MADAM: a parallel exact solver for max-cut based on semidefinite programming and ADMM

Timotej Hrga¹ · Janez Povh^{1,2}

Received: 11 September 2020 / Accepted: 18 August 2021 / Published online: 26 August 2021
© The Author(s) 2021

Abstract

We present MADAM, a parallel semidefinite-based exact solver for Max-Cut, a problem of finding the cut with the maximum weight in a given graph. The algorithm uses the branch and bound paradigm that applies the alternating direction method of multipliers as the bounding routine to solve the basic semidefinite relaxation strengthened by a subset of hypermetric inequalities. The benefit of the new approach is a less computationally expensive update rule for the dual variable with respect to the inequality constraints. We provide a theoretical convergence of the algorithm as well as extensive computational experiments with this method, to show that our algorithm outperforms state-of-the-art approaches. Furthermore, by combining algorithmic ingredients from the serial algorithm, we develop an efficient distributed parallel solver based on MPI.

Keywords Semidefinite programming · Alternating direction method of multipliers · Maximum cut problem · Parallel computing

1 Introduction

1.1 Motivation

The Max-Cut problem is a classical NP-hard optimization problem [1, 2] on graphs with the quadratic objective function and unconstrained binary variables. During the last decades, it has attracted the interest of many researchers, from its theoretical and algorithmic perspective to its applicability in different fields, for instance

✉ Timotej Hrga
timotej.hrga@lecad.fs.uni-lj.si

Janez Povh
janez.povh@lecad.fs.uni-lj.si

¹ Faculty of Mechanical Engineering, University of Ljubljana, Ljubljana, Slovenia

² Institute of Mathematics, Physics and Mechanics Ljubljana, Ljubljana, Slovenia

mathematics, physics, and computer science [3–5]. Several exact algorithmic approaches have been proposed in the literature, including the BiqMac [5] and BiqCrunch solvers [6], which are among the best solvers and are based on semidefinite programming.

Lassere [7] has proved that the Max-Cut problem can be considered as a canonical model of linearly constrained linear and quadratic 0/1 programs. The transformation is based on the exact penalty approach and was further explored and advanced by Gusmeroli and Wiegele [8]. Even for instances of moderate size, it is considered a computational challenge to solve the Max-Cut to optimality. In practice, we typically solve such problems only approximately by using a heuristic or an approximation algorithm [9, 10]. However, to compare these algorithms and evaluate their performance, we still require the optimum solutions. Considering all this, solving increasingly large instances of Max-Cut to optimality on parallel computers is highly needed in scientific computing.

1.2 Problem formulation and notations

The central problem that we consider is the Max-Cut problem, which can be defined as follows. For a given undirected graph $G = (V, E)$ on $n = |V|$ vertices and with edge weights w_e for $e \in E$, the Max-Cut problem asks to find a bipartition of the vertices such that the sum of the weights of the edges across the bipartition is maximized. Let $A = (a_{ij})$ denote the weighted adjacency matrix with $a_{ij} = a_{ji} = w_e$ for edge $e = \{i, j\} \in E$ and zero otherwise. Encoding the partitions by vectors $z = \{0, 1\}^n$, we obtain the following unconstrained binary quadratic optimization problem formulation for Max-Cut:

$$\text{maximize} \quad \sum_{i,j}^n a_{ij} z_i (1 - z_j) = z^T L_0 z \quad \text{subject to} \quad z \in \{0, 1\}^n, \quad (1)$$

where L_0 is the Laplacian matrix of the graph defined by $L_0 = \text{diag}(Ae) - A$. By e we denote the vector of all ones. Note that due to the symmetry of the problem, we can fix the last element of vector z to zero and thus remove the last row and column of L_0 to obtain the matrix $\hat{L}_0 \in \mathbb{R}^{(n-1) \times (n-1)}$.

The inner product on the space of symmetric matrices is given by $X \bullet Y = \langle X, Y \rangle = \text{tr}(XY) = \sum_{i,j} X_{ij} Y_{ij}$ and the associated Frobenius norm is defined by $\|X\|_F = \sqrt{\langle X, X \rangle} = \sqrt{\sum_{i,j} X_{ij}^2}$. Using the property of inner product $z^T L_0 z = \langle L_0, zz^T \rangle$, we can reformulate problem (1) as:

$$\begin{aligned} &\text{maximize} \quad \begin{bmatrix} \hat{L}_0 & 0 \\ 0 & 0 \end{bmatrix} \bullet \begin{bmatrix} zz^T & z \\ z^T & 1 \end{bmatrix} \\ &\text{subject to} \quad z \in \{0, 1\}^{n-1}. \end{aligned} \quad (2)$$

We increase the dimension of the problem by one, since the last column of the solution matrix of a semidefinite relaxation will be used for determining the next branching variable.

In the following, we denote the diagonal matrix, which has v on its diagonal, by $\text{Diag}(v)$, and the vector obtained by extracting the main diagonal from the matrix X is denoted by $\text{diag}(X)$. For the given symmetric matrices A_i , $i = 1, \dots, m$, let $\mathcal{A} : \mathcal{S}_n \rightarrow \mathbb{R}^m$ denote the linear operator mapping $n \times n$ symmetric matrices to $\mathbb{R}^m$ with $\mathcal{A}(X)_i = \langle A_i, X \rangle$. Its adjoint is well known to be $\mathcal{A}^T(y) = \sum_i y_i A_i$. For some real number a , we denote its nonnegative part by $a_+ = \max\{a, 0\}$. This definition is extended to vectors as follows: if $x \in \mathbb{R}^n$, then $(x_+)_i = (x_i)_+$, for $i = 1, \dots, n$. We denote the projection of some symmetric matrix X onto the positive semidefinite cone by X_+ and its projection onto the negative semidefinite cone by X_- . More specifically, if the eigendecomposition of X is given by $X = S \text{Diag}(\lambda) S^T$ with the eigenvalues $\lambda \in \mathbb{R}^n$ and orthogonal matrix $S \in \mathbb{R}^{n \times n}$, then we have

$$X_+ = S \text{Diag}(\lambda_+) S^T, \quad X_- = S \text{Diag}(\lambda_-) S^T.$$

1.3 Related work and our contribution

Over the decades, many methods have been proposed for finding exact solutions of Max-Cut. Some of them are linear programming-based methods [4, 11], which work particularly well when the underlying graph is sparse. Other algorithms combine semidefinite programming with the polyhedral approach [12] to strengthen the basic SDP relaxation with cutting planes. Two of them, the BiqMac [5] and BiqCrunch [6] solvers, respectively, have turned out to be the best performing solvers for Max-Cut in the last decade, also when compared to commercial solvers. Both solvers utilize the branch and bound (B&B) paradigm. However, the distinction is in using different algorithms to solve the underlying SDP relaxation. Furthermore, they do not use any parallelization. For other recent computational approaches, the reader is referred to [13].

Recently, we have developed a BiqBin solver [14] with other authors for the class of binary quadratic problems with linear constraints, which includes the Max-Cut problem. This solver uses the exact penalty approach to reformulate every instance of the binary quadratic problem into an instance of the Max-Cut problem and then solves it using an enhanced version of BiqMac where the underlying SDP relaxations are tight due to the inclusion of hypermetric inequalities. Additionally, the B&B part of BiqBin has been improved compared to BiqMac, and the solver has been fully parallelized. Extensive numerical evidence shows that BiqBin outperforms BiqMac, BiqCrunch, GUROBI, and SCIP on Max-Cut instances and also on some classes of binary quadratic problems.

The main contribution of this paper is the introduction of the parallel exact solver MADAM for Max-Cut, based on the alternating direction method of multipliers (ADMM), which is introduced in [15]. More precisely, we:

- Adapt ADMM to solve SDP relaxations of Max-Cut, with a subset of hypermetric inequalities included, and provide a convergence proof. More specifically, we strengthen the basic SDP relaxation with triangle, pentagonal, and heptagonal inequalities. The key idea in solving the underlying relaxation is to introduce a

slack variable in order to eliminate solving a quadratic program over the non-negative orthant when computing the dual variable corresponding to inequality constraints. Instead, a sparse system of linear equations is solved and a projection onto the nonnegative orthant is performed.

- We propose a rounding procedure which rounds the (nearly feasible) solution of the dual problem obtained by ADMM to feasible solutions and thus provides a valid upper bound for the optimum value of Max-Cut.
- We develop a new B&B algorithm for Max-Cut which exploits the valid upper bounds derived by rounding the nearly feasible solution of ADMM.

The B&B is parallelized based on the load coordinator-worker paradigm using the Message Passing Interface (MPI) library. Scalability results for MADAM show good scaling properties when running on a supercomputer.

- We develop a C implementation of all modules and name it MADAM – Max-Cut Alternating Direction Augmented Lagrangian Method. The source code is available at

<https://github.com/HrgaT/MADAM>.

MADAM is tested against BiqMac, BiqCrunch, and BiqBin on several Max-Cut instances from the BiqMac library. Numerical results show that MADAM outperforms other approaches. The reason is that our adaption of ADMM has a better balance between the time needed to solve the SDP relaxation and the quality of the solution than other solvers.

