 \leq r \leq n$ the set

$$
\mathcal{S}_+^n(B, r) := \mathcal{S}_+^n \cap B^{n \times n} \cap \{X \ : \ \text{rank}(X) \leq r\}
$$

as the set of PSD matrices having entries in B with an upper bound on its rank. We also define $\mathcal{S}_+^n(B) := \mathcal{S}_+^n(B, n)$ as the set of integer PSD matrices without the rank-assumption. Since B is bounded, the set $\mathcal{S}_+^n(B, r)$ is finite for all $r \leq n$.

For several well-known integer sets B, simple descriptions of $\mathcal{S}_+^n(B, r)$ or $\mathcal{S}_+^n(B)$ have been derived, e.g., in the form of a finite generating set of integer rank-1 PSD matrices. Table 1 lists several of such sets, including an explicit description of its elements and cardinality, see [5, 33, 40]. More results on the structural properties of PSD matrices having integer entries are derived by De Loera et al. [34].

Set $\mathcal{D}^n$	Description	$|\mathcal{D}^n|$
$\mathcal{S}_+^n(\{0,1\})$	$X = \sum_i x_i x_i^\top$, where $x_i \in \{0,1\}^n$ and each x_i and x_j with $i \neq j$ have non-overlapping support	B_{n+1}
$\mathcal{S}_+^n(\{0,1\}, 1)$	$X = x x^\top$, where $x \in \{0,1\}^n$	2^n
$\mathcal{S}_+^n(\{0,1\}, r)$	$X = \sum_{i=1}^n x_i x_i^\top$, where $x_i \in \{0,1\}^n$ and each x_i and x_j with $i \neq j$ have non-overlapping support	$\sum_{k=0}^{r+1} \left\{ {n+1 \atop k} \right\}$
$\mathcal{S}_+^n(\{0,\pm1\})$	$X = \sum_i x_i x_i^\top$, where $x_i \in \{0,\pm1\}^n$ and each x_i and x_j with $i \neq j$ have non-overlapping support	D_n
$\mathcal{S}_+^n(\{0,\pm1\}, 1)$	$X = x x^\top$, where $x \in \{0,\pm1\}^n$	$1 + \sum_{k=1}^n \binom{n}{k} 2^{k-1}$
$\mathcal{S}_+^n(\{0,\pm1\}, r)$	$X = \sum_{i=1}^r x_i x_i^\top$, where $x_i \in \{0,\pm1\}^n$ and each x_i and x_j with $i \neq j$ have non-overlapping support	$\sum_{k=0}^n S_B(n, k)$
$\mathcal{S}_+^n(\{\pm1\})$	$X = x x^\top$, where $x \in \{\pm1\}^n$	2^{n-1}

Table 1: Descriptions and cardinalities of the sets $\mathcal{S}_+^n(B, r)$ and $\mathcal{S}_+^n(B)$ for different B and r. In the third column, B_n denotes the nth Bell number, $\left\{ {n \atop k} \right\}$ denotes the Stirling number of the second kind, D_n denotes the nth Dowling number and $S_B(n, k)$ denotes the B-type Stirling number of the second kind. See Appendix A for details on these numbers and their relation to integer PSD matrices.

In the sequel, we let $\mathcal{D}^n = \mathcal{S}_+^n(B, r)$, where the rank constraint is implied by the constraints $\mathcal{A}(X) = b$ in combination with $X \in B^{n \times n}$ and $X \succeq \mathbf{0}$. In other words, $\mathcal{D}^n$ contains the feasible set of $(ISDP)$. In

this more specified setting, the problem $(ISDP)$ reduces to the following problem:

$$
\begin{aligned}
\min \quad & \langle C, X \rangle \\
\text{s.t.} \quad & \mathcal{A}(X) = b, \ X \succeq \mathbf{0} \\
& X \in \mathcal{D}^n.
\end{aligned}
\tag{4}
$$

Although the exact structure of $\mathcal{D}^n$ depends on the problem, it can for example be one of the sets described in Table 1. For instance, if $B = \{0,1\}$ and there is no restriction on a rank, we have $\mathcal{D}^n = \mathcal{S}_+^n(\{0,1\})$. If we now dualize $\mathcal{A}(X) = b$ and $X \succeq \mathbf{0}$, what remains is $P = \{0,1\}^{n \times n}$. Since $\text{conv}(P) = [0,1]^{n \times n}$, Theorem 1 implies that z_{LD} will be equal to z_{SDP}. In what follows, we discuss a general approach that can be applied to strengthen the resulting Lagrangian dual bound.

The set P can be further tightened by keeping (a part of) the constraint $X \succeq \mathbf{0}$ in the set P. If we require that all elements in P must be PSD, the set P becomes $\mathcal{D}^n$, which is practically infeasible to optimize over. However, for relatively small values of m, it is possible to optimize over $\mathcal{D}^m$, e.g., by a complete enumeration. Therefore, instead of letting P be the entire set $\mathcal{D}^n$, we can impose a condition on certain submatrices to be $m \times m$ matrices of the form $\mathcal{D}^m$. To that end, assume that $\mathcal{K}^p = \{K_1, \ldots, K_N\}$ denotes a packing on the set $[n]$ where $m_\ell := |K_\ell|$ for $K_\ell \in \mathcal{K}^p$ is such that $m_\ell \le p$ for all $K_\ell \in \mathcal{K}^p$, for some given positive integer p. Now, we consider the following ISDP that is equivalent to (4):

$$
\begin{aligned}
\min \quad & \langle C, X \rangle \\
\text{s.t.} \quad & \mathcal{A}(X) = b, \ X \succeq \mathbf{0} \\
& X[K_\ell] \in \mathcal{D}^{m_\ell} \qquad \forall K_\ell \in \mathcal{K}^p \\
& X \in B^{n \times n}.
\end{aligned}
\tag{5}
$$

Indeed, the equivalence between (4) and (5) follows from the fact that the rank-constraint in the definition of $\mathcal{D}^n$ (if any), is implied by $\mathcal{A}(X) = b$, $X \in B^{n \times n}$ and $X \succeq \mathbf{0}$. After dualizing $X \succeq \mathbf{0}$ and $\mathcal{A}(X) = b$, we obtain the following feasible set of remaining constraints, which we denote by $P(\mathcal{K}^p)$:

$$
P(\mathcal{K}^p) := \left\{ X \in \mathcal{S}^n \cap B^{n \times n} \ : \ X[K_\ell] \in \mathcal{D}^{m_\ell} \text{ for all } K_\ell \in \mathcal{K}^p \right\}.
\tag{6}
$$

Let us now check whether this set $P(\mathcal{K}^p)$ can be effectively applied as the feasible set of the Lagrangian dual function $g(S, \lambda)$. Let $X^* = \arg\min\{\mathcal{L}(X, \hat{S}, \hat{\lambda}) : X \in P(\mathcal{K}^p)\}$ for some $(\hat{S}, \hat{\lambda})$. First, since the index sets in $\mathcal{K}^p$ are mutually disjoint, the submatrices $\{X[K_\ell] : K_\ell \in \mathcal{K}^p\}$ have a pairwise disjoint support in X. Hence, the optimal submatrices $X^*[K_\ell]$ can be obtained via a complete enumeration over $\mathcal{D}^{m_\ell}$ for all $K_\ell \in \mathcal{K}^p$ independently. Moreover, any element X_{ij}^* where i and j do not belong to the same set in $\mathcal{K}^p$ is set to u if $(C - \hat{S} + \mathcal{A}_1^*(\hat{\lambda}))_{ij} < 0$ and to l otherwise. We conclude that the optimization over $P(\mathcal{K}^p)$ is tractable for small p.

It follows from Theorem 1 that the optimal Lagrangian dual value of (5) with dualized constraints $\mathcal{A}(X) = b$ and $X \succeq \mathbf{0}$ equals

$$
\begin{aligned}
z_{LD}^p := \min \quad & \langle C, X \rangle \\
\text{s.t.} \quad & \mathcal{A}(X) = b, \ X \succeq \mathbf{0} \\
& X[K_\ell] \in \mathcal{P}^{m_\ell} \qquad \forall K_\ell \in \mathcal{K}^p \\
& l \le X_{ij} \le u \text{ for all } (i,j) \in [n] \times [n],
\end{aligned}
\tag{7}
$$

where $\mathcal{P}^{m_\ell} := \text{conv}(\mathcal{D}^{m_\ell})$ denotes the convex hull of $\mathcal{D}^{m_\ell}$.

The value of p, i.e., the maximum size of the sets in $\mathcal{K}^p$, has an impact on the quality of (7). If p increases, the proportion of entries that is required to be both integer and PSD in the subproblems becomes larger, leading to improved bounds as provided in the theorem below. However, this comes at the cost of computation time, as the size of $\mathcal{D}^{m_\ell}$ often grows exponentially, see Table 1. Formally, the hierarchy of the Lagrangian dual approach with respect to the value of p can be captured in the following theorem.

Theorem 3. *Let $\mathcal{K}^{p_1}$ and $\mathcal{K}^{p_2}$ with $p_1 \le p_2$ denote two packings on the set $[n]$, where for each two distinct $i, j \in [n]$ we have: if i and j are in the same subset in $\mathcal{K}^{p_1}$, then i and j are also in the same subset in $\mathcal{K}^{p_2}$. Then,*

$$
z_{SDP} \le z_{LD}^{p_1} \le z_{LD}^{p_2} \le z_{ISDP},
$$

where $z_{LD}^{p_i}$ denotes the Lagrangian dual bound obtained by optimizing over $P(\mathcal{K}^{p_i})$ in each subproblem.

Proof. For both packings $\mathcal{K}^{p_i}$, the inequalities $z_{SDP} \leq z_{LD}^{p_i} \leq z_{ISDP}$ follow immediately from Corollary 1. Now, let X be feasible for (7) with $\mathcal{K}^{p_2}$. Now, for each $K_\ell \in \mathcal{K}^{p_1}$, we know that $K_\ell \subseteq K_{\ell^*}$ for some $K_{\ell^*} \in \mathcal{K}^{p_2}$. Hence, $X[K_\ell]$ is a submatrix of $X[K_{\ell^*}]$, which is an element of $\mathcal{P}^{m_{\ell^*}}$. Since the $m_\ell \times m_\ell$ submatrices of all elements in $\mathcal{P}^{m_{\ell^*}}$ are in $\mathcal{P}^{m_\ell}$, it follows that $X[K_\ell] \in \mathcal{P}^{m_\ell}$ for all $K_\ell \in \mathcal{K}^{p_1}$. So, X is feasible for (7) with $\mathcal{K}^{p_1}$, resulting in $z_{LD}^{p_1} \leq z_{LD}^{p_2}$. $\qquad\square$

Semidefinite relaxations including constraints of the form $X[K_\ell] \in \text{conv}\,(\mathcal{D}^{m_\ell})$ with $\mathcal{D}^{m_\ell} = \mathcal{S}_+^{m_\ell}(\{0,1\})$ or $\mathcal{D}^{m_\ell} = \mathcal{S}_+^{m_\ell}(\{\pm 1\})$ have been studied in [1, 16], in which these are referred to as exact subgraph constraints (ESCs). Those papers also consider overlapping submatrices, i.e., $\mathcal{K}^p$ is not required to be a packing, and include the underlying graph structure of the considered problems. Clearly, we can also introduce both ingredients. While overlapping constraints would result in a more difficult set $P(\mathcal{K}^p)$ in $(LR(S, \lambda))$ than the one we propose here, additional graph structure can make the Lagrangian dual problem easier to solve.

4   Approaches for solving the Lagrangian dual problem

In this section we present several approaches to compute Lagrangian dual bounds for problems of the form $(MISDP)$. Although the optimum to the Lagrangian dual (LD) is theoretically the same as the solution to (2), an explicit description of $\text{conv}(P)$ is in most cases unavailable. Therefore, the here presented algorithms are designed to solve the optimization problem (LD). Since the Lagrangian dual bounds of the hierarchy presented in Section 3 also fit in the framework of (LD), the presented algorithms can also be used to compute problems of the form (5).

