

TECHNISCHE UNIVERSITÄT
CHEMNITZ

**LP and SDP Branch-and-Cut
Algorithms for the Minimum
Graph Bisection Problem:
A Computational Comparison**

M. Armbruster C. Helmberg
M. Fügenschuh A. Martin

Preprint 2011-6



Impressum:

Herausgeber:

Der Dekan der
Fakultät für Mathematik
an der Technischen Universität Chemnitz

Sitz:

Reichenhainer Straße 39
09126 Chemnitz

Postanschrift:

09107 Chemnitz

Telefon: (0371) 531-22000

Telefax: (0371) 531-22009

E-Mail: dekanat@mathematik.tu-chemnitz.de

Internet:

<http://www.tu-chemnitz.de/mathematik/>

ISSN 1614-8835 (Print)

LP and SDP Branch-and-Cut Algorithms for the Minimum Graph Bisection Problem: A Computational Comparison *

Michael Armbruster, Christoph Helmberg
Chemnitz University of Technology
Contact: helmberg@mathematik.tu-chemnitz.de

Marzena Fügenschuh
Beuth University of Applied Sciences Berlin
Contact: fuegenschuh@beuth-hochschule.de

Alexander Martin
University of Erlangen-Nuremberg
Contact: alexander.martin@math.uni-erlangen.de

March 15, 2011

Abstract

While semidefinite relaxations are known to deliver good approximations for combinatorial optimization problems like graph bisection, their practical scope is mostly associated with small dense instances. For large sparse instances, cutting plane techniques are considered the method of choice. These are also applicable for semidefinite relaxations via the spectral bundle method, which allows to exploit structural properties like sparsity. In order to evaluate the relative strengths of linear and semidefinite approaches for large sparse instances, we set up a common branch-and-cut framework for linear and semidefinite relaxations of the *minimum graph bisection problem*. It incorporates separation algorithms for valid inequalities of the bisection cut polytope described in a recent study by the authors. While the problem specific cuts help to strengthen the linear relaxation significantly, the semidefinite bound profits much more from separating the cycle inequalities of the cut polytope on a slightly enlarged support. Extensive numerical experiments show that this semidefinite branch-and-cut approach without problem specific cuts is a superior choice to the classical simplex approach exploiting bisection specific inequalities on a clear majority of our large sparse test instances from VLSI design and numerical optimization.

Keywords: branch and cut algorithms, cutting plane algorithms, polyhedral combinatorics, semidefinite programs

1 Introduction

Let $G = (V, E)$ be an undirected graph with node set $V = \{1, \dots, n\}$ and edge set $E \subseteq \{\{i, j\} : i, j \in V, i < j\}$. For an edge $\{i, j\}$ we will also write ij if there is no way of confusion.

*A conference version of this article appeared as [5].

For given vertex weights $f_i \in \mathbb{N} \cup \{0\}$, $i \in V$, and edge costs $w_{ij} \in \mathbb{R}$, $ij \in E$, a partition of the vertex set V into two disjoint clusters S and $V \setminus S$ with sizes $f(S) := \sum_{i \in S} f_i \leq F$ and $f(V \setminus S) \leq F$, where $F \in \mathbb{N} \cap [\frac{1}{2}f(V), f(V)]$, is called a *bisection*. The *minimum bisection problem* (MB) asks for a bisection such that the total cost of edges in the *cut* $\delta(S) := \{ij \in E : i \in S \wedge j \in V \setminus S\}$ is minimal. The problem is known to be NP-hard [17]. The polytope associated with MB and studied in [6] is the *bisection cut polytope*

$$P_B := \text{conv}\{y \in \mathbb{R}^E : y = \chi^{\delta(S)}, S \subseteq V, f(S) \leq F, f(V \setminus S) \leq F\},$$

where $\chi^{\delta(S)}$ is the incidence vector of the cut $\delta(S)$ with respect to the edge set E .

MB as well as P_B are related to other problems and polytopes already known in the literature. Obviously, the bisection cut polytope is contained in the *cut polytope* [7, 11] $P_C := \text{conv}\{y \in \mathbb{R}^{|E|} : y = \chi^{\delta(S)}, S \subseteq V\}$. If $F = f(V)$ then MB is equivalent to the max-cut problem (using the negative cost function) and $P_B = P_C$. For $F = \lceil \frac{1}{2}|f(V)| \rceil$ MB is equivalent to the *equipartition problem* [10] and the bisection cut polytope equals the *equipartition polytope* [8, 9, 31]. Furthermore, MB is a special case of the minimum node capacitated graph partitioning problem (MNCGP) [13, 14], where more than two clusters are available for the partition of the node set and each cluster has a common limited capacity. The objective in MNCGP is the same as in MB, *i.e.*, minimize the total cost of edges having endpoints in distinct clusters. Finally, we mention the *knapsack polytope* [36] $P_K := \text{conv}\{x \in \{0, 1\}^{|V|} : \sum_{v \in V} f_v x_v \leq F\}$, which plays a fundamental role in valid inequalities for P_B . Graph partitioning problems in general have numerous applications, for instance in numerics [18], VLSI-design [32], compiler-design [28] and frequency assignment [12]. A large variety of valid inequalities for the polytope associated with MNCGP is known [7, 8, 9, 13, 31] and, since MB is a special case of MNCGP, all those inequalities are also valid for P_B .

The emphasis of this work is on investigating the computational possibilities and scope offered by basic linear and semidefinite relaxations in combination with polyhedral results for large sparse graphs arising in real world applications. In particular, the three main objectives are to develop a competitive sparse semidefinite branch-and-cut approach suitable for these requirements, to explore the usefulness of the valid inequalities for P_B given in [6] in combination with linear and semidefinite relaxations, and to compare the performance of semidefinite and linear branch-and-cut approaches within a uniform setting.