The main difference between MADAM, BiqMac, and BiqBin is that MADAM applies ADMM to solve the underlying SDP relaxations, while BiqMac and BiqBin use the bundle method. BiqMac strengthens the basic SDP relaxation with triangle inequalities, whereas BiqBin and MADAM also include a subset of pentagonal and heptagonal inequalities, which make these relaxations harder to solve. Using ADMM in MADAM implies that in each B&B node, we have a (nearly) feasible solution for the SDP relaxation available, which is not the case with the bundle method applied to the partial Lagrangian dual. This gives us more information on how to generate feasible solutions and how to branch. Furthermore, the obtained SDP bounds are tighter, and consequently the size of the explored B&B tree is smaller compared to other solvers. Parallelization of B&B in MADAM is very similar to the parallelization of B&B in BiqBin, since both are based on the load coordinator-workers paradigm and use the MPI library for communication between the processes.

The rest of this paper is organized as follows. Section 2 deals with semidefinite relaxations of the Max-Cut used in this paper. We also describe the class of cutting planes that we use. The description of other solution approaches that are based on semidefinite programming is given in Sect. 3. In Sect. 4, we outline our new ADMM algorithm, discuss implementation details, and provide a proof of convergence. Details on how the proposed ADMM method is used within the B&B framework are given in Sect. 5, while Sect. 6 presents numerical results of the serial algorithm. In Sect. 7, we give insights on how parallelization improves the performance of the MADAM solver. We conclude with a computational experience of our method and present scalability results.

2 Semidefinite relaxation of Max-Cut

The semidefinite program (SDP) optimizes a linear objective function of the matrix over the intersection of the cone of positive semidefinite matrices with an affine space [16]. Compared to linear programming, SDP relaxations have been shown to provide tighter bounds on the optimum value for many combinatorial optimization problems. Particularly for the Max-Cut problem, Goemans and Williamson [10] achieve the celebrated 0.879 approximation ratio using the SDP relaxation.

Since we will consider the semidefinite relaxation of (2), it is more appropriate to work with $-1/1$ variables. We apply the change of variables between $z \in \{0, 1\}^{n-1}$ and $x \in \{-1, 1\}^{n-1}$, defined by $z = \frac{1}{2}(x + e)$. This can be written in matrix form as

$$\begin{bmatrix} z \\ 1 \end{bmatrix} = \begin{bmatrix} \frac{1}{2}I & \frac{1}{2}e \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ 1 \end{bmatrix} = U \begin{bmatrix} x \\ 1 \end{bmatrix}$$

and therefore

$$\begin{bmatrix} zz^T & z \\ z^T & 1 \end{bmatrix} = U \begin{bmatrix} xx^T & x \\ x^T & 1 \end{bmatrix} U^T.$$

Observe that for any $x \in \{-1, 1\}^{n-1}$, the matrix xx^T is positive semidefinite and its diagonal is equal to the vector of all ones. Using the above transformation, this allows us to write the Max-Cut problem in $-1/1$ variables as

$$\begin{aligned} &\text{maximize} && \langle L, X \rangle \\ &\text{subject to} && \text{diag}(X) = e \\ &&& X = \begin{bmatrix} xx^T & x \\ x^T & 1 \end{bmatrix} \end{aligned}$$

or equivalently

$$\begin{aligned} &\text{maximize} && \langle L, X \rangle \\ &\text{subject to} && \text{diag}(X) = e \\ &&& X \succeq 0 \\ &&& \text{rank}(X) = 1, \end{aligned} \tag{MC}$$

where

$$L = U^T \begin{bmatrix} \hat{L}_0 & 0 \\ 0 & 0 \end{bmatrix} U = \frac{1}{4} \begin{bmatrix} \hat{L}_0 & \hat{L}_0 e \\ e^T \hat{L}_0 & e^T \hat{L}_0 e \end{bmatrix} \in \mathbb{R}^{n \times n}.$$

By dropping the rank-one constraint, we obtain the basic SDP relaxation

$$\begin{aligned} &\text{maximize} && \langle L, X \rangle \\ &\text{subject to} && \text{diag}(X) = e \\ &&& X \succeq 0. \end{aligned} \tag{MC}_{\text{BASIC}}$$

Interior-point methods (IPM) turned out to be the most prominent algorithms for solving SDPs. There exist several numerical packages to solve semidefinite programs, *e.g.*, CSDP [17], MOSEK [18], SeDuMi [19], and SDPT3 [20]. Many of them are based on some variant of the primal-dual path following interior-point method, see for example [21].

Despite their expressive power, we are in practice limited to solving small and middle-size SDPs. As the size of problems and number of constraints grow, interior-point methods show poor scalability for large-scale problems. However, for special cases like the relaxation (MC_{BASIC}), the interior-point method tailored to this problem scales well, since the matrix determining the Newton direction can be efficiently constructed [21].

The bound from (MC_{BASIC}) is not strong enough to be successfully used within a branch and bound framework in order to solve larger Max-Cut problems to optimality. By adding additional equality or inequality constraints known as cutting planes, which are valid for all feasible solutions of (MC), we strengthen the upper bound. One such class of cutting planes, called triangle inequalities, is obtained as is shown below. Observe that for an arbitrary triangle with vertices $i < j < k$ in the graph, any partition of vertices cuts either 0 or 2 of its edges. Moving to the $\{-1, 1\}$ model, this leads to

$$\begin{aligned} -X_{ij} - X_{ik} - X_{jk} &\leq 1, \\ -X_{ij} + X_{ik} + X_{jk} &\leq 1, \\ X_{ij} - X_{ik} + X_{jk} &\leq 1, \\ X_{ij} + X_{ik} - X_{jk} &\leq 1, \end{aligned}$$

(see [12] and [5]). The inequalities are collected as $\mathcal{B}_{\text{TRI}}(X) \leq e$ and we end up with the following SDP relaxation:

$$\begin{aligned} &\text{maximize} && \langle L, X \rangle \\ &\text{subject to} && \text{diag}(X) = e \\ & && \mathcal{B}_{\text{TRI}}(X) \leq e \\ & && X \geq 0. \end{aligned} \quad (\text{MC}_{\text{TRI}})$$

To exploit a stronger relaxation of (MC_{TRI}), the bound is further strengthened by using pentagonal and heptagonal inequalities, which belong to the family of hypermetric inequalities, valid for any matrix X from the convex hull of rank-one matrices xx^T , for $x \in \{-1, 1\}^n$. For every integer vector b such that $e^T b$ is odd, the inequality $|x^T b| \geq 1$ holds for all $x \in \{-1, 1\}^n$. Therefore, we have $\langle bb^T, xx^T \rangle \geq 1$. The hypermetric inequality specified by the vector b is the inequality

$$\langle bb^T, X \rangle \geq 1.$$

In MADAM, we consider the subset of hypermetric inequalities generated by choosing b with $b_i \in \{-1, 0, 1\}$ and by fixing the number of non-zero entries to 3, 5, or 7. These correspond to triangle, pentagonal, and heptagonal inequalities, respectively, and we collect them as $\mathcal{B}(X) \leq e$. We obtain the following strengthening of (MC_{TRI}):

$$\begin{aligned}
& \text{maximize} && \langle L, X \rangle \\
& \text{subject to} && \text{diag}(X) = e \\
& && \mathcal{B}(X) \leq e \\
& && X \geq 0.
\end{aligned} \tag{MC}_{\text{HYP}}$$

There is a very large number of triangle, pentagonal, and heptagonal inequalities in (MC_{HYP}) . Solving this relaxation directly is intractable even for moderate values of n . In MADAM, we iteratively identify a subset of the proposed cutting planes by separating the most violated inequalities using the current approximate solution. There are $4\binom{n}{3}$ triangle inequalities, but for a given X , we can enumerate all of them and identify the most violated ones. Due to a large number of pentagonal and heptagonal inequalities, the computational overhead to evaluate all of them is prohibitive for larger instances and the separation has to be done heuristically. We use the idea proposed in BiqBin, where the separation problem is reformulated as a quadratic assignment problem of the form

$$\begin{aligned}
& \text{minimize} && \langle H, PXP^T \rangle \\
& && P \in \Pi,
\end{aligned}$$

where Π is the set of all $n \times n$ permutation matrices and H determines the type of pentagonal or heptagonal inequality specified by the vector b . The problem is then approximately solved by using simulated annealing to obtain inequality with a potentially large violation. The reader is referred to [14] for more details.

3 Other solution approaches based on semidefinite programming

In this section, we summarize other approaches for solving the Max-Cut problem to optimality, especially the BiqMac [5], BiqBin [14], and BiqCrunch [6] solvers, which are widely considered as the most powerful solvers for unconstrained binary quadratic problems.

3.1 BiqMac and BiqBin

The starting point of BiqMac is the strengthened SDP relaxation (MC_{TRI}) . By dualizing only the triangle inequality constraints, the authors obtain the convex nonsmooth partial dual function

$$f(\gamma) = e^T \gamma + \max_{\substack{\text{diag}(X) = e, \\ X \geq 0}} \langle L - \mathcal{B}_{\text{TRI}}^T(\gamma), X \rangle,$$

where γ is the nonnegative dual variable. Evaluating the dual function and computing a subgradient amounts to solving an SDP of the form $(\text{MC}_{\text{BASIC}})$, which can be

efficiently computed using an interior-point method tailored to this problem. The approximate minimizer of the dual problem

$$\begin{aligned} & \text{minimize} && f(\gamma) \\ & \text{subject to} && \gamma \geq 0 \end{aligned}$$

is then computed using the bundle method [22]. In BiqBin, the bound is further strengthened by also dualizing pentagonal and heptagonal inequalities, i.e. the bundle method is used to minimize the partial dual function of (MC_{HYP}) .