The Lagrangian dual function $g(S, \lambda)$, see $(LR(S, \lambda))$, is a piecewise linear concave function. Due to its non-differentiability, its optimization relies on the use of subgradients. A subgradient of $g(\cdot, \cdot)$ at a point $(S^*, \lambda^*) \in \mathcal{S}_+^n \times \mathbb{R}^{m_1}$ is a pair $(\Gamma, \gamma) \in \mathcal{S}^n \times \mathbb{R}^{m_1}$ such that

$$g(S, \lambda) \leq g(S^*, \lambda^*) + \langle \Gamma, S - S^* \rangle + \gamma^\top (\lambda - \lambda^*) \qquad \text{for all } (S, \lambda) \in \mathcal{S}_+^n \times \mathbb{R}^{m_1}.$$

Now, let $X^* := \arg\min\{\mathcal{L}(X, S^*, \lambda^*) \,:\, X \in P\}$. Then,

$$g(S, \lambda) \leq g(S^*, \lambda^*) + \langle -X^*, S - S^* \rangle + (\mathcal{A}_1(X^*) - b_1)^\top (\lambda - \lambda^*). \tag{8}$$

This shows that $(\Gamma, \gamma) = (-X^*, \mathcal{A}_1(X^*) - b_1)$ is a subgradient of $g(\cdot, \cdot)$ at (S^*, λ^*). All here presented algorithms exploit subgradients to compute the Lagrangian dual bound (LD).

4.1   A projected-deflected subgradient algorithm

We present a projected-deflected subgradient algorithm for computing Lagrangian dual bounds. The subgradient method was developed by Shor in the 1960s as an extension of the gradient method to non-differentiable functions, see e.g., [53]. The subgradient method is a first-order method in which the slow rate of convergence is compensated by the low complexity of each iteration. The results in [14] show that if extensive tuning is performed, the subgradient method can be competitive with more sophisticated approaches, when the required tolerance is not too high. To the best of our knowledge, we are the first to consider a subgradient algorithm for MISDP.

The projected-deflected subgradient algorithm starts with an initial set of dual multipliers (S^0, λ^0) $\in \mathcal{S}_+^n \times \mathbb{R}^{m_1}$. Then, it iteratively obtains $g(S^\ell, \lambda^\ell)$ by minimizing $\mathcal{L}(X, S^\ell, \lambda^\ell)$ over $X \in P$, see (1), yielding an optimal solution X^ℓ. As indicated before, we assume this optimization to be tractable, e.g., for the problem (5) we exploit the structure of the sets provided in Section 3. The algorithm now computes the subgradient $(\Gamma^\ell, \gamma^\ell) := (-X^\ell, \mathcal{A}_1(X^\ell) - b_1)$. The dual multipliers are updated by a step update in the direction $(D^\ell, d^\ell) \in \mathcal{S}^n \times \mathbb{R}^{m_1}$, where this direction is based on the subgradient $(\Gamma^\ell, \gamma^\ell)$. This can be done in several ways, as we discuss below. Then, we set

$$S^{\ell+1} \leftarrow \mathcal{P}_{\mathcal{S}_+^n}(S^\ell + \alpha^\ell D^\ell) \quad \text{and} \quad \lambda^{\ell+1} \leftarrow \lambda^\ell + \beta^\ell d^\ell, \tag{9}$$

where α^ℓ and β^ℓ are appropriate stepsize parameters. The resulting dual matrix $S^\ell + \alpha^\ell D^\ell$ is projected onto the PSD cone in order to stay dual feasible.

Different strategies for the choice of the dual updates (D^ℓ, d^ℓ) and the stepsize parameters α^ℓ and β^ℓ are proposed in the literature for the case of integer linear programming. For an overview of such strategies, see e.g., [14] and references therein. In our preliminary study [9], we considered several of

those strategies and altered them to the case of MISDP. In particular, the standard subgradient update with Polyak's stepsize [47], the deflected subgradient algorithm using the update scheme from Camerini et al. [6], the deflected subgradient algorithm using the update scheme from Sherali and Ulular [52], and the conditional subgradient algorithm with the Polyak update [47] were implemented. The preliminary results showed that the deflected subgradient algorithm with the update scheme from [52] is the most efficient among the tested algorithms. Therefore, we further consider only that variant of the subgradient algorithm.

The deflected subgradient algorithm [52] constructs the dual update vector as a linear combination between the subgradient vector and the previous dual update vector, i.e.,

$$D^\ell = \Gamma^\ell + \varphi_1^\ell D^{\ell-1} \quad \text{and} \quad d^\ell = \gamma^\ell + \varphi_2^\ell d^{\ell-1}, \tag{10}$$

where $\varphi_1^\ell, \varphi_2^\ell \geq 0$ are deflection parameters. This choice of the dual update prevents the dual multipliers from a zigzagging pattern. Sherali and Ulular [52] propose to take

$$\varphi_1^\ell = \frac{||\Gamma^\ell||_F}{||D^\ell||_F} \quad \text{and} \quad \varphi_2^\ell = \frac{||\gamma^\ell||}{||d^\ell||}, \tag{11}$$

in order to let the new dual update vector bisect the angle between the current subgradient Γ^ℓ (resp. γ^ℓ) and the previous dual update $D^{\ell-1}$ (resp. $d^{\ell-1}$).

The stepsize parameters α^ℓ and β^ℓ can be chosen according to the Polyak update [47], that is

$$\alpha^\ell := \frac{\mu_1^\ell (U^* - g(S^\ell, \lambda^\ell))}{||\Gamma^\ell||_F^2} \quad \text{and} \quad \beta^\ell := \frac{\mu_2^\ell (U^* - g(S^\ell, \lambda^\ell))}{||\gamma^\ell||^2}. \tag{12}$$

Here U^* is an upper bound on the optimal Lagrangian dual value z_{LD}, which can be obtained by a heuristic on the primal problem, and $0 < \mu_1^\ell, \mu_2^\ell \leq 2$ are stepsize parameters. The justification for this step length is its theoretical convergence when U^* is set equal to z_{LD}, see [47].

A pseudo-code of the projected-deflected subgradient algorithm is given in Algorithm 1.

Algorithm 1 Projected-deflected subgradient algorithm for solving (LD)

Input: $C, \mathcal{A}_1, \mathcal{A}_2, b_1, b_2,$ and B_{ij} for all $(i,j) \in \mathcal{J}$.
1: Initialize dual pair $(S^0, \lambda^0) \in \mathcal{S}_+^n \times \mathbb{R}^{m_1}$ and set $\ell = 0$, $\hat{z} = -\infty$.
2: **while** stopping criteria are not met **do**
3:     Compute $g(S^\ell, \lambda^\ell)$ and obtain $X^\ell = \arg\min\{\mathcal{L}(X, S^\ell, \lambda^\ell) : X \in P\}$.
4:     **if** $\hat{z} < g(S^\ell, \lambda^\ell)$ **then**
5:         $\hat{z} \leftarrow g(S^\ell, \lambda^\ell)$
6:     **end if**
7:     Compute subgradients $\Gamma^\ell = -X^\ell$, $\gamma^\ell = \mathcal{A}_1(X^\ell) - b_1$.
8:     Update stepsize parameters α^ℓ, β^ℓ using (12), and dual updates D^ℓ, d^ℓ using (10)–(11).
9:     Update $S^{\ell+1} \leftarrow \mathcal{P}_{\mathcal{S}_+^n}(S^\ell + \alpha^\ell D^\ell)$, $\lambda^{\ell+1} \leftarrow \lambda^\ell + \beta^\ell d^\ell$.
10:     $\ell \leftarrow \ell + 1$
11: **end while**
Output: $\hat{z}$

The subgradient algorithm stops when the difference between consecutive dual multipliers is small, i.e., $||S^\ell - S^{\ell-1}||_F < \varepsilon_1$ and $||\lambda^\ell - \lambda^{\ell-1}|| < \varepsilon_2$ for some predefined parameters $\varepsilon_1, \varepsilon_2 > 0$. Moreover, we implement a stagnation criterion: the algorithm stops if there has been no improvement in the value of the objective function for the last N_{stag} iterations.

4.2 A projected-accelerated subgradient algorithm

We consider here a version of the accelerated gradient method introduced by Nesterov [45]. Namely, we adopt the approach from [45] and adjust it for the case of MISDP. The resulting projected-accelerated subgradient algorithm computes the Lagrangian dual bound (LD).

The projected-accelerated subgradient algorithm starts with an initial set of dual multipliers $(S^0, \lambda^0) \in \mathcal{S}_+^n \times \mathbb{R}^{m_1}$ and introduces auxiliary variables $Y^0 := S^0$ and $\delta^0 := \lambda^0$. Then, the algorithm computes $g(Y^0, \delta^0) = \min_{X \in P} \mathcal{L}(X, Y^0, \delta^0)$, yielding an optimal solution X^0 and corresponding subgradient (Γ^0, γ^0). Next, the algorithm iteratively performs updates of dual multipliers by employing the following two-step iteration

$$S^{\ell+1} \leftarrow \mathcal{P}_{\mathcal{S}_+^n}(Y^\ell + \alpha^\ell \Gamma^\ell), \quad Y^{\ell+1} = S^{\ell+1} + \frac{\eta_\ell - 1}{\eta_{\ell+1}}(S^{\ell+1} - S^\ell)$$

$$\lambda^{\ell+1} \leftarrow \delta^\ell + \beta^\ell \gamma^\ell, \quad \delta^{\ell+1} = \lambda^{\ell+1} + \frac{\eta_\ell - 1}{\eta_{\ell+1}}(\lambda^{\ell+1} - \lambda^\ell), \quad \ell = 0, 1, \ldots,$$

where α^ℓ, β^ℓ are stepsize parameters chosen with respect to the Polyak update [47], see also (12), and $\{\eta_\ell\}_{\ell=0}^\infty$ the sequence generated as follows

$$\eta_0 = 0 \quad \text{and} \quad \eta_{\ell+1} = \frac{1 + \sqrt{1 + 4\eta_\ell^2}}{2}, \quad \ell = 0, 1, \ldots.$$

Iteratively, the algorithm performs a step of gradient ascent to go from Y^ℓ (resp. δ^ℓ) to $S^{\ell+1}$ (resp. $\lambda^{\ell+1}$), followed by a correction in the auxiliary variables, i.e., the so-called look-ahead step. Nesterov's two-point step iteration approach from [45] results in an accelerated gradient algorithm that achieves an improved convergence rate with respect to a simple gradient algorithm. Matrix $(Y^\ell + \alpha_\ell \Gamma^\ell)$ is projected onto the PSD cone so that the dual variable stays dual feasible. A pseudo-code of projected-accelerated subgradient algorithm is presented in Algorithm 2.

We use the same stopping criteria for this algorithm as for Algorithm 1, where we monitor a stagnation criterion with respect to Y^ℓ instead of S^ℓ.

Algorithm 2 Projected-accelerated subgradient algorithm for solving (LD)

Input: $C, \mathcal{A}_1, \mathcal{A}_2, b_1, b_2,$ and B_{ij} for all $(i,j) \in \mathcal{J}$
1: Initialize dual pair $(S^0, \lambda^0) \in \mathcal{S}_+^n \times \mathbb{R}^{m_1}$.
2: Set $Y^0 = S^0$, $\delta^0 = \lambda^0$, $\eta_0 = 0$ and $\ell = 0$.
3: **while** stopping criteria are not met **do**
4:    Obtain $X^\ell = \arg\min\{\mathcal{L}(X, Y^\ell, \delta^\ell) : X \in P\}$.
5:    Compute subgradients $\Gamma^\ell = -X^\ell$, $\gamma^\ell = \mathcal{A}_1(X^\ell) - b_1$.
6:    Update stepsize parameters α^ℓ, β^ℓ using (12).
7:    Update $S^{\ell+1} \leftarrow \mathcal{P}_{\mathcal{S}_+^n}(Y^\ell + \alpha^\ell \Gamma^\ell)$, $\quad \lambda^{\ell+1} \leftarrow \delta^\ell + \beta^\ell \gamma^\ell$.
8:    Compute $\eta_{\ell+1} = \frac{1}{2}\left(1 + \sqrt{1 + 4\eta_\ell^2}\right)$.
9:    Update $Y^{\ell+1} \leftarrow S^{\ell+1} + \frac{\eta_\ell - 1}{\eta_{\ell+1}}(S^{\ell+1} - S^\ell)$, $\quad \delta^{\ell+1} \leftarrow \lambda^{\ell+1} + \frac{\eta_\ell - 1}{\eta_{\ell+1}}(\lambda^{\ell+1} - \lambda^\ell)$.
10:    $\ell \leftarrow \ell + 1$
11: **end while**
12: Compute $\hat{z} = g(S^\ell, \lambda^\ell)$.
Output: $\hat{z}$

4.3 A projected bundle algorithm

We present a projected bundle algorithm for solving the Lagrangian dual problem (LD). The bundle method was introduced in the 1970's [29, 32] as a method for optimizing non-smooth functions. The bundle method has already been exploited for computing SDP bounds, however, our algorithm differs from SDP-based bundle algorithms from the literature. Namely, function evaluations in our algorithm are over a polyhedral set, while in other bundle algorithms function evaluations are over a spectrahedron. We consider here a version of the proximal bundle algorithm [28].

The bundle algorithm starts with an initial set of dual multipliers $(S^0, \lambda^0) \in \mathcal{S}_+^n \times \mathbb{R}^{m_1}$ and obtains X^0 from $g(S^0, \lambda^0) = \min\{\mathcal{L}(X, S^0, \lambda^0) : X \in P\}$, see $(LR(S, \lambda))$. Then, the algorithm computes a subgradient of $g(\cdot, \cdot)$ at (S^0, λ^0), that is, $(\Gamma^0, \gamma^0) = (-X^0, \mathcal{A}_1(X^0) - b_1)$. The bundle algorithm is an iterative algorithm that maintains in each iteration the best current approximation $(\widehat{S}, \widehat{\lambda})$ to the maximizer of $g(\cdot, \cdot)$, (some of) previously computed points $X^i \in P$ and the corresponding subgradients $(\Gamma^i, \gamma^i) = (-X^i, \mathcal{A}_1(X^i) - b_1)$. The set that contains feasible points X^i and corresponding subgradients is called *the bundle* and denoted by $\mathcal{B}$. We denote by $J_\mathcal{B}$ its index set.

One iteration of the algorithm is as follows. We assume to have the bundle $\mathcal{B} = \{(X^i, \Gamma^i, \gamma^i) : i \in J_\mathcal{B}\}$. To compute the next trial point $(S_{\text{trial}}, \lambda_{\text{trial}})$, the bundle method combines the following two concepts. The first concept is to approximate the function $g(S, \lambda)$ in the neighborhood of the previous iterates. That is accomplished by the following subgradient model:

$$g_{\text{appr}}(S, \lambda) := \min\left\{\mathcal{L}(X, S, \lambda) : X \in \text{conv}\left(\{X^i : i \in J_\mathcal{B}\}\right)\right\},$$

which can be rewritten as follows

$$g_{\text{appr}}(S, \lambda) = \min_{\sigma \in \Delta} \sum_{i \in J_\mathcal{B}} \sigma_i \left(\langle C, X^i \rangle + \langle \Gamma^i, S \rangle + \lambda^\top \gamma^i\right), \tag{13}$$

where $\Delta := \{\sigma \in \mathbb{R}^{|J_\mathcal{B}|} : \mathbf{1}^\top \sigma = 1, \ \sigma \geq 0\}$. The second concept is to stay in the vicinity of the center point $(\widehat{S}, \widehat{\lambda})$. This can be achieved by using the proximal point approach that exploits the Moreau–Yosida regularization of $g_{\mathrm{appr}}(\cdot, \cdot)$. The combination of the two ideas results in the following concave problem whose solution gives a new trial point:

$$
\begin{aligned}
\max \quad & g_{\mathrm{appr}}(S, \lambda) - \frac{t}{2}(||S - \widehat{S}||_F^2 + ||\lambda - \widehat{\lambda}||^2) \\
\text{s.t.} \quad & S \succeq \mathbf{0}, \ \lambda \in \mathbb{R}^{m_1}.
\end{aligned}
\tag{14}
$$

Here $t > 0$ is known as the proximity parameter. Since the above problem is a convex quadratic optimization problem, the complexity of each bundle iteration is more costly than a subgradient iteration. Moreover, the above problem is a quadratic semidefinite programming problem. Nevertheless, we proceed in the line of the bundle literature that commonly proposes to solve the dual problem of (14). That is, we aim to solve the following optimization problem:

$$
\begin{aligned}
\min_{Q \succeq 0} \ \max_{\substack{S \in \mathcal{S}^n \\ \lambda \in \mathbb{R}^{m_1}}} \ \min_{\sigma \in \Delta} \ & \sum_{i \in J_\mathcal{B}} \sigma_i \left(\langle C, X^i \rangle + \langle \Gamma^i, S \rangle + \lambda^\top \gamma^i \right) - \frac{t}{2}(||S - \widehat{S}||_F^2 + ||\lambda - \widehat{\lambda}||^2) + \langle S, Q \rangle \\
= \min_{\substack{\sigma \in \Delta \\ Q \succeq 0}} \ \max_{\substack{S \in \mathcal{S}^n \\ \lambda \in \mathbb{R}^{m_1}}} \ & \sum_{i \in J_\mathcal{B}} \sigma_i \left(\langle C, X^i \rangle + \langle \Gamma^i, S \rangle + \lambda^\top \gamma^i \right) - \frac{t}{2}(||S - \widehat{S}||_F^2 + ||\lambda - \widehat{\lambda}||^2) + \langle S, Q \rangle.
\end{aligned}
\tag{15}
$$

Indeed, we may swap around the minimization w.r.t. σ and the maximization w.r.t. (S, λ), due to the set Δ being compact and the objective function being concave in (S, λ) and linear in σ. From the first-order optimality conditions for the inner maximization problem we have

$$
S = \widehat{S} + \frac{1}{t}\left(Q + \sum_{i \in J_\mathcal{B}} \sigma_i \Gamma^i \right), \quad \lambda = \widehat{\lambda} + \frac{1}{t} \sum_{i \in J_\mathcal{B}} \sigma_i \gamma^i.
$$

Now, we incorporate this into (15) and obtain the following minimization problem:

$$
\begin{aligned}
\min_{\sigma \in \Delta, Q \succeq 0} \ & \frac{1}{2t} \langle Q, Q \rangle + \frac{1}{2t} \left\| \sum_{i=1}^\ell \sigma_i \gamma^i \right\|^2 + \frac{1}{2t} \left\langle \sum_{i=1}^\ell \sigma_i X^i, \sum_{i=1}^\ell \sigma_i X^i \right\rangle \\
& + \left\langle \sum_{i=1}^\ell \sigma_i X^i, C - \widehat{S} - \frac{1}{t} Q \right\rangle + \left(\sum_{i=1}^\ell \sigma_i \gamma^i \right)^\top \widehat{\lambda} + \langle \widehat{S}, Q \rangle.
\end{aligned}
\tag{16}
$$

This problem is a convex quadratic optimization problem, and we solve it approximately by keeping one set of the variables constant. In particular, keeping σ fixed we solve for Q, i.e.,

$$
Q = \mathcal{P}_{\mathcal{S}_+^n}\left(\sum_{i \in J_\mathcal{B}} \sigma_i X^i - tS \right),
$$

and keeping Q constant results in a convex quadratic problem over the set Δ that can be efficiently solved. Thus, we start with $Q = \mathbf{0}$ and solve for σ, which we thereafter keep constant to solve for Q, and iterate this process several times to get (approximate) solutions to (16). In practice, one has to make a few iterations to obtain a good approximate solution. However, the most expensive computation is projection onto the PSD cone. Using the final estimates of Q and σ, say $\widetilde{Q}$ and $\widetilde{\sigma}$, we compute the new trial point as follows

$$
S_{\mathrm{trial}} = \mathcal{P}_{\mathcal{S}_+^n}\left(\widehat{S} + \frac{1}{t}\left(\widetilde{Q} + \sum_{i \in J_\mathcal{B}} \widetilde{\sigma}_i \Gamma^i \right) \right) \quad \text{and} \quad \lambda_{\mathrm{trial}} = \widehat{\lambda} + \frac{1}{t} \sum_{i \in J_\mathcal{B}} \widetilde{\sigma}_i \gamma^i.
\tag{17}
$$

To finish one iteration of the bundle algorithm, we still need to evaluate $(LR(S, \lambda))$ at $(S_{\mathrm{trial}}, \lambda_{\mathrm{trial}})$, which results in the matrix $X_{\mathrm{trial}} \in P$. Note that if $\widetilde{\sigma}_i = 0$ for some i, then the corresponding subgradients have no influence in the optimization and therefore we remove them from the bundle $\mathcal{B}$ and thus, from further computations.

The bundle algorithm now needs to decide whether or not $(S_{\mathrm{trial}}, \lambda_{\mathrm{trial}})$ becomes the next center point $(\widehat{S}, \widehat{\lambda})$. The algorithm first checks whether the improvement of the objective function in the trial

point is at least $\rho \in (0, 1)$ fraction of the improvement that the subgradient model provides. If it does, $(S_{\text{trial}}, \lambda_{\text{trial}})$ becomes the next center point, and the bundle $\mathcal{B}$ is updated with X_{trial} and the corresponding subgradient. In the literature this is called a serious step. Otherwise, the algorithm keeps the last center point, but nevertheless updates $\mathcal{B}$ with X_{trial} and the corresponding subgradient, resulting into a so-called null step.

For updating the proximity parameter t, we use an update similar to the one proposed by Kiwiel [28]. The projected bundle algorithm stops if the difference of the function value at the current center point and the function value of the subgradient model at the trial point is smaller than some prescribed tolerance. The algorithm also stops if it reaches the maximum number of iterations. A pseudo-code of the projected bundle algorithm is given in Algorithm 3.

Algorithm 3 Projected bundle algorithm for solving (LD)

Input: $C, \mathcal{A}_1, \mathcal{A}_2, b_1, b_2,$ and B_{ij} for all $(i,j) \in \mathcal{J}$
1: Initialize dual pair $(S^0, \lambda^0) \in \mathcal{S}^n_+ \times \mathbb{R}^{m_1}$, t and ρ, and set $\ell = 0$.
2: Compute $X^0 = \arg\min\{\mathcal{L}(X, S^0, \lambda^0) : X \in P\}$ and the subgradient (Γ^0, γ^0).
3: Set the bundle to $\mathcal{B} = \{(X^0, \Gamma^0, \gamma^0)\}$.
4: **while** stopping criteria are not met **do**
5:     Solve (15) to obtain $\widetilde{Q}$ and $\widetilde{\sigma}$.
6:     Determine next trial point $(S_{\text{trial}}, \lambda_{\text{trial}})$ using (17).
7:     Compute $g(S_{\text{trial}}, \lambda_{\text{trial}})$ and obtain $X_{\text{trial}} = \arg\min\{\mathcal{L}(X, S_{\text{trial}}, \lambda_{\text{trial}}) : X \in P\}$.
8:     Obtain subgradients $\Gamma_{\text{trial}} = -X_{\text{trial}}, \gamma_{\text{trial}} = \mathcal{A}_1(X_{\text{trial}}) - b_1$.
9:     Decide whether $(S_{\text{trial}}, \lambda_{\text{trial}})$ becomes $(\widehat{S}, \widehat{\lambda})$ (serious step) or not (null step).
10:    Update the bundle $\mathcal{B}$ and the parameter t.
11:    $\ell \leftarrow \ell + 1$
12: **end while**
Output: $g(\widehat{S}, \widehat{\lambda})$

---

5 The max-k-cut problem

The max-k-cut problem is the problem of partitioning the vertex set of a graph into k subsets such that the total weight of edges joining different sets is maximized. For the case that $k = 2$, the max-k-cut problem is known as the max-cut problem. The max-k-cut problem is $\mathcal{NP}$-hard [3] and has many applications such as VLSI design, frequency planning, statistical physics, sports team scheduling, see e.g. [4, 44].