3.2 BiqCrunch

Due to the diagonal constraint in (MC_{TRI}) , it can easily be proven that all feasible matrices of this problem satisfy the condition $n^2 - \|X\|_F^2 \geq 0$, in which the equality holds if and only if the matrix has rank 1. The idea of BiqCrunch is to add a multiple of this quadratic regularization term to the objective function. They obtain the following regularized problem:

$$\begin{aligned} & \text{maximize} && \langle L, X \rangle + \frac{\alpha}{2} (n^2 - \|X\|_F^2) \\ & \text{subject to} && \text{diag}(X) = e \\ & && \mathcal{B}_{TRI}(X) \leq e \\ & && X \geq 0. \end{aligned} \tag{3}$$

By using quadratic regularization, the simplified dual problem

$$\begin{aligned} & \text{minimize} && \frac{1}{2\alpha} \left\| [L - \text{Diag}(y) - \mathcal{B}_{TRI}^T(z)]_+ \right\|_F^2 + e^T y + e^T z + \frac{\alpha}{2} n^2 \\ & [0.5em] \text{subject to} && y \text{ free}, z \geq 0 \end{aligned}$$

produces a family of upper bounds dependent on the penalty parameter α . Note that the objective function is convex and differentiable, see [3]. The function value and gradient are evaluated by computing partial spectral decomposition. This enables that for a fixed α , the dual function is minimized using the L-BFGS-B algorithm [23] from the family of quasi-Newton methods. Reducing the penalty parameter α increases the tightness of the upper bound. This is due to the decreased impact of the regularization term in the primal objective function of (3).

4 Alternating direction method of multipliers

Several algorithmic alternatives to interior-point methods have been proposed in the literature to solve semidefinite programs [24–26]. The alternating direction method of multipliers (ADMM) has been studied over the last decades due to its wide range of applications in a number of areas and its capability of solving large-scale problems [15, 27, 28]. ADMM is an algorithm that solves convex optimization problems by breaking them into smaller pieces, which are easier to

handle. This is achieved by an alternating optimization of the augmented Lagrangian function. The method deals with the following problems with only two blocks of functions and variables:

$$\begin{aligned} & \text{minimize} && f_1(x_1) + f_2(x_2) \\ & \text{subject to} && A_1x_1 + A_2x_2 = b \\ & && x_1 \in \mathcal{X}_1, x_2 \in \mathcal{X}_2. \end{aligned} \quad (4)$$

Let

$$\mathcal{L}_\rho(x_1, x_2, y) = f_1(x_1) + f_2(x_2) + y^T(A_1x_1 + A_2x_2 - b) + \frac{\rho}{2} \|A_1x_1 + A_2x_2 - b\|_2^2$$

be the augmented Lagrangian function of (4) with the Lagrange multiplier y and a penalty parameter $\rho > 0$. Then ADMM consists of iterations

$$\begin{aligned} x_1^{k+1} &= \arg \min_{x_1 \in \mathcal{X}_1} \mathcal{L}_\rho(x_1, x_2^k, y^k) \\ x_2^{k+1} &= \arg \min_{x_2 \in \mathcal{X}_2} \mathcal{L}_\rho(x_1^{k+1}, x_2, y^k) \\ y^{k+1} &= y^k + \rho(A_1x_1^{k+1} + A_2x_2^{k+1} - b). \end{aligned} \quad (5)$$

For details the reader is referred to [27] and references therein.

ADMM has already been successfully applied for solving semidefinite relaxations of combinatorial optimization problems. The Boundary Point Method [28] is an augmented Lagrangian method applied to the dual SDP in standard form. In the implementation, alternating optimization is used since only one iteration is performed in the inner loop. The method is currently one of the best algorithms for computing the theta number of a graph [29]. However, the drawback of the method is that it can only solve equality constrained semidefinite programs. Wen et al. [15] present ADMM for general semidefinite programs with equality and inequality constraints. Compared to our approach, the drawback of their method is in solving the quadratic program over the nonnegative orthant when computing the dual variable corresponding to inequality constraints. Furthermore, they use multi-block ADMM for which the proof of convergence is not presented.

4.1 ADMM method for (MC_{HYP})

By introducing the slack variable s , the problem (MC_{HYP}) is equivalent to

$$\begin{aligned} & \text{maximize} && \langle L, X \rangle \\ & \text{subject to} && \text{diag}(X) = e \\ & && \mathcal{B}(X) + s = e \\ & && X \succeq 0, s \geq 0. \end{aligned} \quad (6)$$

One can easily verify that its dual problem can be written as

$$\begin{aligned}
& \text{minimize} && e^T y + e^T t \\
& \text{subject to} && L - \text{Diag}(y) - \mathcal{B}^T(t) + Z = 0 \\
& && u - t = 0 \\
& && y, t \text{ free, } Z \geq 0, u \geq 0,
\end{aligned} \tag{7}$$

where y and t are dual variables associated respectively with the equality constraints, whereas u and Z are dual multipliers to the conic constraints.

Let m denote the number of hypermetric inequalities considered in (MC_{HYP}) . For fixed $\rho > 0$, consider the augmented Lagrangian L_ρ for (7):

$$\begin{aligned}
L_\rho(y, t, Z, u; X, s) &= e^T y + e^T t + \langle L - \text{Diag}(y) - \mathcal{B}^T(t) + Z, X \rangle \\
&\quad + \frac{\rho}{2} \|L - \text{Diag}(y) - \mathcal{B}^T(t) + Z\|_F^2 + \langle u - t, s \rangle + \frac{\rho}{2} \|u - t\|_2^2 \\
&= e^T y + e^T t + \frac{\rho}{2} \|L - \text{Diag}(y) - \mathcal{B}^T(t) + Z + X/\rho\|_F^2 \\
&\quad + \frac{\rho}{2} \|u - t + s/\rho\|_2^2 - \frac{1}{2\rho} \|X\|_F^2 - \frac{1}{2\rho} \|s\|_2^2.
\end{aligned}$$

The alternating direction method of multipliers for problem (7) consists of the alternating minimization of L_ρ with respect to one variable while keeping others fixed to get y , Z , t , and u . Then the primal variables X and s are updated using the following rules:

$$y^{k+1} = \arg \min_{y \in \mathbb{R}^n} L_\rho(y, t^k, Z^k, u^k; X^k, s^k) \tag{8a}$$

$$t^{k+1} = \arg \min_{t \in \mathbb{R}^m} L_\rho(y^{k+1}, t, Z^k, u^k; X^k, s^k) \tag{8b}$$

$$Z^{k+1} = \arg \min_{Z \geq 0} L_\rho(y^{k+1}, t^{k+1}, Z, u^k; X^k, s^k) \tag{8c}$$

$$u^{k+1} = \arg \min_{s \in \mathbb{R}_+^m} L_\rho(y^{k+1}, t^{k+1}, Z^{k+1}, u; X^k, s^k) \tag{8d}$$

$$X^{k+1} = X^k + \rho(L - \text{Diag}(y^{k+1}) - \mathcal{B}^T(t^{k+1}) + Z^{k+1}) \tag{8e}$$

$$s^{k+1} = s^k + \rho(u^{k+1} - t^{k+1}). \tag{8f}$$

Let us closely look at the subproblems (8a)–(8f). The first order optimality condition for problem (8a) is

$$\nabla_y L_\rho = e + \rho y - \rho \text{diag}(L - \mathcal{B}^T(t) + Z + X/\rho) = 0.$$

Hence the computation of y is trivial

$$y = -\frac{e}{\rho} + \text{diag}(L - \mathcal{B}^T(t) + Z + X/\rho).$$

Similarly for (8b), we compute the gradient with respect to t

$$\nabla_t L_\rho = e - s - \rho(u - t) - \rho \mathcal{B}(L - \text{Diag}(y) - \mathcal{B}^T(t) + Z + X/\rho) = 0.$$

Therefore, t is the solution of the following linear system

$$(\mathcal{B}\mathcal{B}^T + I)t = \mathcal{B}(L - \text{Diag}(y) + Z + X/\rho) + u + \frac{s - e}{\rho}.$$

Note that the $\mathcal{B}\mathcal{B}^T + I$ matrix is sparse and positive definite. Hence the above linear system can be efficiently solved using the sparse Cholesky factorization.

By defining $M = L - \text{Diag}(y) - \mathcal{B}^T(t) + X/\rho$, the subproblem (8c) can be formulated as

$$\arg \min_{Z \geq 0} \|M + Z\|_F^2,$$

where the solution

$$Z = (-M)_+ = -M_- \quad (9)$$

is the projection of the $-M$ matrix onto the positive semidefinite cone.

The subproblem (8d) can be written as

$$\arg \min_{u \geq 0} \|u - t + s/\rho\|_2^2.$$

It asks for the nonnegative vector that is closest to $v = t - s/\rho$. Hence the solution is

$$u = \max(0, v) = v_+,$$

the nonnegative part of vector v .

After all variables for the dual problem (7) have been updated in alternating manner, we compute primal variables X and s using the expressions (8e) and (8f). As already observed in the Boundary Point Method [28], the update rule for X can be simplified and also computed from the spectral decomposition of matrix M . Using (8e), we get

$$\begin{aligned} X^{k+1} &= X^k + \rho(L - \text{Diag}(y^{k+1}) - \mathcal{B}^T(t^{k+1}) + Z^{k+1}) \\ &= \rho(L - \text{Diag}(y^{k+1}) - \mathcal{B}^T(t^{k+1}) + X^k/\rho + Z^{k+1}) \\ &= \rho(M - M_-) = \rho M_+. \end{aligned}$$

Similarly, we can simplify the formula for (8f)

$$\begin{aligned} s^{k+1} &= s^k + \rho(u^{k+1} - t^{k+1}) \\ &= \rho(u^{k+1} - (t^{k+1} - s^k/\rho)) \\ &= \rho(v_+ - v) = -\rho v_-. \end{aligned}$$

Hence by construction, the matrix X is positive semidefinite and the vector s is nonnegative.

The overall complexity of one iteration of the method is solving a linear system with matrix $\mathcal{B}\mathcal{B}^T + I$ and computing the partial eigenvalue decomposition of M . Compared to interior-point methods where the coefficient matrix changes in each iteration, matrix $\mathcal{B}\mathcal{B}^T + I$ remains constant throughout the algorithm and its factorization can be cached at the beginning to efficiently solve the linear system in each iteration.

The difference between Algorithm 1 and the method proposed by Wen et al. [15] lies in the update rule of the dual variable t corresponding to inequality constraints. The authors directly apply ADMM on the dual of (MC_{HYP}) . In our approach, we introduce the slack variable s resulting in an unconstrained optimization problem for variable t in (8b). This reduces the overall complexity of the algorithm by eliminating the need to solve a convex quadratic program of order m over the nonnegative orthant, especially since multiple hypermetric inequalities are added to strengthen the bound. Instead, a sparse system of linear equations is solved and a projection onto the nonnegative orthant is used.

4.2 Implementation

The above update rules ensure that during the algorithm, nonnegativity of vectors u and s , and conic constraints for matrices X and Z are maintained, as well as complementarity conditions $u^T s = 0$ and $ZX = 0$. Hence, once primal and dual feasibility are reached, the method converges to the optimal solution. To measure the accuracy of primal and dual feasibility, we use

$$r_p = \frac{\|\text{diag}(X) - e\|_2 + \|\max(\mathcal{B}(x) - e, 0)\|_2}{1 + \sqrt{n}},$$

$$[0.5em]r_D = \frac{\|L - \text{Diag}(y) - \mathcal{B}^T(t) + Z\|_F + \|u - t\|_2}{1 + \|L\|_F}.$$

We terminate our algorithm when $\max\{r_p, r_D\} < \varepsilon$, for prescribed tolerance $\varepsilon > 0$.

The performance of the method is dependent on the choice of the penalty parameter ρ . Numerical experiments show that for the problems we consider, the starting value of $\rho = 1$ or $\rho = 1.6$ is a good choice and the value is dynamically tuned during the algorithm in order to improve the practical convergence. A simple strategy to adjust the value of ρ is observing the residuals:

$$\rho^{k+1} = \begin{cases} \tau \rho^k & \text{if } \log\left(\frac{r_D}{r_p}\right) > \mu \\ \rho^k / \tau & \text{if } \log\left(\frac{r_p}{r_D}\right) > \mu \\ \rho^k & \text{otherwise,} \end{cases} \quad (10)$$

for some parameters μ and τ . In our numerical tests, we use $\mu = 0.5$ and $\tau = 1.001$. The idea behind this penalty parameter update scheme is trying to keep the primal

and dual residual norms in the same order of magnitude as they both converge to zero.

The computational time of our ADMM method is essentially determined by the number of partial eigenvalue decompositions and the efficiency of the sparse Cholesky solver, since these are the most computationally expensive steps. For obtaining positive eigenvalues and corresponding eigenvectors, we use the LAPACK [30] routine DSYEVR. For factoring the $\mathcal{B}\mathcal{B}^T + I$ matrix and then performing back-solves to get t , we use the sparse direct solver from CHOLMOD [31], a high performance library for the sparse Cholesky factorization. CHOLMOD is part of the SuiteSparse linear algebra package [32].

In the following two subsections, we elaborate on two potential issues when using ADMM, and how to resolve them. These are obtaining a safe upper bound, which can be used within the B&B algorithm, and the convergence of multi-block ADMM.

4.3 Safe upper bound

To safely use the proposed upper bound within the B&B algorithm, we need a certificate that the value of the dual function is indeed a valid upper bound for the original problem (1). To achieve this, the quadruplet (y, t, Z, u) has to be dual feasible, i.e. by assigning $t \leftarrow u$, the equation

$$L - \text{Diag}(y) - \mathcal{B}^T(u) + Z = 0$$

has to be satisfied. However, since we only approximately solve the primal-dual pair of semidefinite programs to some precision, the dual feasibility is not necessarily reached when the algorithm terminates. Note that variable y is unconstrained, whereas the conic conditions on u and Z are satisfied by construction. In the following, we describe the post-processing step we do after each computation of the bound.

Proposition 1 *Let y , t , u , and Z be the output variables computed with iteration scheme (8a) – (8f). Let $\lambda_{\min}$ denote the smallest eigenvalue of matrix $\hat{Z} := \text{Diag}(y) + \mathcal{B}^T(u) - L$. If $\lambda_{\min} \geq 0$, then the value $e^T y + e^T u$ is a valid upper bound for problem (1). Otherwise the value $e^T \hat{y} + e^T u$, where $\hat{y} = y - \lambda_{\min} e$, provides an upper bound on the largest cut. Furthermore, it always holds that the value $e^T \hat{y} + e^T u$ is larger than $e^T y + e^T u$.*

Proof After the ADMM method terminates, and by assigning $t \leftarrow u$, the quadruplet (y, t, Z, u) satisfies all the constraints of (7) within machine accuracy, only the linear constraint

$$L - \text{Diag}(y) - \mathcal{B}^T(u) + Z = 0$$

may not be satisfied. By replacing Z with $\hat{Z}$, we satisfy this linear constraint, but the positive semidefiniteness of $\hat{Z}$ might be violated. If the smallest eigenvalue $\lambda_{\min}$

of $\hat{Z}$ is nonnegative, then the new quadruplet $(y, t, \hat{Z}, u)$ is feasible for (7) and the $e^T y + e^T u$ value provides an upper bound on the size of the largest cut.

If $\lambda_{\min}$ is negative, we adjust the matrix $\hat{Z}$ to be positive semidefinite by using

$$\hat{Z} \leftarrow \hat{Z} - \lambda_{\min} I.$$

To maintain dual equality constraints, we correct the unconstrained variable as $\hat{y} = y - \lambda_{\min} e$. It is obvious that the value $e^T \hat{y} + e^T u = e^T y + e^T u - n\lambda_{\min}$ is larger than $e^T y + e^T u$. $\square$

We summarize the ADMM-based algorithm for solving the SDP relaxation (MC_{HYP}) in Algorithm 1.

Algorithm 1: ADMM method for solving SDP relaxation (MC_{HYP}) of Max-Cut

Output: Upper bound on the largest cut of a graph.

Select $\rho > 0$, $\varepsilon > 0$;

$k = 0$, $u^k = 0$, $s^k = 0$, $X^k = 0$, $Z^k = 0$;

Compute sparse Cholesky factorization of matrix $\mathcal{B}\mathcal{B}^T + I$;

for $k = 0, 1, \dots$ **do**

 Compute $y^{k+1} = -\frac{e}{\rho} + \text{diag}(L + Z^k + X^k/\rho)$;

 Solve for t^{k+1} by using cached factorization:

$$(\mathcal{B}\mathcal{B}^T + I)t^{k+1} = \mathcal{B}(L + Z^k + X^k/\rho) + u^k + \frac{s^k - e}{\rho};$$

$$v = t^{k+1} - s^k/\rho;$$

$$M = L - \text{Diag}(y^{k+1}) - \mathcal{B}^T(t^{k+1}) + X^k/\rho;$$

$$Z^{k+1} = -M_-;$$

$$u^{k+1} = v_+;$$

$$X^{k+1} = \rho M_+;$$

$$s^{k+1} = -\rho v_-;$$

 Compute $\delta = \max\{r_P, r_D\}$;

if $\delta < \varepsilon$ **then**

break;

 Update penalty parameter ρ according to (10).