Let $G = (V, E)$ be an undirected graph with n vertices and $W = (w_{ij})$ a weight matrix such that $w_{ij} = 0$ if $\{i, j\} \notin E$. Let $Z \in \{0, 1\}^{n \times k}$ denote the characteristic matrix of a partitioning of V, where $Z_{ij} = 1$ if vertex i is in subset j and $Z_{ij} = 0$ otherwise. Then, $Y := ZZ^\top$ is a binary $n \times n$ PSD matrix of rank at most k that satisfies $\text{diag}(Y) = \mathbf{1}_n$. Following [40, Corollary 2], the rank constraint can be enforced by a linear matrix inequality, leading to the following ISDP formulation for the max-k-cut problem:

$$
\begin{aligned}
\max \quad & \frac{1}{2}\langle L, Y \rangle \\
\text{s.t.} \quad & \text{diag}(Y) = \mathbf{1}_n, \quad \begin{pmatrix} k & \mathbf{1}_n^\top \\ \mathbf{1}_n & Y \end{pmatrix} \succeq \mathbf{0} \\
& Y \in \mathcal{S}^n, \ Y \in \{0, 1\}^{n \times n},
\end{aligned}
\tag{18}
$$

where $L := \text{Diag}(W\mathbf{1}_n) - W$ denotes the weighted Laplacian matrix of G.

A more compact reformulation of the max-k-cut problem can be obtained by rewriting the linear matrix inequality of (18) to $Y - \frac{1}{k}\mathbf{J} \succeq \mathbf{0}$ using the Schur complement lemma. By defining the new variable $X := \frac{k}{k-1}(Y - \frac{1}{k}\mathbf{J})$, the formulation (18) is equivalent to the following discrete program:

$$
\begin{aligned}
\max \quad & \frac{k-1}{2k}\langle L, X \rangle \\
\text{s.t.} \quad & \text{diag}(X) = \mathbf{1}_n, \quad X \succeq \mathbf{0} \\
& X \in \left\{\frac{-1}{k-1}, 1\right\}^{n \times n}.
\end{aligned}
\tag{19}
$$

Here, $X_{ij} = -1/(k-1)$ if vertices i, j are in different subsets of the partition and $X_{ij} = 1$ otherwise. One can easily show that a matrix X that is feasible for (19) is of rank at most $k - 1$. The formulation (19) is well-known in the literature, see [15], where also a geometric meaning of the program is provided.

The following well-known SDP relaxation for the max-k-cut problem is obtained from (19) by relaxing $X_{ij} \in \left\{ \frac{-1}{k-1}, 1 \right\}$ to $\frac{-1}{k-1} \le X_{ij} \le 1$:

$$
\begin{aligned}
\max \quad & \frac{k-1}{2k} \langle L, X \rangle \\
\text{s.t.} \quad & \operatorname{diag}(X) = \mathbf{1}_n, \quad X \succeq \mathbf{0} \\
& X_{ij} \ge \frac{-1}{k-1}, \quad \forall i, j \in [n].
\end{aligned}
\tag{20}
$$

The constraints $X_{ij} \le 1$ for all i, j are omitted since they are redundant due to $X \succeq \mathbf{0}$ and $\operatorname{diag}(X) = \mathbf{1}_n$. For $k = 2$, also the constraint $X \ge -\mathbf{J}_n$ is redundant, and the resulting problem is equivalent to the basic SDP relaxation for the max-cut problem, see e.g., [21].

The SDP relaxation can be further tightened by adding valid inequalities such as the triangle and the clique inequalities, see e.g., [19]. The triangle inequalities ensure that if any two vertices i and j are in the same partition, and vertices j and h are in the same partition, then also vertices i and h have to be in the same partition. The resulting $3\binom{n}{3}$ triangle inequalities are of the form:

$$
X_{ij} + X_{jh} - X_{ih} \le 1, \quad i, j, h \in V.
\tag{21}
$$

The clique inequalities ensure that for each subset of $k + 1$ vertices, at least two of the vertices belong to the same partition. Those inequalities are of the form:

$$
\sum_{i,j \in S, \ i<j} X_{ij} \ge -\frac{k}{2}, \quad \forall S \subseteq V \quad \text{where} \quad |S| = k + 1.
\tag{22}
$$

There are $\binom{n}{k+1}$ clique inequalities. While separation of triangle inequalities can be done efficiently, the exact separation of clique inequalities is $\mathcal{NP}$-hard in general. There exist other types of inequalities that one may add to the SDP relaxation (20), such as the general clique inequalities, wheel inequalities and bicycle wheel inequalities, see e.g., [10].

In the sequel we follow the idea of Section 3 to derive bounds for the max-k-cut problem. We first consider the following discrete SDP that is equivalent to (19):

$$
\begin{aligned}
\max \quad & \frac{k-1}{2k} \langle L, X \rangle \\
\text{s.t.} \quad & \operatorname{diag}(X) = \mathbf{1}_n, \quad X \succeq \mathbf{0} \\
& X[K_\ell] \in \widehat{\mathcal{D}}^{m_\ell} \quad \forall K_\ell \in \mathcal{K}^p \\
& X \in \left\{ \frac{-1}{k-1}, 1 \right\}^{n \times n},
\end{aligned}
\tag{23}
$$

where

$$
\widehat{\mathcal{D}}^{m_\ell} := \left\{ X \in \mathcal{S}_+^{m_\ell} \cap \left\{ \frac{-1}{k-1}, 1 \right\}^{m_\ell \times m_\ell} : \operatorname{diag}(X) = \mathbf{1}_{m_\ell}, \operatorname{rank}(X) \le k - 1 \right\},
\tag{24}
$$

and p is a given integer such that $3 \le p \le n$. The set $\widehat{\mathcal{D}}^{m_\ell}$ is the image of the elements in $\mathcal{S}_+^{m_\ell}(\{0,1\}, k)$ with diagonal elements equal to one under the mapping $Y \to \frac{k}{k-1}(Y - \frac{1}{k}\mathbf{J})$. Therefore, we can fully enumerate over its elements by exploiting Table 1.

The collection $\mathcal{K}^p$ is a collection of subsets of $[n]$, each of size at most p. Observe that the exact sub-matrix constraint $X[K_\ell]$ for some set K_ℓ only concerns the off-diagonal elements of $X[K_\ell]$. Consequently, the submatrix constraints for two sets K_1 and K_2 that intersect on a single index, i.e., $|K_1 \cap K_2| = 1$, can be evaluated independently. Thus, we no longer require $\mathcal{K}$ to consist of mutually disjoint sets, it is sufficient to require that the elements in $\mathcal{K}$ do not pairwise intersect in more than one index. We call a collection of subsets satisfying this property an *edge-packing* of $[n]$. We construct the Lagrangian dual by dualizing the constraints $X \succeq \mathbf{0}$, which results in the Lagrangian dual bound, z_{LD}^p.

13

6 Computational results for the max-k-cut problem

In this section we provide a computational study on the strength of the Lagrangian dual bounds obtained from the discrete SDP formulation of the max-k-cut problem discussed in Section 5. The design of the computational tests is presented in Section 6.1. Computational results for various k and instance types are presented in Sections 6.2 –6.5.

6.1 Design of the computational experiments

In this section we present the design of our computational experiments. More specifically, we present the considered instances, the various bounding approaches that we take into account and several implementation details of the algorithms for computing Lagrangian dual bounds.

6.1.1 Description of instances

We test our approaches on several instance classes from the literature, mainly inspired by the max-cut problem or partition problems. These instance classes are as follows:

- **Rudy instances**: These instances are randomly generated using the machine-independent graph generator rudy [50]. The instances 'g05_n' are unweighted graphs on n vertices, where each edge is included with probability $1/2$. The instance classes 'pm1d_n' and 'pm1s_n' contain graphs on n vertices with edge densities 0.9 and 0.1, respectively, having edge weights chosen uniformly at random from $\{0, \pm 1\}$. The graphs 'pw$d$_n' are defined on n vertices with edge density $d \in \{0.1, 0.5, 0.9\}$, where the weights are integers from $\{0, \ldots, 10\}$ chosen uniformly at random. Finally, the class 'wd_n' is defined similarly, except for the weights being chosen as integers from $\{-10, \ldots, 10\}$. Each instance class consists of 10 randomly generated instances.

  For $60 \leq n \leq 100$, the resulting test set consists of 130 instances. The authors of [22, 24] consider similar instances for larger graphs with $n = 180$, leading to an additional set of 90 instances.

- **Spinglass instances**: Toroidal 2D-grid graphs arise in physics when computing ground states for Ising spinglasses, see e.g., [54]. Those instances are generated using the rudy graph generator [50]. In particular, spinglass2pm$_{n_t, n_r}$ is a toroidal 2D-grid for a spinglass model with weights $\{-1, 1\}$. The grid has size $n_t \times n_r$. The percentage of edges with negative weights is 50%.

- **Band instances:** These instances are considered in [12, 23]. Given the number of vertices n and the value of k, we let the edge set of a band graph be defined as $E := \{\{i, j\} : j - i \leq k + 1, \ 1 \leq i < j \leq n\}$. In these graphs, 50% of edge weights are set to -1 and the others are set to 1.

6.1.2 Bounding approaches

In Section 5 we have shown that the max-k-cut problem can be formulated as the discrete SDP (23). The Lagrangian dual problem obtained from this formulation can be solved by any of the alternative approaches presented in Section 4. Note that these approaches were designed for minimization problems, so we multiply the objective function of (23) by -1 in order to view it as a minimization problem. In our numerical tests, we consider the following bounding approaches for the max-k-cut problem:

- **SDP-B**: This approach refers to solving the basic SDP relaxation of the max-k-cut problem, i.e., (20), using the alternating direction method of multipliers (ADMM). We follow a similar implementation as described in [41, Section 3.1] for the graph partition problem.

- **LD**: By LD we refer to the Lagrangian dual bound obtained from the discrete SDP (23) after dualizing $X \succeq \mathbf{0}$. Observe that (23) fits in the form (5), where $B = \{\frac{-1}{k-1}, 1\}$ and $\widehat{\mathcal{D}}^{m_\ell}$ as given in (24). To obtain an edge-packing $\mathcal{K}^p$ that is likely to provide strong bounds, we use the heuristic that is described in Section 6.1.3. The resulting Lagrangian dual problem is now solved using the three approaches described in Section 4:

  - **LD-DSG**: This approach refers to solving the Lagrangian dual problem using the projected-deflected subgradient algorithm discussed in Section 4.1. For the deflection parameter φ_1^ℓ, see (11), (observe that φ_2^ℓ does not exist as we linearize only the PSD constraint), we use the update scheme of Sherali and Ulular [52]. Moreover, we use Polyak's stepsize update [47], see (12), where U^* is computed using the heuristic introduced in Section 6.1.4. For the value of

μ_1^ℓ we use the approach of Held and Karp [35], implying that $\mu_1^0 = 1$ and we halve its value each time the obtained bound did not increase for $N_{\text{step}} = 40$ subsequent iterations. As starting point S^0, we use the (approximate) dual solution that we obtain from the implementation of the ADMM mentioned above. For the stopping criteria, we use $N_{\text{stag}} = 100$ and $\varepsilon_1 = 0.01$ (observe that ε_2 does not exist). These values are based on preliminary experiments.

- **LD-ASG**: This approach refers to solving the Lagrangian dual problem using the projected-accelerated subgradient algorithm discussed in Section 4.2. The parameters in this approach are α^ℓ, η_ℓ and μ_1^ℓ (again, β^ℓ does not exist, since we do only dualize the PSD constraint). The sequence $\{\eta_\ell\}_\ell^\infty$ is fully determined after taking $\eta_0 = 0$. The parameter α^ℓ is computed via (12), where U^* is chosen as described in Section 6.1.4. With respect to the parameter μ_1^ℓ, we discriminate between two cases: We take $\mu_1^\ell = 0.025$ for $n \leq 100$, whereas $\mu_1^\ell = 0.012$ if $n > 100$. The reason to lower the value of μ_1^ℓ for larger instances is to reduce the oscillating behavior at the start of the algorithm that is heavier for larger instances. We use the same starting point S^0 as has been used for LD-DSG. For the stopping criteria, we take $\varepsilon_1 = 1e-3$ and $N_{\text{stag}} = 15$, which are based on preliminary tests.