Post-processing:

$$\hat{Z} = \text{Diag}(y) + \mathcal{B}^T(t) - L;$$

Compute the smallest eigenvalue $\lambda_{\min}(\hat{Z})$;

if $\lambda_{\min}(\hat{Z}) \geq 0$ **then**

return $e^T y + e^T u$;

else

return $e^T y + e^T u - n\lambda_{\min}(\hat{Z})$;

4.4 Convergence of the method

It has been recently shown [33] that multi-block ADMM is not necessarily convergent. In the theorem presented below, we show that due to the special structure of operators in the case of semidefinite relaxation of Max-Cut, we can achieve a convergent scheme by reducing it to a 2-block method. For the sake of completeness, we include the proof of convergence. We also note that Chen et al. in [33] prove the same result in a more general setting.

Theorem 1 *The sequence $\{(X^k, s^k, y^k, t^k, Z^k, u^k)\}$ generated by Algorithm 1 from any starting point $(X^0, s^0, y^0, t^0, Z^0, u^0)$ converges to solutions (X^*, s^*) , (y^*, t^*, Z^*, u^*) of the primal-dual pair of semidefinite programs (6) and (7).*

Proof The convergence of our multi-block ADMM is guaranteed due to the orthogonality relations of the operators diag and $\mathcal{B}$ and their adjoints:

$$\mathcal{B}(\text{Diag}(y)) = 0, \quad \text{diag}(\mathcal{B}^T(t)) = 0, \quad \text{for any vectors } y \in \mathbb{R}^n \text{ and } t \in \mathbb{R}^m. \quad (11)$$

In this case the multi-block ADMM reduces to a special case of the original method (5). To see this, note that orthogonality in (11) implies that the first order optimality conditions for variables y and t in (8a) and (8b) reduce to

$$\begin{aligned} \nabla_y L_\rho &= e + \rho y - \rho \text{diag}(L + Z + X/\rho) = 0 \\ \nabla_t L_\rho &= e - s - \rho(u - t) - \rho \mathcal{B}(L - \mathcal{B}^T(t) + Z + X/\rho) = 0, \end{aligned}$$

meaning that y and t are independent and are jointly minimized by regarding (y, t) as one variable. Similarly, update rules for variables Z and u , as well as for primal pair X and s , ensure that they can also be jointly minimized. By separately regarding (Z, u) and (X, s) as one variable, the iterate scheme (8a) – (8f) can be written as:

$$\begin{aligned} (y^{k+1}, t^{k+1}) &= \arg \min_{(y,t) \in \mathbb{R}^n \times \mathbb{R}^m} L_\rho((y, t), (Z^k, u^k); (X^k, s^k)) \\ [0.5em] (Z^{k+1}, u^{k+1}) &= \arg \min_{(Z,u) \in S_n^+ \times \mathbb{R}_+^m} L_\rho((y^{k+1}, t^{k+1}), (Z, u); (X^k, s^k)) \\ [0.5em] (X^{k+1}, s^{k+1}) &= (X^k, s^k) + \rho(L - \text{Diag}(y^{k+1}) - \mathcal{B}^T(t^{k+1}) + Z^{k+1}, u^{k+1} - t^{k+1}). \end{aligned}$$

Thus the convergence of Algorithm 1 is implied by the analysis in [15], which looks at 2-block ADMM as a fixed point method. $\square$

5 Branch and bound

The prominent bounds obtained by using the ADMM method for solving (MC_{HVP}) motivate us to use this method within a branch and bound framework. We call the outcome of this work the MADAM solver – Max-Cut Alternating Direction Augmented Lagrangian Method.

The Branch and Bound algorithm (B&B) is one of the standard enumerative methods for computing the global optimum of an NP-hard problem. It consists of the following ingredients:

- the bounding procedure, which provides an upper bound for the optimum value for each instance of a problem;
- the branching procedure, which splits the current problem into more problems of smaller dimensions by fixing some variables;
- a heuristic for generating a feasible solution providing a lower bound.

5.1 Bounding routine

The bounding routine of MADAM applies the ADMM algorithm, described in Algorithm 1, which for a given set of triangle, pentagonal, and heptagonal inequalities minimizes the dual function (7). In order to obtain a tight upper bound, we use the cutting-plane approach where multiple hypermetric inequalities are iteratively added and purged after each computation of the upper bound. First, the optimal solution of the basic semidefinite relaxation (MC_{BASIC}) is computed using the interior-point method. Then we separate the initial set of triangle inequalities. After the minimizer of problem (7) is obtained, we purge all inactive constraints and new violated triangles are added. The problem with an updated set of inequalities is solved and the process is iterated as long as the decrease of the upper bound is sufficiently large. After separating the most violated triangle inequalities, we add pentagonal and heptagonal inequalities in order to further decrease the upper bound. We monitor the maximum violation of triangle inequalities r_{TRI} and as soon as the number is sufficiently small, we use the heuristic from Sect. 2 to add some strongly violated pentagonal inequalities to the relaxation. Similarly, as the maximum violation of pentagonal inequalities r_{PENT} drops below some threshold, new heptagonal inequalities are separated and added to the relaxation. In our numerical tests, we use the thresholds $r_{\text{TRI}} < 0.2$ and $r_{\text{PENT}} < 0.4$, and each time separate at most $10 \cdot n$ violated triangle inequalities. For pentagonal and heptagonal inequalities, a different strategy is used. The initial number of these cutting planes is small, but each time the node is not pruned, this number is increased by 200. A description of our bounding routine is in Algorithm 2.

Algorithm 2: Semidefinite bounding procedure of MADAM

```

solve the basic SDP relaxation ( $\text{MC}_{\text{BASIC}}$ );
separate initial set of triangle inequalities;
while upper bound decreases significantly do
    use ADMM method to obtain approximate maximizer  $X$  of ( $\text{MC}_{\text{HYP}}$ );
    remove hypermetric inequalities that are not active;
    add new hypermetric inequalities that are violated;

```

The warm-start strategy is the standard approach taken with ADMM methods to ensure faster convergence. After solving (MC_{HYP}) approximately at each iteration, the inactive hypermetric inequalities are removed and new hypermetric inequalities are added. This yields a new relaxed instance with a different set of inequality constraints. Thus, we can save the variables from the output of the ADMM at the k -th iteration, and then use them to initialize the variables for the ADMM at the $(k + 1)$ -th iteration. The previous ADMM iterates give a good enough approximation to result in far fewer iterations to compute the updates than if the method were started at zero or some other initial values. This is especially the case when the variables are near optimal and the method has almost converged. Furthermore, this strategy works particularly well in our case since only the variables s and u corresponding to cutting planes need to be adjusted to fit with the new operator $\mathcal{B}$ associated with a different set of inequality constraints.

We keep the values s_i and u_i the same if the i -th inequality stays, otherwise we initialize the values with zero.

In order to improve the performance of MADAM, we can stop the bounding routine when we detect that we will not be able to prune the current B&B node. BiqCrunch terminates the bounding routine if the difference between consecutive bounds is less than some prescribed parameter and if the gap between the upper and lower bounds is still large. In MADAM, we borrow an idea from BiqMac. After some cutting plane iterations, we make a linear forecast to decide whether it is worth doing more iterations. If the gap cannot be closed, we terminate the bounding routine, branch the current node, and start evaluating new subproblems.

5.2 Branching rules

In MADAM, we use two branching strategies which are based on the approximate solution matrix X obtained from Algorithm 1. Once the ADMM method terminates, the last column of X , with the exception of the last element, is extracted. Due to the diagonal constraint and positive semidefiniteness of feasible matrices of (MC_{HYP}) , all the entries lie in the interval $[-1, 1]$. By using the transformation $z = \frac{1}{2}(x + e)$, a vector with entries in $[0, 1]$ is obtained. We can follow different strategies based on the information we get from z . Similarly to BiqCrunch and BiqBin, the decision on which variable z_i to branch on is based on the following two strategies:

- i most-fractional: we branch on the vertex i for which the variable z_i is closest to 0.5;
- ii least-fractional: we branch on the vertex i for which the variable z_i is furthest from 0.5;

With this, we create two new subproblems, one where z_i is fixed to 0, and the other where it is fixed to 1. The subproblems correspond to nodes in the B&B tree and are added to the priority queue of unexplored problems. Priority is based on the upper bound computed by the ADMM method. When selecting the next subproblem, a node with the largest upper bound is evaluated first.

5.3 Generating feasible solutions

For generating high quality feasible solutions of the Max-Cut problem, we use the optimum solution X of (MC_{HYP}) in the Goemans-Williamson rounding hyperplane technique [10]. By using factorization $X = V^T V$ with column vectors v_i of V and selecting some random vector r , one side of the cut is obtained as $\{i \mid v_i^T r \geq 0\}$. The cut vector x obtained from this heuristic is then further improved by flipping the vertices. Note that in our case the factorization of X is already available from the bounding routine, which computes the eigenvalue decomposition of matrix M . To summarize, for generating good cuts, we use the following scheme:

- i Use the Goemans-Williamson rounding hyperplane technique to generate cut vector x from minimizer of (MC_{HYP}) .
- ii The cut x is locally improved by checking all possible moves of a single vertex to the opposite partition block.

The process is repeated several times with different random vector r , due to its low computational cost (in practice we repeat it n times, where n is the size of the input graph). Numerical experiments show that by using the proposed heuristic, the optimum solution is usually found already in the root node.

5.4 Strategy for faster enumeration of the B&B tree

To obtain tight upper bounds on maximum cuts, we strengthen the basic SDP relaxation with a subset of hypermetric inequalities. Adding these cutting planes to the model and solving the relaxation in each B&B node is computationally expensive and not necessary, especially if one can not prune the node. We use the strategy proposed in BiqBin, where we check (before including the hypermetric inequalities) whether we will be able to prune the current node, or whether it is better to branch.

Firstly, in the root node, we cache the bound OPT_{SDP} of the basic SDP relaxation (MC_{BASIC}) and compute the bound OPT_{HYP} by iteratively adding violated hypermetric inequalities. Let

$$diff = OPT_{SDP} - OPT_{HYP}$$

denote the difference between optimal values of both relaxations and let lb denote the current lower bound. Secondly, at all other nodes, we only compute the basic SDP relaxation. If the condition

$$OPT_{SDP} \leq lb + diff + 1 \quad (12)$$

is satisfied, meaning we are close enough to the lower bound, we estimate that by adding the cutting planes to compute the tighter bound OPT_{HYP} , we will prune the node. This idea helps efficiently traverse the B&B tree, and we only invest time into the bounding routine when it is really needed. Numerical experiments show that overall, this strategy produces more B&B nodes than necessary, but the performance of the algorithm is immensely improved, since in the first few levels of the B&B tree, only the basic SDP relaxation needs to be computed.

6 Numerical results: serial algorithm

We compare our serial version of MADAM with three state-of-the-art semidefinite based exact solvers for Max-Cut, BiqMac [5], BiqCrunch [6], and BiqBin [14]. All solvers are written in the C programming language. Several dense graphs available in the BiqMac library [34] with 100 vertices and integer edge weights are used to test the algorithms.

First, we take the pw05_100.0 problem and plot the convergence curves for the bounding routines in the BiqCrunch, BiqMac, BiqBin, and MADAM solvers in the root nodes of the corresponding B&B trees. Figure 1 depicts the decrease of the dual function values in the course of the bound computation. Note that we have moved the BiqMac curve slightly to the right to make a better distinction between the BiqBin and the BiqMac curves, since they were strongly overlapping at the beginning. We also omit the first part of the BiqCrunch curve, since at the beginning, the penalty parameter α is large, and the bounding routine produces weak bounds.

The improvements of the serial version of MADAM over the other three solvers are due to using the SDP relaxation that produces a tighter upper bound, and in using the ADMM method instead of the bundle method or the quasi-Newton method. BiqMac and BiqCrunch use only triangle inequalities to strengthen the basic SDP relaxation, whereas BiqBin and MADAM also add pentagonal and heptagonal inequalities. Consequently, the bound on the maximum cut is tighter. The exact minimizer of the nonsmooth partial dual function of BiqBin is difficult to reach using the bundle method. On the other hand, it is often the case that ADMM converges to modest accuracy (which is sufficient for our application) in a small number of iterations. Although BiqBin and MADAM use the same SDP relaxation to compute the upper bound, there is a fundamental difference in the routine that separates new violated inequalities. After using the ADMM method, we obtain the approximate maximizer X of (MC_{hyp}) . This matrix is used to add new cutting planes to the relaxation. However, this information is not available from the bundle method. Instead, a convex combination of the bundle matrices is used as the input of the separation routine. For this reason, the MADAM solver is more successful in identifying promising cutting planes and consequently the upper bound is tighter. Furthermore, in our approach we reduce the complexity of ADMM by introducing a slack variable. Without it, the main cost in each iteration would

Fig. 1 Comparison of the bounding routine of BiqCrunch, BiqMac, BiqBin, and MADAM in the root node on the pw05_100.0 problem

consist of the projection onto the positive semidefinite cone, and solving a quadratic program over the nonnegative orthant with m variables, where m denotes the number of cutting planes. Instead of solving a quadratic program, we need to compute the solution of the sparse linear system and take the nonnegative part of the linear combination of the obtained vector and slack variable. This refinement further increases the performance of ADMM.

Next, we compare the running times of the solvers. Let lb denote the best lower bound given by the heuristic, and ub is the upper bound from the bounding procedure. Using the fact that we know that the optimal cut is an integer, we can prune the node of the branch and bound tree if the condition

$$ub < lb + 1$$

is satisfied. For BiqMac, BiqCrunch, and BiqBin, we used their default settings. Each test instance is solved with both branching rules that the solvers provide. The one that produced the faster time of execution is used in the tables presenting the numerical results. For MADAM, the experiments show that the most-fractional rule dominates over the least-fractional rule, i.e. produces faster execution times for most of the graphs, and is therefore set as the default strategy. We also run the algorithm with different penalty parameter values ρ found in the literature [15]. The value of the penalty parameter is set to 1.6 for graphs from families g05 and pm1d, while for graphs from families w05, w09, pw05, and pw09 this value is decreased to 1, since we have found this works well in our code. All the computations are done on a cluster consisting of 24-core nodes featuring two Intel Xeon E5-2680V3 running at 2.5 GHz with 64 GB of RAM. The code is compiled against OpenBLAS and LAPACK. In Tables 1, 2, 3, 4, we report the needed number of B&B nodes and running time to solve the problem to optimality. Observe that our solver completely outperforms other approaches. For many instances, we significantly reduce the running time of the current best solver. The reason for efficiency is the simplicity and speed of the proposed bounding routine.

Table 1 CPU times (s) and B&B nodes to solve g05 problems

Problem	BiqMac		BiqCrunch		BiqBin		MADAM - serial	
	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes
g05_100.0	555.16	531	355.58	329	166.93	269	98.33	195
g05_100.1	3547.17	3643	1933.51	1751	955.21	964	494.10	705
g05_100.2	115.87	127	125.57	95	52.76	83	40.07	43
g05_100.3	1308.85	1215	602.59	643	324.56	655	129.60	497
g05_100.4	71.03	69	44.35	33	18.30	15	9.71	11
g05_100.5	116.16	129	111.12	109	37.37	49	28.63	31
g05_100.6	177.22	193	139.52	99	51.05	67	29.52	47
g05_100.7	332.35	337	282.33	329	145.12	121	75.31	73
g05_100.8	291.28	275	198.62	163	72.40	73	35.78	67
g05_100.9	321.10	277	232.88	229	104.52	149	47.34	101

Table 2 CPU times (s) and B&B nodes to solve pm1d problems

Problem	BiqMac		BiqCrunch		BiqBin		MADAM - serial	
	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes
pm1d_100.0	917.88	867	821.70	1011	317.62	507	138.36	287
pm1d_100.1	2191.20	2113	1302.31	1073	510.81	611	267.08	383
pm1d_100.2	1458.67	1357	1003.45	837	369.62	707	185.09	585
pm1d_100.3	298.95	285	273.41	281	103.10	141	51.05	81
pm1d_100.4	1188.23	1113	756.02	607	294.30	473	142.01	329
pm1d_100.5	260.25	255	294.19	211	90.58	73	57.65	61
pm1d_100.6	248.40	217	182.62	127	72.72	117	36.80	59
pm1d_100.7	80.00	109	104.40	103	36.80	29	25.80	17
pm1d_100.8	107.25	97	57.18	39	21.66	23	12.79	21
pm1d_100.9	382.85	347	294.32	243	112.32	135	55.55	97

Table 3 CPU times (s) and B&B nodes to solve w problems

Problem	BiqMac		BiqCrunch		BiqBin		MADAM - serial	
	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes
w05_100.0	735.17	667	524.39	401	194.08	367	114.82	171
w05_100.1	182.31	187	157.68	95	69.11	61	41.82	31
w05_100.2	137.86	139	104.17	63	43.36	31	27.14	23
w05_100.3	532.43	473	383.18	439	155.61	213	120.45	125
w05_100.4	921.26	849	564.00	383	240.62	273	168.03	191
w05_100.5	969.22	861	532.27	337	199.36	147	150.84	155
w05_100.6	185.23	165	120.55	81	53.42	67	41.32	45
w05_100.7	138.66	117	81.11	41	39.22	23	28.79	13
w05_100.8	799.63	693	514.31	387	194.68	237	150.44	193
w05_100.9	98.13	97	58.88	31	23.88	19	17.70	9
w09_100.0	474.91	419	393.72	255	129.50	143	119.88	109
w09_100.1	2538.66	2553	2447.42	1667	829.93	811	635.93	641
w09_100.2	1043.74	973	787.44	531	289.20	449	222.66	311
w09_100.3	1242.63	1151	1098.87	819	336.90	619	263.81	435
w09_100.4	575.05	541	575.57	361	194.62	183	138.72	155
w09_100.5	35.84	17	17.29	7	9.86	7	10.64	1
w09_100.6	81.73	101	110.93	53	27.36	23	35.35	15
w09_100.7	318.65	269	287.99	187	87.26	83	76.59	75
w09_100.8	174.95	199	188.65	89	87.18	51	76.31	83
w09_100.9	504.25	449	358.77	263	119.36	149	113.38	207