- **LD-Bundle**: This approach refers to solving the Lagrangian dual problem using the projected bundle algorithm described in Section 4.3. As starting point for the proximity parameter t, we use $t = \kappa |z_{SDP} - U^*| / \|X^*\|_F$, where $\kappa > 0$ is a constant, X^* is an approximate solution to the SDP relaxation (20), z_{SDP} its corresponding objective value and U^* is the bound as computed by the heuristic of Section 6.1.4. We take κ equal to 0.1 if $n \leq 100$ and $\kappa = 0.2$ otherwise. The parameter ρ for deciding whether we take a serious or a null step is set to $\frac{1}{2}$. The problem (15) is solved iteratively, that is, we iteratively solve with respect to either σ or Q, while keeping the other variable fixed. We continue doing this until the norm of the difference between consecutive values of σ is less than $1e-4$, or after the maximum number of 20 iterations are reached. The entire bundle algorithm terminates whenever the function value of the current center point and the function value of the subgradient model at the trial point is smaller than $1e-3$ or if the maximum number of 600 iterations is reached.

6.1.3 Edge-packing heuristic

Recall that the Lagrangian dual bound depends on the edge-packing $\mathcal{K}^p$. To determine an edge-packing that is likely to provide a strong Lagrangian dual bound, we apply the following heuristic. Let X^* denote an (approximate) solution to the SDP relaxation (20). For each triple of vertices $i, j, h \in V$, we define

$$v_{ijh}^\Delta := (X_{ij}^* + X_{jh}^* - X_{ih}^* - 1)^+ + (X_{ij}^* - X_{jh}^* + X_{ih}^* - 1)^+ + (-X_{ij}^* + X_{jh}^* + X_{ih}^* - 1)^+$$

and for each $S \subseteq V$ with $|S| = k + 1$, we define

$$v_S^C := \left(-\frac{k}{2} - \sum_{i,j \in S, i < j} X_{ij}^* \right)^+,$$

where $(\cdot)^+ = \max(0, \cdot)$. Indeed, v_{ijh}^Δ and v_S^C provide measures of the violation of the triangle- and clique-inequalities (21) and (22), respectively. When a triple (i, j, h) or a subset S violates these inequalities by a large amount, we have an incentive to include this triple/set within the edge-packing $\mathcal{K}^p$, as we expect its inclusion to have a positive effect on the resulting bound. Now, we proceed as follows. We let $\overline{E}$ denote the set of edges that are already included in the packing, where we initially set $\overline{E} = \emptyset$. Now, we start with a set K containing the three vertices i, j, h that have the largest positive value v_{ijh}^Δ. Now, we greedily add to K the vertex t that maximizes the total violation

$$\sum_{i,j \in K, i < j} v_{tij}^\Delta + \sum_{\substack{S \subseteq [n]:|S| = k+1, \\ t \in S, S \setminus t \subseteq K}} v_S^C$$

associated with vertex t and the vertices already in K. We restrict ourselves to the vertices t such that $\{(t, k) : k \in K\}$ does not intersect with $\overline{E}$, preventing an edge from being packed twice. We continue adding vertices to K until $|K| = p$ or no more vertex with positive violation can be found. Then, we add K to the edge-packing $\mathcal{K}^p$ and add to $\overline{E}$ all edges in the subgraph induced by K. Now, we repeat this process starting from three vertices whose induced subgraph is not yet covered by $\overline{E}$, until $\mathcal{K}^p$ contains a maximum number of subsets. This maximum is in our computations set equal to $5n$.

6.1.4 Computation of near-optimal k-partitions

To apply Polyak's stepsize update [47], see (12), one needs a lower bound U^* on the optimal value of the max-k-cut problem. Hence, we aim to find near-optimal feasible solutions to the max-k-cut problem. We proceed as follows. Let X^* denote an (approximate) solution to the SDP relaxation (20) of the max-k-cut problem. The approximation algorithm introduced by Frieze and Jerrum [15], in the sequel denoted by the FJ algorithm, applies a rounding strategy on X^* to obtain a feasible max-k-cut. As the algorithm is probabilistic, we perform 500 independent trials of the FJ algorithm and take its best solution.

Next, we try to improve this feasible solution using a variable depth search [27]. Let $V_0 = V$ and compute for each vertex $v \in V_0$ the solution we obtain from moving v to its best alternative cluster (i.e., the one that leads to the largest cut value). Let v_1 denote the vertex which movement leads to the best improvement (even if the cost improvement is negative) and delete v_1 from V_0. Now, we repeat the procedure until V_0 is empty, which leads to an ordered list of vertices $\{v_1, \ldots, v_n\}$. Finally, we compute the value of $T \in [n]$ such that the simultaneous movement of the vertices $\{v_1, \ldots, v_T\}$ to their best alternative cluster results in the best possible cut. The resulting feasible k-partition can be decoded by a matrix $Y \in \{\frac{-1}{k-1}, 1\}^{n \times n}$ such that

$$Y_{ij} = \begin{cases} \frac{-1}{k-1} & \text{if } i \text{ and } j \text{ are in different subsets of the partition,} \\ 1 & \text{otherwise.} \end{cases}$$

Inspired by the heuristic of Rendl et al. [48], we now replace X^* by $\omega X^* + (1 - \omega)Y$, where $\omega \in (0, 1)$, and repeat the entire process again using this new solution X^*. We continue doing so until no better solution can be found by either the FJ algorithm or the variable-depth search. Based on preliminary tests, we take $\omega = 0.8$.

6.2 Computational results for the max-cut problem

In this section we present computational results for the max-cut problem.

Table 2 shows average bound values for the max-cut problem on several types of Rudy instances. The column 'n' denotes the number of vertices, whereas the columns 'SDP-B' and 'OPT' denote the SDP bound (20) and the known optima (as presented in [24, 57]), respectively, averaged over all 10 instances of that class. For each instance, we performed the approaches LD-DSG, LD-ASG and LD-Bundle to compute the Lagrangian dual bound. Although the resulting bounds obtained by these approaches slightly differ, the difference is always within 1% (relative to the minimum of the three). The columns 'LD' provide this minimal value. The values in the columns 'Rel. gap closed' are computed by the formula $100 \cdot (\text{SDP} - \text{LD})/(\text{SDP} - \text{OPT})$ and denote the percentage of the gap between the SDP bound and the optimum that is closed by the Lagrangian dual bound. In Table 2 we only present results for $p = 3$ and $p = 7$, and results for p ranging between $p = 3$ and $p = 23$ are given in Figure 1. The corresponding computation times (in wall clock time) are presented in Table 3. Here we do make a distinction between the approaches LD-ASG, LD-ASG and LD-Bundle.

Observe that Table 2 and 3 only include a subset of the instance types of the Rudy instances. The reason for this is that for the large instances with $n = 180$, optima are known for only six out of nine instance types [24]. To keep the presentation balanced, we only present results for six instance types of the small instances as well.

It can be observed from Table 2 that the Lagrangian dual bound provides a significant improvement over the bound (20). The relative gap closed by the Lagrangian dual for the instances with $n \leq 100$ is in the range 41%-57%. For the larger instances with $n = 180$, this gap is in the range 28%-31%. This drop falls within our expectations, since if n increases, the number of subsets of $[n]$ with size at most p also increases. Hence, the relative number of subsets $K_\ell \subseteq [n]$ for which the constraint $X[K_\ell] \in \widehat{\mathcal{D}}^{m_\ell}$ is included in the discrete SDP becomes lower as n increases, leading to a weaker bound. When p increases, we observe that the Lagrangian dual bound becomes stronger, which is indeed in line with Theorem 3. This comes at the cost of computation time, since the evaluation of the dual function $g(S, \lambda)$ becomes more costly. Indeed, this evaluation involves a full enumeration over the elements in $\widehat{\mathcal{D}}^{m_\ell}$, whose cardinality increases (drastically) with p. This effect is clearly visible in Table 3.

When comparing the various approaches for computing the Lagrangian dual bound, we observe from Table 3 that LD-ASG is the fastest approach on almost all instance types, although LD-DSG follows on a short distance. Apparently, the accelerated two-step approach by Nesterov [45] pays off, which is reflected by a lower number of iterations needed to converge. The projected-bundle method LD-Bundle yields the

Instance Class	n	SDP-B	OPT	LD bound for p = 7		LD bound for p = 17	
				LD	Rel. gap closed (%)	LD	Rel. gap closed (%)
g05_80	80	950.4	929.1	940.2	47.7	939.5	51.4
pm1d_80	80	297.3	254.5	276.9	47.6	274.9	52.3
pm1s_100	100	139.7	122.6	131.7	46.5	131.1	50.1
pw05_100	100	8364.7	8147.1	8273.8	41.8	8268.1	44.4
w01_100	100	790.7	699.9	742.3	53.3	739.1	56.9
w09_100	100	2544.8	2176.9	2381.8	44.3	2372.3	46.9
g05_180	180	4582.6	4494.6	4557.7	28.2	4557.8	28.2
pm1d_180	180	1042.3	870.5	990.2	30.3	989.0	31.0
pw05_180	180	25 952.2	25 347.3	25 771.8	29.8	25 765.7	30.8
pw09_180	180	43 515.2	42 964.7	43 355.1	29.1	43 350.5	29.9
w05_180	180	4671.1	3913.1	4447.6	29.5	4441.5	30.3
w09_180	180	6173.5	5167.8	5888.6	28.3	5879.4	29.3

Table 2: Average bound values (SDP-B, LD and optimum) for the max-cut problem on Rudy instances. Each reported value is the average over 10 randomly generated instances of that type.

| Instance Class	n	SDP-B	LD bound for p = 7			LD bound for p = 17		
			LD-DSG	LD-ASG	LD-Bundle	LD-DSG	LD-ASG	LD-Bundle
g05_80	80	0.19	1.05	0.99	1.56	12.07	11.42	5.87
pm1d_80	80	0.10	1.10	0.87	1.49	13.33	10.81	5.98
pm1s_100	100	0.92	2.52	2.64	3.21	17.68	18.45	8.57
pw05_100	100	0.28	2.82	1.61	5.96	26.89	15.27	16.15
w01_100	100	0.36	3.21	2.55	5.61	26.81	22.42	14.28
w09_100	100	0.46	3.23	1.94	6.31	30.46	16.13	16.68
g05_180	180	2.74	10.65	10.16	14.16	68.83	55.45	28.51
pm1d_180	180	1.54	11.10	9.42	12.88	84.46	63.38	30.54
pw05_180	180	0.96	12.62	9.03	27.29	101.44	58.83	50.01
pw09_180	180	0.96	11.65	8.09	27.95	96.89	55.70	50.84
w05_180	180	0.92	12.50	9.16	23.90	99.04	59.71	49.86
w09_180	180	0.91	12.76	8.96	25.63	105.31	54.41	50.78

Table 3: Average computation times (wall clock time in seconds) of Lagrangian dual approaches for the max-cut problem on Rudy instances. Each reported value is the average over 10 randomly generated instances of that type.

largest computation times for $p = 7$. We observe in our experiments that the number of iterations needed by the bundle method is much lower than for the subgradient methods, but one iteration is in general more costly, leading to a larger total computation time. At the same time, it seems that the bundle method is more robust against (small) changes in the parameters, whereas such changes cause a large effect on the performance of the subgradient methods. This is a well-known short-coming of first-order methods, see e.g., [14]. This robustness is also reflected by our experiments for larger values of p. As explained before, p has a negative effect on the computation time of the Lagrangian dual bound. This effect is mainly due to an increase in the computation time per iteration due to a full enumeration. Since the bundle method needs the least number of iterations, the approach LD-Bundle becomes the favored method when p becomes large.

To further study the effect of p on the Lagrangian dual bound, we present in Figure 1 the (average) relative gap closed and the (average) computation time per instance type for different values of p. For $n \leq 100$, we test for $3 \leq p \leq 23$. For p larger than 23, we observe that we can no longer perform the computations within one hour. For the large Rudy instances with $n = 180$, we observe that this frontier is reached earlier, hence we present results for $3 \leq p \leq 19$. Different from Table 2 and 3, we now take into account all instance types for which optimal values are known. The Lagrangian dual bounds are computed via the approach LD-DSG, but results obtained by the other two algorithms are similar.