Table 4 CPU times (s) and B&B nodes to solve pw problems

Problem	BiqMac		BiqCrunch		BiqBin		MADAM - serial	
	Time	Nodes	Time	Nodes	Time	Nodes	Time	Nodes
pw05_100.0	1657.05	1553	1155.51	999	487.22	805	274.04	629
pw05_100.1	481.93	427	384.57	327	162.50	243	101.95	145
pw05_100.2	709.90	687	454.90	359	194.99	269	125.52	203
pw05_100.3	132.51	121	129.83	93	40.96	47	30.94	33
pw05_100.4	793.85	777	585.08	411	222.35	263	133.66	179
pw05_100.5	190.30	175	123.76	85	51.56	73	33.14	47
pw05_100.6	1401.88	1385	976.62	711	312.49	355	217.01	269
pw05_100.7	386.94	357	276.98	179	86.87	71	57.71	45
pw05_100.8	98.62	121	101.21	49	25.15	17	23.05	13
pw05_100.9	323.26	307	321.72	327	113.98	129	65.91	97
pw09_100.0	579.02	527	457.12	323	171.56	227	107.15	163
pw09_100.1	821.02	753	706.71	543	248.79	311	153.76	231
pw09_100.2	292.42	275	248.28	175	88.34	109	46.98	127
pw09_100.3	153.66	153	170.37	113	52.44	39	34.71	27
pw09_100.4	394.53	373	310.72	213	109.10	147	73.68	93
pw09_100.5	489.02	519	513.91	321	295.02	145	129.74	111
pw09_100.6	390.50	337	278.82	241	113.95	131	72.65	77
pw09_100.7	1010.22	963	824.16	589	271.59	293	182.11	193
pw09_100.8	305.12	295	287.91	171	134.71	75	86.93	59
pw09_100.9	260.41	255	193.98	119	72.06	69	68.10	57

7 Parallel algorithm

In this section, we describe how the algorithmic ingredients from the serial B&B algorithm are combined into a parallel solver which utilizes the distributive memory parallelism. OpenMP (Open Multi-Processing) and MPI (Message Passing Interface) are two widely used paradigms available for parallel computing [35, 36]. Algorithms based on B&B can be parallelized on different levels. For example, shared memory parallelization using OpenMP can be applied to speed up the computation of bounds, and secondly, exploration of different branches of the B&B tree can be done in parallel using the distributive scheme. In this case, several workers, each with its own memory, are evaluating B&B nodes concurrently and using MPI to share important information via messages. This is the approach we have taken.

The load coordinator–worker paradigm with distributed work pools is applied, in which one process becomes the master process carefully managing the status of each worker, while the remaining processes concurrently explore the branches of the B&B tree. Each worker has its own local priority queue of subproblems and the work is shared when one of them becomes idle. The master process keeps track of the status of each worker and acts as a load coordinator receiving

messages and, based on their content, acts in the appropriate manner. In order to distinguish between different workers, each process is given a fundamental identifier, called rank of the worker.

At the start of the algorithm, the master process reads and broadcasts the graph instance to the workers. It is important that every process has the knowledge of the original graph, since construction of subproblems via branching or via received MPI messages is done based on this information. All the data about the B&B nodes is encoded as an MPI structure, which is used in communication between different workers in order to efficiently exchange and construct the subproblems.

Next, the load coordinator evaluates the root node and distributes the current best lower bound. After the bounding step, two new subproblems are generated, which are then sent to the first two idle processes. Afterwards, the load coordinator monitors the status of workers (idle or busy), counts the number of B&B nodes, sends the ranks of free workers, and distributes the best solution found thus far.

After the initialization phase, the load coordinator waits for different types of messages sent by the workers. Firstly, if the worker's local queue is empty, the message is sent informing the master that the process is idle and it can receive further work. Secondly, during the algorithm, working processes generate multiple candidates for optimal solutions. The master node keeps track of the current best optimal value and solution. When a new solution is received, the value is compared, updated if necessary, and distributed back during the communication phase. And thirdly, the master process helps the computing processes with work sharing by sending messages containing the ranks of idle workers. After a worker computes the lower and upper bounds, it compares these values to see if this branch of the B&B tree can be safely pruned or if further branching is needed, and the construction of new subproblems takes place. In the latter case, a request message is sent to the load coordinator, asking for idle processes to share with one of the new generated subproblems and/or subproblems left in the queue from previous branching processes. If no idle worker is available, the generated subproblems remain in the worker's queue and the work continues locally. Otherwise, subproblems are encoded and sent to available idle workers. This is also where exchange of the current best lower bound happens.

When all the workers become idle (all local queues are empty), the master process sends a message to finish, and the algorithm terminates.

Lastly, we elaborate on how efficient distribution of the subproblems to the workers is achieved and how the idle time is reduced. If the number of available workers is large, we need to reach a certain depth of the B&B tree in order for the processes to receive some work. Until this happens in the algorithm, workers are idle and resources are wasted. To fully exploit all the available resources, we need a strategy for the worker processes to start evaluating the nodes of the B&B tree as soon as possible. This is again where the strategy using the variable *diff* from Sect. 5 benefits the algorithm. After the computation of the upper bound in the root node, the load coordinator distributes the value *diff* to the workers. When the first two idle processes evaluate the generated subproblems, the value of the basic SDP relaxation is typically such that the condition (12) is not satisfied. Thus, on the first few levels of the B&B tree, the workers compute only the basic SDP bound to faster evaluate the node, and idle processes quickly receive the generated subproblems.

Table 5 Numerical results obtained with the parallel version of MADAM

Problem	MADAM - serial	MADAM - parallel				
		Number of workers				
		3	6	12	24	48
g05_100.1	494.10	206.37	109.97	68.53	37.25	23.85
pm1d_100.1	267.08	100.23	53.51	30.98	17.53	10.92
w09_100.1	635.93	226.02	125.19	72.43	41.37	26.46
w09_100.2	222.66	75.64	43.17	26.80	14.64	11.02
w09_100.3	263.81	106.76	57.96	34.42	19.75	15.17
pw05_100.0	274.04	95.95	50.59	29.16	17.09	10.71
pw05_100.6	217.01	84.11	51.39	26.36	18.97	15.46

Fig. 2 The plots depict scalability of MADAM solver for the w09_100.1 problem. Due to increased number of workers, the computational time drops. However, larger number of workers also increases communication costs, which results in the divergence of the curve from the ideal scaling curve (right)

7.1 Numerical results: parallel algorithm

We present the numerical results comparing the serial with the parallel version of the MADAM solver. As a Message Passing Interface library, we use Open MPI to exchange the data between the computing cores. Among the graphs that are used to test the sequential algorithm, we consider only those for which the solver needs more than 200 seconds to compute the optimal solution. We solve the same problem

with increasing number of cores and compare the times against the serial algorithm. The results are reported in Table 5.

For the hardest instance w09_100.1 for the serial solver, we also present the scalability plot to demonstrate the efficiency of the parallel version. Due to the increased number of workers, the computational time drops. However, the larger number of workers also increases communication costs since multiple request are sent to the load coordinator. This is shown in Fig. 2 as the curve diverges from the ideal scaling curve.

Next, we use `rudy` [37], a machine independent graph generator, to construct new Max-Cut instances with 180 vertices, similar to the ones in the `BiqMac` library. The parameters for generating the graphs can be found in [38] in Appendix B. Table 6 presents numerical results obtained by applying the parallel solver using 240 cores for all instances. We report the maximum cut values and the number of B&B nodes needed to solve the problems to optimality.

8 Conclusions and future work

In this paper, we have presented an efficient exact solver `MADAM` for the Max-Cut problem that applies the alternating direction method of multipliers to efficiently compute high-quality SDP-based upper bounds. This is due to the small computational complexity per iteration of the proposed method, since it essentially consists of solving one sparse system of linear equations and projection onto the nonnegative orthant and positive semidefinite cone. Furthermore, tight upper bounds are obtained due to the inclusion of a subset of hypermetric inequalities. Numerical results show that the `MADAM` solver outperforms state-of-the-art approaches. By combining the elements of the serial algorithm and the new strategy for a faster exploration of the B&B tree, we have developed a parallel solver based on MPI that utilizes the load coordinator-worker paradigm. We have shown that the `MADAM` solver scales very well and that we can greatly reduce the time needed to solve the Max-Cut problem to optimality and increase the size of instances that can be solved in a routine way.

In the future, we intend to increase the performance of the algorithm by using one-sided MPI communication, thus enabling a complete removal of the master process from the algorithm and a more efficient exchange of the messages. We will also further improve the performance of `MADAM` by passing cutting planes and the output of the bounding routine from parent to children nodes. We believe that this will help accelerate the ADMM method and the search for prominent inequalities.

Table 6 Numerical results obtained with the parallel version of MADAM for graphs with 180 vertices

Problem	MADAM - parallel		
	Solution	Time	Nodes
g05_180.0	4507	671.23	148617
g05_180.1	4500	670.63	137665
g05_180.2	4503	1116.24	281215
g05_180.3	4495	3706.61	786457
g05_180.4	4492	5209.59	1556485
g05_180.5	4488	8964.00	2333997
g05_180.6	4489	10542.41	2298681
g05_180.7	4490	8952.15	2168339
g05_180.8	4490	7431.42	2413043
g05_180.9	4492	4092.88	1083549
pm1d_180.0	884	9907.77	1227367
pm1d_180.1	930	1098.64	164965
pm1d_180.2	865	374.01	54539
pm1d_180.3	866	442.40	73171
pm1d_180.4	843	1567.66	215171
pm1d_180.5	785	7697.99	1916281
pm1d_180.6	747	963.60	206105
pm1d_180.7	933	631.29	49671
pm1d_180.8	902	5146.23	958197
pm1d_180.9	950	3892.91	1377653
w05_180.0	4126	4615.81	243631
w05_180.1	4160	19464.53	2826185
w05_180.2	3697	2708.77	377159
w05_180.3	3718	1894.69	235077
w05_180.4	3777	6718.74	830525
w05_180.5	3883	4640.78	534895
w05_180.6	4233	10021.08	1806407
w05_180.7	3830	9354.87	1178125
w05_180.8	3801	4440.65	550265
w05_180.9	3906	3715.36	576693
w09_180.0	5550	13585.11	1204445
w09_180.1	5397	8165.05	905925
w09_180.2	4710	6399.73	613239
w09_180.3	5003	16577.92	1781313
w09_180.4	5267	8915.48	1920259
w09_180.5	4977	3402.87	237069
w09_180.6	5744	7312.46	954457
w09_180.7	5388	33339.01	4722991
w09_180.8	4907	52024.30	12930415
w09_180.9	4735	14737.37	1777259
pw05_180.0	25536	22309.20	3324025
pw05_180.1	25301	1111.58	146393

Table 6 (continued)

Problem	MADAM - parallel		
	Solution	Time	Nodes
pw05_180.2	25260	6354.55	1036951
pw05_180.3	25514	411.39	30617
pw05_180.4	25128	9534.08	1185213
pw05_180.5	25338	4751.92	481553
pw05_180.6	25200	20387.66	3626241
pw05_180.7	25373	707.48	131639
pw05_180.8	25215	3118.33	395689
pw05_180.9	25608	20099.60	2309945
pw09_180.0	43070	9750.36	2214985
pw09_180.1	42820	2441.80	462807
pw09_180.2	42736	2820.30	195949
pw09_180.3	42880	3489.30	473039
pw09_180.4	42784	1619.40	197847
pw09_180.5	42969	9087.99	748497
pw09_180.6	42934	5775.62	729985
pw09_180.7	43210	13250.48	2430053
pw09_180.8	42806	26826.61	4179147
pw09_180.9	43438	322.01	39889

Acknowledgements The authors would like to thank the anonymous referees for their careful reading of the paper and for their constructive comments, which are greatly appreciated. This work was supported by the Slovenian Research Agency: projects N1-0057, N1-0071, J1-1691, J1-8130, J1-8132. The first author also acknowledges the financial support from the Slovenian Research Agency under the program for young researchers (MR+).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Garey, M.R., Johnson, D.S.: *Computers and intractability*, volume 29. wh freeman New York, (2002)
2. Karp, R.M.: Reducibility among combinatorial problems. In: *Complexity of computer computations*, pp. 85–103. Springer (1972)
3. Krislock, N., Malick, J., Roupin, F.: Improved semidefinite bounding procedure for solving max-cut problems to optimality. *Math. Program.* **143**(1–2), 61–86 (2014)

4. Liers, F., Jünger, M., Reinelt, G., Rinaldi, G.: Computing exact ground states of hard ising spin glass problems by branch-and-cut. *New Optim. Algorithms Phys.* **50**(47–68), 6 (2004)
5. Rendl, F., Rinaldi, G., Wiegele, A.: Solving max-cut to optimality by intersecting semidefinite and polyhedral relaxations. *Math. Program.* **121**(2), 307 (2010)
6. Krislock, N., Malick, J., Roupin, F.: BiqCrunch: A semidefinite branch-and-bound method for solving binary quadratic problems. *ACM Trans. Math. Softw. (TOMS)* **43**(4), 1–23 (2017)
7. Lasserre, J.B.: A MAX-CUT formulation of 0/1 programs. *Oper. Res. Lett.* **44**(2), 158–164 (2016)
8. Gusmeroli, N., Wiegele, A.: EXPEDIS: An Exact Penalty Method over Discrete Sets. *arXiv pre-print arXiv:1912.09739*, (2019)
9. Dunning, I., Gupta, S., Silberholz, J.: What works best when? A systematic evaluation of heuristics for Max-Cut and QUBO. *INFORMS J. Comput.* **30**(3), 608–624 (2018)
10. Goemans, M.X., Williamson, D.P.: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM (JACM)* **42**(6), 1115–1145 (1995)
11. De Simone, C., Rinaldi, G.: A cutting plane algorithm for the max-cut problem. *Optim. Methods Softw.* **3**(1–3), 195–214 (1994)
12. Helmberg, C., Rendl, F.: Solving quadratic (0, 1)-problems by semidefinite programs and cutting planes. *Math. Program.* **82**(3), 291–315 (1998)
13. Palagi, L., Piccialli, V., Rendl, F., Rinaldi, G., Wiegele, A. (eds.): : Computational approaches to max-cut. In: *Handbook on semidefinite, conic and polynomial optimization*, pp. 821–847. Springer (2012)
14. Gusmeroli, N., Hrga, T., Lužar, B., Povh, J., Siebenhofer, M., Wiegele, A.: BiqBin: a parallel branch-and-bound solver for binary quadratic problems with linear constraints. *arXiv pre-print arXiv:2009.06240*, (2020)
15. Wen, Z., Goldfarb, D., Yin, W.: Alternating direction augmented lagrangian methods for semidefinite programming. *Math. Program. Comput.* **2**(3–4), 203–230 (2010)
16. Wolkowicz, H., Saigal, R., Vandenberghe, L.: *Handbook of semidefinite programming: theory, algorithms, and applications*, vol. 27. Springer Science & Business Media (2012)
17. Borchers, B.: CSDP, A C library for semidefinite programming. *Optim. Methods Softw.* **11**(1–4), 613–623 (1999)
18. ApS, M.: The MOSEK optimization toolbox for MATLAB manual, (2015)
19. Sturm, J.F.: Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optim. Methods Softw.* **11**(1–4), 625–653 (1999)
20. Toh, K.C., Todd, M.J., Tütüncü, R.H.: SDPT3-a MATLAB software package for semidefinite programming, version 1.3. *Optim. Methods Softw.* **11**(1–4), 545–581 (1999)
21. Helmberg, C., Rendl, F., Vanderbei, R.J., Wolkowicz, H.: An interior-point method for semidefinite programming. *SIAM J. Optim.* **6**(2), 342–361 (1996)
22. Kiwiel, K.C.: A survey of bundle methods for nondifferentiable optimization. *Math. Appl. Jan. Ser.* **6**(6), 263–282 (1989)
23. Byrd, R.H., Peihuang, L., Nocedal, J., Zhu, C.: A limited memory algorithm for bound constrained optimization. *SIAM J. Sci. Comput.* **16**(5), 1190–1208 (1995)
24. Burer, S., Monteiro, R.D.C.: A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Program.* **95**(2), 329–357 (2003)
25. Fischer, I., Gruber, G., Rendl, F., Sotirov, R.: Computational experience with a bundle approach for semidefinite cutting plane relaxations of max-cut and equipartition. *Math. Program.* **105**(2–3), 451–469 (2006)
26. Malick, J., Povh, J., Rendl, F., Wiegele, A.: Regularization methods for semidefinite programming. *SIAM J. Optim.* **20**(1), 336–356 (2009)
27. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J., et al.: Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends® Mach. Learn.* **3**(1), 1–122 (2011)
28. Povh, J., Rendl, F., Wiegele, A.: A boundary point method to solve semidefinite programs. *Computing* **78**(3), 277–286 (2006)
29. Lovász, L.: On the shannon capacity of a graph. *IEEE Trans. Inf. Theory* **25**(1), 1–7 (1979)
30. Anderson, E., Bai, Z., Bischof, C., Blackford, L.S., Demmel, J., Dongarra, J., Du Croz, J., Greenbaum, A., Hammarling, S., McKenney, A., et al.: *LAPACK Users' guide*. SIAM, (1999)

31. Chen, Y., Davis, T.A., Hager, W.W., Rajamanickam, S.: Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate. *ACM Trans. Math. Softw. (TOMS)* **35**(3), 1–14 (2008)
32. Davis, T.A., Yifan, H.: The University of Florida sparse matrix collection. *ACM Trans. Math. Softw. (TOMS)* **38**(1), 1–25 (2011)
33. Chen, C., He, B., Ye, Y., Yuan, X.: The direct extension of ADMM for multi-block convex minimization problems is not necessarily convergent. *Math. Program.* **155**(1–2), 57–79 (2016)
34. Wiegele, A.: Biq Mac Library - A collection of Max-Cut and quadratic 0 – 1 programming instances of medium size. Preprint (2007)
35. Dagum, L., Menon, R.: OpenMP: an industry-standard API for shared-memory programming. *Comput. Sci. Eng.* **1**, 46–55 (1998)
36. Forum, M.P.: MPI: A Message-Passing Interface Standard. Technical report. , Knoxville, TN, USA (1994)
37. Rinaldi, G.: Rudy, 1998. <http://www-user.tu-chemnitz.de/~helmberg/rudy.tar.gz>
38. Wiegele, A.: *Nonlinear optimization techniques applied to combinatorial optimization problems* (2006)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.