Figure 1 confesses that the value of p has a positive impact on the quality of the bound, resulting in a significant reduction in the gap between the obtained bounds and the optimal solution. We observe that the marginal improvement over p is, however, diminishing. That is, an improvement in p from 3 to 5 has a large effect on the quality of the Lagrangian dual bounds, whereas this positive effect tails off when p

becomes larger. At the same time, the computation times remain relatively small for values of p up to 20 (resp. 15) for the small (resp. large) instance. For larger values of p, we observe that the computation times rapidly grow due to the extremely large cardinality of the involved sets $\widehat{D}^{m_\ell}$, see Appendix A.
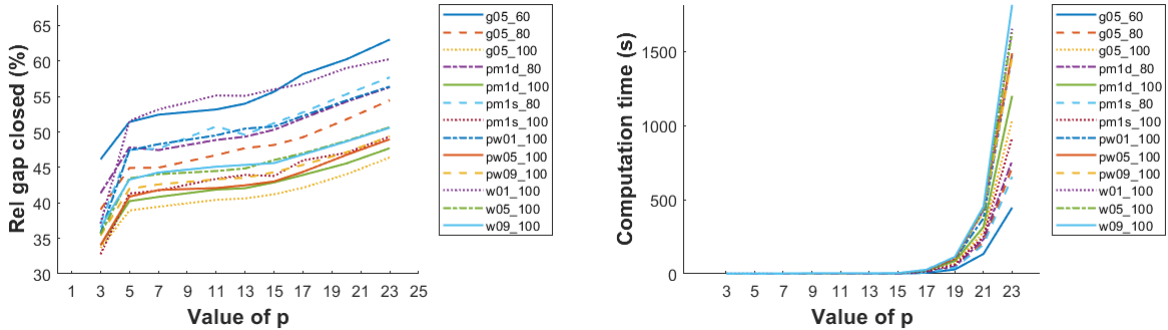
(a) Rel. gap closed for the max-cut problem on small Rudy instances

(b) Computation times for the max-cut problem on small Rudy instances
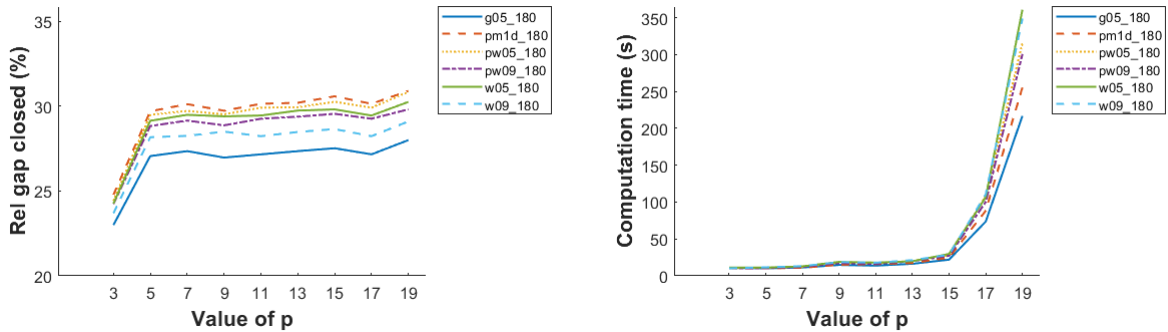
(c) Rel. gap closed for the max-cut problem on large Rudy instances

(d) Computation times for the max-cut problem on large Rudy instances

Figure 1: Average relative gap (between SDP (20) and optimum) closed and average computation times of the approach LD-DSG for different values of p. Results are presented for the max-cut problem on small and large Rudy instances.

6.3 Computational results for the max-3-cut problem

We now discuss our findings for the max-3-cut problem. For this problem we perform experiments on the same instances as for the max-cut problem, leading to Table 4 and 5 and Figure 2. Although the set-up of the experiments and the explanation of the columns in the tables is the same, there are two differences. First, due to the larger computation times of the Lagrangian dual bounds for the max-3-cut problem, we test up to $p \leq 15$ (resp. $p \leq 13$) for the smaller (resp. larger) instances. For that reason, we report results for $p = 13$ instead of $p = 17$ in Table 4 and 5. These larger computation times can be explained from the fact that the sets $\widehat{\mathcal{D}}^{m_\ell}$ have a (significant) larger cardinality when k increases, see Appendix A. Moreover, since optimal solutions for the max-3-cut problem are not available in the literature, we instead report the lower bounds as obtained by the procedure explained in Section 6.1.4. Consequently, the relative gap closed is computed by the formula $100 \cdot (\text{SDP} - \text{LD})/(\text{SDP} - \text{LB})$.

When considering Table 4 and 5, we draw very similar conclusions as for the max-cut problem. Again, the Lagrangian dual bounds provide a substantial improvement compared to the basic SDP bound (20), although the relative gaps closed are smaller than for the max-cut problem. This can be explained by the fact that we compare to a lower bound instead of the optimum.

The projected-accelerated version of the subgradient algorithm seems to perform best for small values of p, whereas the projected-bundle method becomes favorable for large values of p. The conclusions drawn from Figure 2 are similar to those drawn from Figure 1.

Instance Class	n	SDP-B	LB	LD bound for p = 7		LD bound for p = 13	
				LD	Rel. gap closed (%)	LD	Rel. gap closed (%)
g05_80	80	1251.5	1215.0	1246.4	13.8	1245.8	15.4
pm1d_80	80	368.8	296.8	356.5	17.0	355.7	18.2
pm1s_100	100	172.6	142.2	167.5	16.6	167.0	18.2
pw05_100	100	11 007.3	10 620.2	10 962.5	11.6	10 957.9	12.8
w01_100	100	954.8	807.9	922.2	22.2	918.9	24.4
w09_100	100	3171.1	2568.4	3076.4	15.7	3068.0	17.1
g05_180	180	6074.1	5906.5	6061.2	7.7	6060.6	8.1
pm1d_180	180	1316.2	997.6	1287.5	9.0	1285.9	9.5
pw05_180	180	34 320.1	33 191.3	34 230.3	8.0	34 223.9	8.5
pw09_180	180	57 789.1	56 744.6	57 699.9	8.5	57 694.8	9.0
w05_180	180	5923.4	4617.3	5789.4	10.3	5783.3	10.7
w09_180	180	7838.3	5988.2	7671.3	9.0	7660.4	9.6

Table 4: Average bound values (SDP-B, LD and LB) for the max-3-cut problem on Rudy instances. Each reported value is the average over 10 randomly generated instances of that type.

Instance Class	n	SDP-B	LD bound for p = 7			LD bound for p = 13		
			LD-DSG	LD-ASG	LD-Bundle	LD-DSG	LD-ASG	LD-Bundle
g05_80	80	0.07	1.50	0.85	1.44	41.73	22.46	17.26
pm1d_80	80	0.05	1.81	0.80	1.81	48.40	22.85	17.55
pm1s_100	100	0.34	2.74	2.26	3.17	55.62	43.95	27.23
pw05_100	100	0.23	4.10	1.56	5.09	96.14	33.50	38.28
w01_100	100	0.20	4.17	2.36	5.01	85.00	51.36	32.39
w09_100	100	0.41	4.86	1.95	5.19	110.85	34.80	39.82
g05_180	180	0.81	11.51	7.35	13.66	352.37	209.35	147.64
pm1d_180	180	0.57	13.82	7.79	14.63	414.38	210.28	159.96
pw05_180	180	0.94	14.58	7.78	17.21	500.36	204.30	162.29
pw09_180	180	0.84	15.75	7.36	17.58	526.15	197.02	168.00
w05_180	180	1.18	16.64	8.44	17.76	504.47	220.80	160.83
w09_180	180	1.48	17.98	8.77	19.77	567.58	215.09	180.73

Table 5: Average computation times of Lagrangian dual approaches for the max-3-cut problem on Rudy instances. Each reported value is the average over 10 randomly generated instances of that type.

6.4 Computational results for the max-4-cut problem

For the max-4-cut problem, we test our approaches on the Band instances [12, 23]. Table 6 shows the bound values, the optima as reported in [12] and the relative gap closed by the Lagrangian dual bound for various values of p. The column names are similar as in previous sections. Table 7 and Figure 3 show the corresponding computation times of the Lagrangian dual approach, where we restrict ourselves to the approach LD-ASG, as this turned out to be the most efficient procedure for small values of p, see Sections 6.2 and 6.3.

We observe that the Lagrangian dual bounds are significantly improving over the basic SDP bound SDP-B, with a relative gap closed by the Lagrangian dual bound ranging between 19–38% for small values of p ($p = 3, 5$) to 65–74% for $p = 13$. Interestingly, these gaps seem not to depend much on the value of n and are high also for larger instances. Also, the diminishing marginal effect over p that we observed for the rudy instances, seems not to be present for these type of instances. Indeed, Figure 3 suggests that the relative gap closed is improving at a fairly constant rate over p.

## 6.5 Computational results for the Spinglass instances

The final class of instances for which we present results are the Spinglass instances [54]. In Table 8 we present bounds, optima and relative gap closed for $p = 3$ and $p = 9$ on a set of the Spinglass instances. Also, computation times for the approach LD-ASG are shown. We refer to the previous sections for an explanation of the columns.
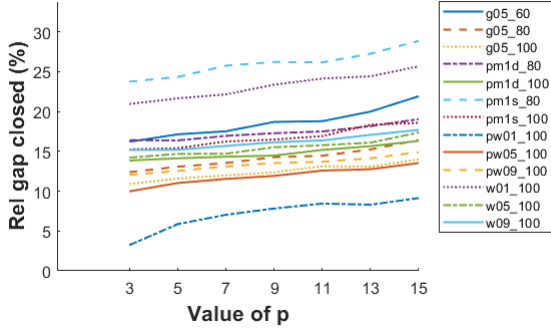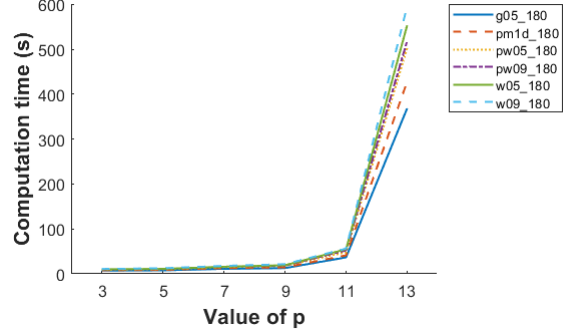
(a) Rel. gap closed for the max-3-cut problem on small Rudy instances
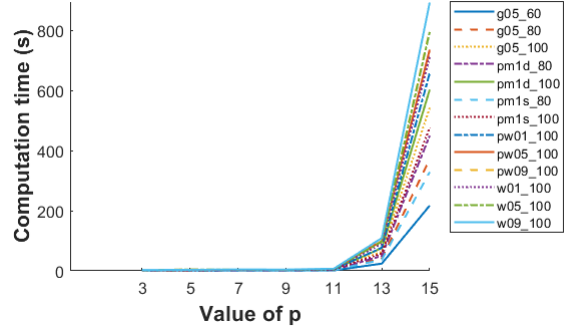
(b) Computation times for the max-3-cut problem on small Rudy instances
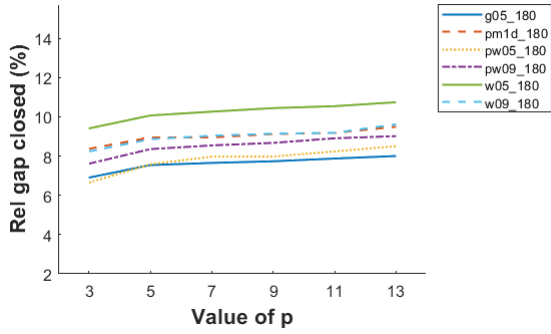
(c) Rel. gap closed for the max-3-cut problem on large Rudy instances

(d) Computation times for the max-3-cut problem on large Rudy instances

Figure 2: Relative gap closed and computation times of the approach LD-DSG for different values of p. Results are presented for the max-3-cut problem on small and large Rudy instances.

We observe that our Lagrangian dual bounds perform very well on these special structured instances. Already for $p = 3$, we observe that the relative gap closed is at least about 40%, and in some cases as large as 58%. For $p = 9$, we observe a significant improvement, leading to a relative gap closed within the range 64-80%. These bounds can be computed relatively fast, with computation times below a minute for almost all considered instances.

Instance Class	n	SDP-B	OPT	Relative gap closed by LD bound (%)					
				p = 3	p = 5	p = 7	p = 9	p = 11	p = 13
band50_4	50	68.9	59	37.5	34.2	41.0	46.9	62.1	73.6
band100_4	100	138.7	117	34.2	28.8	35.8	49.9	51.5	65.0
band150_4	150	208.6	175	20.2	29.0	35.9	48.7	53.9	64.4
band200_4	200	278.4	234	20.3	28.4	33.9	48.9	53.0	67.1
band250_4	250	348.3	292	19.9	30.3	35.3	46.5	54.6	66.4

Table 6: Bound values (SDP-B and optimum) and relative gap closed by the Lagrangian dual bound for various values of p for the max-4-cut problem on Band instances.

Instance Class	n	Computation times (s) for LD-ASG					
		p = 3	p = 5	p = 7	p = 9	p = 11	p = 13
band50_4	50	0.27	0.21	0.28	0.43	8.52	172.10
band100_4	100	1.41	1.28	1.56	1.86	22.80	377.78
band150_4	150	2.61	5.33	5.81	6.95	63.29	994.69
band200_4	200	4.63	10.77	10.55	12.87	102.51	1764.86
band250_4	250	7.33	16.63	18.10	20.68	153.30	2616.91

Table 7: Computation times in seconds of Lagrangian dual approach LD-ASG for different values of p for the max-4-cut problem on Band instances.

7 Conclusion

In this work we consider the Lagrangian duality theory for MISDPs and show its potential to provide strong bounds for problems that can be modeled as a MISDP.

Starting from a MISDP in standard form, we introduce the Lagrangian relaxation and the associated Lagrangian dual problem, see (LD). Since this problem exploits both positive semidefiniteness and integrality, we show that the resulting Lagrangian dual bound is always at least as strong as the continuous SDP relaxation, see Theorem 1 and Corollary 1. Also, we consider conditions under which the Lagrangian dual bound attains its maximum and minimum value, see Theorem 2. For the case of pure integer SDPs, we extend on the aforementioned theory, and introduce a hierarchy of Lagrangian dual bounds by exploiting the matrix structure of integer PSD matrices, see Table 1. The level-p bound is obtained as the Lagrangian dual of an ISDP, see (5), having additional constraints on a subset of its p-by-p submatrices.

In Section 4 we present three algorithms for solving the Lagrangian dual problem, each of them finding its roots in non-smooth optimization. To the best of our knowledge, the projected-deflected and the projected-accelerated subgradient algorithms, see Algorithms 1 and 2, are the first implementations of subgradient methods to obtain SDP bounds. We also present a projected bundle algorithm in Algorithm 3, which is more efficient than bundle methods in the literature. Namely, instead of solving an SDP in each iteration, our algorithm enumerates over a discrete set, which can be done efficiently for moderate sizes of p.

The three proposed algorithms are exploited to obtain Lagrangian dual bounds for the max-k-cut problem in Section 6. The Lagrangian dual bounds are stronger than their continuous SDP counterparts on all considered instances and often by a significant amount. The bounds become stronger for larger levels in the hierarchy, although the marginal improvement diminishes over p. The relative gap closed by the Lagrangian dual bound (with $p = 13$) compared to the continuous SDP bound can be as large as 70% for certain graph types. When comparing the three algorithms, we observe that the projected-accelerated subgradient algorithm is favored for small values of p, whereas the projected bundle algorithm becomes superior when p becomes large. This computational advantage of the bundle algorithm upon the subgradient algorithms is evident from the number of iterations required to achieve the given accuracy. Since the computation times of the three proposed algorithms is moderate, this opens perspectives for exploiting Lagrangian dual bounds in a branching framework.

References

[1] E. Adams, M.F. Anjos, F. Rendl, and A. Wiegele. A hierarchy of subgraph projection-based semidefinite relaxations for some NP-hard graph optimization problems. *INFOR: Inf. Syst. Oper. Res.*,
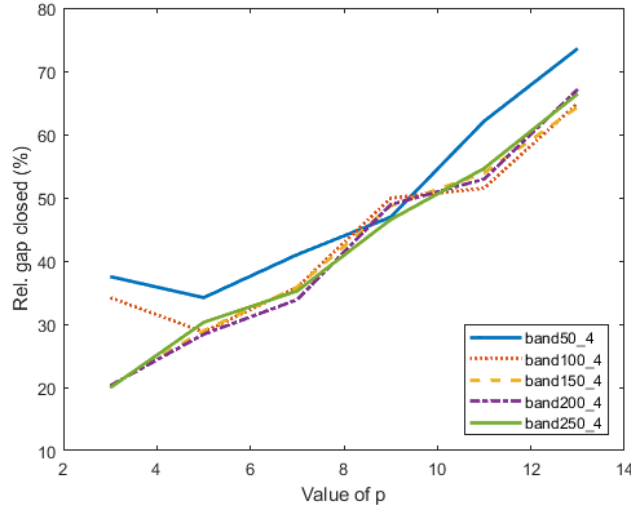
Figure 3: Relative gap closed (in %) by Lagrangian dual approach LD-ASG for the max-4-cut problem on Band instances for different values of p.

Instance Class ($n_t \times n_r$)	n	m	k	SDP-B	OPT	LD bound for p = 3			LD bound for p = 9		
						LD	Rel. gap closed (%)	Comp. time (s)	LD	Rel. gap closed (%)	Comp. time (s)
12×12	144	288	2	113.6	104	108.0	57.9	13.51	106.5	73.7	15.16
12×13	169	338	2	128.1	114	121.1	49.4	22.92	118.2	70.4	23.31
14×14	196	392	2	150.0	132	141.4	47.8	35.34	137.2	71.3	39.05
15×15	225	450	2	167.1	146	156.5	50.0	38.72	152.0	71.5	40.78
16×16	256	521	2	197.2	178	187.1	52.7	55.11	183.3	72.6	63.02
12×12	144	288	3	129.3	120	124.8	48.5	18.06	121.9	79.2	22.86
12×13	169	338	3	150.5	138	145.1	43.7	24.15	142.4	64.7	27.64
14×14	196	392	3	175.6	161	169.2	44.0	29.07	164.7	74.8	28.66
15×15	225	450	3	197.8	179	189.3	45.3	36.70	184.4	71.3	41.14
16×16	256	521	3	229.4	211	221.1	45.1	47.83	216.5	70.2	74.92
12×12	144	288	4	130.2	120	126.2	39.5	10.84	123.0	70.6	14.92
12×13	169	338	4	152.1	139	146.8	39.9	10.37	143.6	64.6	15.82
14×14	196	392	4	177.4	162	171.3	39.5	16.17	166.6	70.0	25.75
15×15	225	450	4	199.8	180	191.7	40.9	22.24	185.1	74.3	35.53
16×16	256	512	4	231.8	213	224.2	40.3	39.01	219.2	67.2	41.54

Table 8: Bounds (SDP-B, LD and optimum), relative gap closed by the Lagrangian dual bound and computation times for LD-ASG for the max-4-cut problem on the Spinglass instances.

53(1):40–48, 2015.

[2] V.E. Adler. Set partitions and integrable hierarchies. *Theor. Math. Phys*, 187(3):842–870, 2016.

[3] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and hardness of approximation problems. *Proceedings., 33rd Annual Symposium on Foundations of Computer Science, IEEE Computer Society Press*, pages 14–23, 1992.

[4] F. Barahona, M. Grötschel, M. Jünger, and G. Reinelt. An application of combinatorial optimization to statistical physics and circuit layout design. *Oper. Res.*, 36(3):493–513, 1988.

[5] A. Berman and C. Xu. {0,1} completely positive matrices. *Linear Algebra Appl.*, 399:35–51, 2005.

[6] P. M. Camerini, L. Fratta, and F. Maffioli. *Nondifferentiable Optimization*, chapter On improving relaxation methods by modified gradient techniques, pages 26–34. Springer Berlin Heidelberg, Berlin, Heidelberg, 1975.

[7] J.H. Conway and R.K. Guy. *Famous Families of Numbers*, pages 91–126. Springer, New York, NY, 1996.

[8] D. Cvetković, M. Čangalović, and V. Kovačević-Vujčić. Semidefinite programming methods for the symmetric traveling salesman problem. In G. Cornujols, R.E. Burkard, and G.J. Woeginger, editors, *Integer programming and Combinatorial Optimization (IPCO 1999)*, volume 1610 of *Lecture Notes in Computer Science*. Springer, Berlin, Heidelberg, 1999.

[9] F. de Meijer. *Integrality and cutting planes in semidefinite programming approaches for combinatorial optimization*. PhD thesis, Tilburg Univeristy, 2023.

[10] V.J. Rodrigues de Sousa, M.F. Anjos, and S. Le Digabel. Computational study of valid inequalities for the maximum k-cut problem. *Ann. Oper. Res.*, 265(1):5–27, 2018.

[11] A. Eisenblätter. The semidefinite relaxation of the k-partition polytope is strong. In William J. Cook and Andreas S. Schulz, editors, *Integer Programming and Combinatorial Optimization*, pages 273–290. Springer Berlin Heidelberg, 2002.

[12] R. Fakhimi, H. Validi, I.V. Hicks, T. Terlaky, and L.F Zuluaga. The max k-cut problem on classical and quantum solvers. https://api.semanticscholar.org/CorpusID:248220524, 2022.

[13] I. Fischer, G. Gruber, F. Rendl, and R. Sotirov. Computational experience with a bundle approach for semidefinite cutting plane relaxations of max-cut and equipartition. *Math. Program.*, 105:451–469, 2006.

[14] A. Frangioni, B. Gendron, and E. Gorgone. On the computational efficiency of subgradient methods: a case study with Lagrangian bounds. *Math. Program. Comput.*, 9:573 – 604, 2017.

[15] A. M. Frieze and M. Jerrum. Improved approximation algorithms for max-k-cut and max bisection. *Algorithmica*, 18:67–81, 1995.

[16] E. Gaar and F. Rendl. A computational study of exact subgraph based SDP bounds for max-cut, stable set and coloring. *Math. Program.*, 183:283–308, 2020.

[17] T. Gally, M.E. Pfetsch, and S. Ulbrich. A framework for solving mixed-integer semidefinite programs. *Optim. Methods Softw.*, 33(3):594–632, 2018.

[18] A. M. Geoffrion. Lagrangean relaxation for integer programming. In *In M.L. Balinski, editor, Approaches to Integer Programming*, pages 82–114. Springer Berlin Heidelberg, Berlin, Heidelberg, 1974.

[19] B. Ghaddar, M. F. Anjos, and F. Liers. A branch-and-cut algorithm based on semidefinite programming for the minimum k-partition problem. *Ann. Oper. Res.*, 188:155–174, 2011.

[20] W. Gil-González, A. Molina-Cabrera, O.D. Montoya, and L.F. Grisales-Noreña. An MI-SDP model for optimal location and sizing of distributed generators in dc grids that guarantees the global optimum. *Appl. Sci.*, 10(21):7681, 2020.

[21] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995.

[22] N. Gusmeroli, T. Hrga, B. Lužar, J. Povh, M. Siebenhofer, and A. Wiegele. BiqBin: A parallel branch-and-bound solver for binary quadratic problems with linear constraints. *ACM Trans. Math. Softw.*, 48(2), July 2022.

[23] H. Hijazi and G. Wang. Exploiting sparsity for the min k-partition problem. *Math. Program. Comput.*, 12(109):109–130, 2020.

[24] T. Hrga and J. Povh. MADAM: a parallel exact solver for max-cut based on semidefinite programming and ADMM. *Comput. Optim. Appl.*, 80:347–375, 2021.

[25] H. Hu and R. Sotirov. On solving the quadratic shortest path problem. *INFORMS J. Comput.*, 32(2):219–233, 2020.

[26] OEIS Foundation Inc. The On-Line Encyclopedia of Integer Sequences. https://oeis.org, 2024.

[27] B. W. Kernighan and S. Lin. An efficient heuristic procedure for partitioning graphs. *The Bell System Technical Journal*, 49(2):291–307, 1970.

[28] K. C. Kiwiel. Proximity control in bundle methods for convex nondifferentiable minimization. *Math. Program.*, 46:105–122, 1990.

[29] K.C. Kiwiel. *Methods of descent for nondiferentiable optimization*. Springer, Berlin, 1985.

[30] K. Kobayashi and Y. Takano. A branch-and-cut algorithm for solving mixed-integer semidefinite optimization problems. *Comput. Optim. Appl.*, 75:493–513, 2020.

[31] B. Kocuk and R.D.A. Morán. On subadditive duality for conic mixed-integer programs. *SIAM Journal on Optimization*, 29(3):2320–2336, 2019.

[32] C. Lemaréchal. Nonsmooth optimization and descent methods. Technical report, International Institute for Applied Systems Analysis, 1978.

[33] A.N. Letchford and M.M. Sørensen. Binary positive semidefinite matrices and associated integer polytopes. *Math. Program. Series A*, 131:253–271, 2012.

[34] J.A. De Loera, B. Marsters, L. Xu, and Shixuan Zhang. Integer points in arbitrary convex cones: The case of the PSD and SOC cones. Preprint, 2024.

[35] Held M. and R.M. Karp. The travelling salesman problem and minimum spanning trees: Part II. *Math. Program.*, 1:6–25, 1971.

[36] F. Matter and M.E. Pfetsch. Presolving for mixed-integer semidefinite optimization. *INFORMS J. Optim.*, 5(2):131–154, 2022.

[37] F. de Meijer, M. Siebenhofer, R. Sotirov, and A. Wiegele. Spanning and splitting: Integer semidefinite programming for the quadratic minimum spanning tree problem. https://arxiv.org/abs/2410.04997, 2024.

[38] F. de Meijer and R. Sotirov. SDP-based bounds for the quadratic cycle cover problem via cutting plane augmented Lagrangian methods and reinforcement learning. *INFORMS J. Comput.*, 33(4):1262–1276, 2021.

[39] F. de Meijer and R. Sotirov. The Chvátal-Gomory procedure for integer SDPs with applications in combinatorial optimization (accepted). *Math. Program.*, 2024.

[40] F. de Meijer and R. Sotirov. On integrality in semidefinite programming for discrete optimization. *SIAM J. Optim.*, 34(1):1071–1096, 2024.

[41] F. de Meijer, R. Sotirov, A. Wiegele, and S. Zhao. Partitioning through projections: Strong SDP bounds for large graph partition problems. *Comp. & Oper. Res.*, 151:106088, 2023.

[42] P. Meurdesoif. Strengthening the Lovász bound for graph coloring. *Math. Program.*, 102:577 – 588, 2005.

[43] R.R. Meyer. On the existence of optimal solutions to integer and mixed integer programming problems. *Math. Program.*, 7:223–235, 1974.

[44] J. E. Mitchell. Realignment in the National Football League: Did they do it right? *Nav. Res. Logist.*, 50(7):683–701, 2003.

[45] Y. Nesterov. A method for solving the convex programming problem with convergence rate $o(1/k^2)$. *Proceedings of the USSR Academy of Sciences*, 269:543–547, 1983.

[46] D. E. Oliveira, H. Wolkowicz, and Y. Xu. ADMM for the SDP relaxation of the QAP. *Math. Program. Comput.*, 10:631–658, 2018.

[47] B.T. Polyak. Minimization of unsmooth functionals. *USSR Comput. Math. Math. Phys.*, 9:14–29, 1969.

[48] F. Rendl, G. Rinaldi, and A. Wiegele. Solving max-cut to optimality by intersecting semidefinite and polyhedral relaxations. *Math. Program.*, 121(2):307, 2010.

[49] F. Rendl and R. Sotirov. Bounds for the quadratic assignment problem using the bundle method. *Math. Program.*, 109:505–524, 2007.

[50] G. Rinaldi. Rudy: a generator for random graphs, 1998.

[51] B.E. Sagan and J.P. Swanson. q-Stirling numbers in type B. *Eur. J. Comb.*, 118:103899, 2024.

[52] H.D. Sherali and O. Ulular. A primal-dual conjugate subgradient algorithm for specially structured linear and convex programming problems. *Appl. Math. Optim.*, 20:193–221, 1989.

[53] N.Z. Shor, Krzysztof C. Kiwiel, and A. Ruszcayǹski. *Minimization methods for non-differentiable functions.* Springer-Verlag, New York, NY, USA, 1985.

[54] C. de Simone, M. Diehl, M. Jünger, P. Mutzel, G. Reinelt, and G. Rinaldi. Exact ground states of Ising spin glasses: New experimental results with a branch-and-cut algorithm. *J. Stat. Phys.*, pages 487–496, 1995.

[55] M. Sion. On general minimax theorems. *Pac. J. Math.*, 8(1):171–176, 1958.

[56] D. Sun, K. C. Toh, Y. Yuan, and X. Y. Zhao. SDPNAL +: A Matlab software for semidefinite programming with bound constraints (version 1.0). *Optim. Methods Softw.*, 35(1):87–115, 2020.

[57] A. Wiegele. Biq Mac Library - A collection of max-cut and quadratic 0-1 programming instances of medium size. Preprint, 2007.

[58] K. Yonekura and Y. Kanno. Global optimization of robust truss topology via mixed integer semidefinite programming. *Optim. Eng.*, 11(3):355–379, 2010.

[59] X. Zheng, H. Chen, Y. Xu, Z. Li, Z. Lin, and Z. Liang. A mixed-integer SDP solution to distributionally robust unit commitment with second order moment constraints. *CSEE J. Power Energy Syst.*, 6(2):374–383, 2020.

A   Set partitions and their cardinalities

In Section 3 we introduced the sets $\mathcal{S}_+^n(B, r)$, where we claim that for small values of n we can fully enumerate over these sets. This enumeration involves the combinatorial structure of so-called set partitions, i.e., partitions of the set of integers from 1 to n. As a result, the cardinalities of $\mathcal{S}_+^n(B, r)$ can be expressed in terms of some well-known combinatorial quantities. These quantities are presented below.

Stirling number of the second kind   A Stirling number of the second kind, see e.g., [7], counts the number of possibilities to partition a set of n elements into k non-empty subsets. It is denoted by $\left\{ {n \atop k} \right\}$. In particular, $\left\{ {0 \atop 0} \right\} = 1$, $\left\{ {0 \atop k} \right\} = 0$ for all $k > 0$ and $\left\{ {n \atop 1} \right\} = 1$ for all $n \geq 1$. For higher values of n and k, the recurrence relation $\left\{ {n+1 \atop k} \right\} = k \left\{ {n \atop k} \right\} + \left\{ {n \atop k-1} \right\}$ can be applied.

Bell number   The nth Bell number, denoted by B_n, equals the number of partitions of a set of n elements [7]. We have $B_0 = B_1 = 1$ and higher Bell numbers can be obtained by the recurrence relation $B_{n+1} = \sum_{k=0}^{n} \binom{n}{k} B_k$. Since the Bell number takes into account all possible partitions of a set (irrespective of the number of subsets), the nth Bell number is the sum over all Stirling numbers of the second kind with n fixed, i.e.,

$$B_n = \sum_{k=0}^{n} \left\{ {n \atop k} \right\}.$$

B-type Stirling number of the second kind   Now, let us consider partitions of the extended set of integers $\{-n, \ldots, n\}$. We call a partition π of $\{-n, \ldots, n\}$ a B-type (or signed) partition [51] if (1) $\pi = -\pi$, i.e., for any subset $B \in \pi$, it holds that $-B = \{-b : b \in B\} \in \pi$, and (2) π contains just one subset B_0 such that $B_0 = -B_0$, which is called the 0-subset. For example, the following partition of $\{-5, \ldots, 5\}$ is a B-type partition:

$$\{\{1, -5\}, \{2\}, \{-4, -3, 0, 3, 4\}, \{-2\}, \{-1, 5\}\},$$

that consists of 5 subsets. The B-type Stirling number of the second kind $S_B(n, k)$ counts the number of B-type partitions of $\{-n, \ldots, n\}$ that consist of exactly k pairs of subsets (excluding the 0-subset). Indeed, $S_B(n, k)$ is the B-type equivalent of the Stirling number of the second kind $\left\{ {n \atop k} \right\}$.

Dowling number   The nth Dowling number, denoted by D_n, equals the total number of B-type partitions of the set $\{-n,\ldots,n\}$, see e.g.,[2]. Since the number of subset pairs in any B-type partition of $\{-n,\ldots,n\}$ can range from 0 to n, we have

$$D_n = \sum_{k=0}^{n} S_B(n,k).$$

Thus, the nth Dowling number has the same relationship with the B-type Stirling number of the second kind as the nth Bell number has with the regular Stirling number of the second kind.

Relationship with $\mathcal{S}_+^{\mathbf{n}}(\mathbf{B},\mathbf{r})$ We now couple the above-mentioned quantities to the cardinality of the sets $\mathcal{S}_+^n(B,r)$. For $B = \{0,1\}$ and $1 \le r \le n$, the elements in $\mathcal{S}_+^n(\{0,1\},r)$ are of the form $X = \sum_{i=1}^{r} x_i x_i^\top$, where $x_i \in \{0,1\}^n$ and each x_i and x_j have non-overlapping support. Each x_i can be thought of as the characteristic vector of a subset of $\{1,\ldots,n\}$. If we add to this set an additional element 0, an element $X \in \mathcal{S}_+^n(\{0,1\},r)$ can be seen as a partition of the set $\{0,\ldots,n\}$, where the set that contains the 0-element corresponds to all integers that are not in the support of one of the x_i's. Hence, there exists a bijection between the elements in $\mathcal{S}_+^n(\{0,1\},r)$ and all partitions of a set of $n+1$ elements that consists of at most $r+1$ subsets. We conclude

$$|\mathcal{S}_+^n(\{0,1\},r)| = \sum_{k=0}^{r+1} \left\{ {n+1 \atop k} \right\}.$$

In particular, when $r = n$, we have

$$|\mathcal{S}_+^n(\{0,1\})| = |\mathcal{S}_+^n(\{0,1\},n)| = \sum_{k=0}^{n+1} \left\{ {n+1 \atop k} \right\} = B_{n+1}.$$

Now, let us consider $B = \{0,\pm 1\}$. The elements in $\mathcal{S}_+^n(\{0,\pm 1\},r)$ are of the form $X = \sum_{i=1}^{r} x_i x_i^\top$, where $x_i \in \{0,\pm 1\}^n$ and each x_i and x_j have non-overlapping support. Each x_i can be seen as the characteristic vector of a subset of $\{-n,\ldots,n\}$, where we include $-u$ if $(x_i)_u = -1$ and u if $(x_i)_u = 1$ (and we do not include u or $-u$ if $(x_i)_u = 0$). Observe that the characteristic vector $-x_i$ leads to the *same* partition of $\{-n,\ldots,n\}$. A matrix $X = \sum_{i=1}^{r} x_i x_i^\top$ then corresponds to a partition of $\{-n,\ldots,n\}$, where the elements that are not in the support of any x_i are in the same subset as 0. Therefore, each matrix X corresponds to a B-type partition of $\{-n,\ldots,n\}$. Thus, there exists a bijection between the elements in $\mathcal{S}_+^n(\{0,\pm 1\},r)$ and all B-type partitions of the set $\{-n,\ldots,n\}$ consisting of at most r subset pairs. We conclude

$$|\mathcal{S}_+^n(\{0,\pm 1\},r)| = \sum_{k=0}^{r} S_B(n,k).$$

For $r = n$, we obtain

$$|\mathcal{S}_+^n(\{0,\pm 1\})| = |\mathcal{S}_+^n(\{0,\pm 1\},n)| = \sum_{k=0}^{n} S_B(n,k) = D_n.$$

We finalize this section by presenting how the cardinalities of the above-mentioned sets evolve over n. Table 9 contains the values of $|\mathcal{S}_+^n(B,r)|$ for different B, r and n, see also [26].

Set	Cardinality						
	$n=3$	$n=4$	$n=5$	$n=6$	$n=7$	$n=8$	$n=9$
$|\mathcal{S}_+^n(\{0,1\},3)|$	15	51	187	715	2795	11051	35550
$|\mathcal{S}_+^n(\{0,1\})|$	15	52	203	877	4140	21147	115975
$|\mathcal{S}_+^n(\{0,\pm 1\},3)|$	24	115	622	3656	22724	146565	968922
$|\mathcal{S}_+^n(\{0,\pm 1\})|$	24	116	648	4088	28640	219920	1832224

Table 9: Number of matrices in $\mathcal{S}_+^n(B,r)$ for different B, r and n.