In order to use the same separation routines for semidefinite and linear relaxations, the basic linear relaxation is designed to work directly with the edge formulation of P_B . For this purpose we assume, without loss of generality (w.l.o.g.), that G contains a star $K_{1,n-1}$ with central node $1 \in V$ by adding edges $1j$ of cost zero, if necessary. Because the cycle inequalities [7] together with integrality describe all cuts in G , our initial linear relaxation is based on the formulation

$$\begin{aligned} & \text{minimize} && w^T y \\ & \text{subject to} && \sum_{i=2}^n f_i y_{1i} \leq F, \\ & && f_1 + \sum_{i=2}^n f_i (1 - y_{1i}) \leq F, \\ & && \sum_{ij \in D} y_{ij} - \sum_{ij \in C \setminus D} y_{ij} \leq |D| - 1, \quad D \subseteq C \subseteq E, |D| \text{ odd, } C \text{ cycle in } G, \\ & && y \in \{0, 1\}^E. \end{aligned} \tag{1}$$

The semidefinite relaxation is based on [34], which is also the basis of the famous approximation algorithms for max-cut of Goemans and Williamson [19] and for equipartition [15]. For its formulation, consider $x \in \{-1, 1\}^V$ as indicator vector for the partition of the node set where $x_i = x_j$ if and only if the nodes i and j are on the same side of the cut. Then a feasible bisection vector x of MB satisfies $|f^T x| \leq 2F - f(V)$ and the cut vector induced by this bisection is $y_{ij} = \frac{1-x_i x_j}{2}$, *i.e.*, it has value 0 if $ij \in E$ is not in the cut and value 1 if ij is in the cut. Let $L \in \mathbb{R}^{V \times V}$, $L_{ii} = \sum_{ij \in E} w_{ij}$ for $i \in V$, $L_{ij} = -w_{ij}$ for $ij, ji \in E$ and $L_{ij} = 0$ otherwise, be the weighted Laplacian of G . The canonical semidefinite relaxation of MB is obtained by replacing xx^T by a symmetric, positive semidefinite matrix $X \succeq 0$ having its diagonal equal to the vector of all ones,

$$\begin{aligned} & \text{minimize} && \langle \tfrac{1}{4}L, X \rangle \\ & \text{subject to} && \text{diag}(X) = \mathbf{1}, \\ & && \langle ff^T, X \rangle \leq [2F - f(V)]^2, \\ & && X \succeq 0. \end{aligned} \tag{2}$$

Here and in the following $\langle \cdot, \cdot \rangle$ denotes the canonical matrix inner product, *i.e.*, for matrices $A, B \in \mathbb{R}^{m \times n}$, $\langle A, B \rangle := \sum_{i=1}^m \sum_{j=1}^n A_{ij} B_{ij}$. Interpreting X_{ij} as $x_i x_j$ the linear relation

$$y_{ij} = \frac{1 - X_{ij}}{2}$$

allows to go back and forth between $\{0, 1\}$ and $\{-1, 1\}$ representations of the cuts without loss of sparsity or structural properties like low rank representations (see, *e.g.*, [20]). In particular, any valid inequality $a^T y \leq \beta$ for P_B is easily translated into an equally sparse matrix constraint $\langle A, X \rangle \leq \beta$. While the primal semidefinite variable X is not sparse, its dual matrix variable in the dual semidefinite program to (2) inherits all structural properties of the cost and constraint matrices. The spectral bundle method [26] was designed to exploit these structural properties in solving the dual. It does not need the primal X and is thus able to handle large sparse instances. Due to its good warm start possibilities, the spectral bundle method allows an efficient primal cutting plane approach [22]; the high quality of the bounds obtained there by improving relaxation (2) with cycle inequalities encouraged this work. It was also observed in [22] that these bounds improve sometimes considerably, if the cycle inequalities are not only separated with respect to the elements X_{ij} for $ij \in E$ but with respect to a slightly enlarged support as might be obtained by adding a few appropriate virtual edges of weight 0. As this *extended support* relaxation includes only one bisection-specific inequality, namely $\langle ff^T, X \rangle \leq [2F - f(V)]^2$, it is one of the major surprise outcomes of this work that on our instances the new bisection specific cuts did *not* yield significant further improvement, quite in contrast to the strong improvements in the linear case, see Section 5.

Our instances for comparing linear and semidefinite branch-and-cut approaches are typically too large to be solved exactly, therefore we will mostly consider the size of the gap achieved after a given time limit. In the majority of the cases the semidefinite approach is the clear winner, not always because of better bounds but quite often also because of better solutions. While our implementation does not keep the full matrix X , the spectral bundle method allows to maintain a low rank approximation of X that is a good basis for efficient random hyperplane rounding methods of the Goemans'-Williamson's type, see Section 3. Almost as important as these are the possibilities offered by X to develop new criteria for branching decisions (Section 4).

A recent successful study of a combined semidefinite polyhedral branch-and-cut approach for max-cut is [35]. It considers triangle inequalities exclusively and uses a polyhedral bundle method with a semidefinite oracle that solves a dense semidefinite max-cut-relaxation by interior point methods. Their approach is suitable for rather dense graphs with up to 400 nodes, while ours is designed to be applicable for sparse graphs with up to 2000 nodes.

The paper is structured as follows. In [6] we gave a detailed analysis of P_B including several classes of new and facet-defining inequalities. We briefly summarize these results and those from the literature in Section 2 and sketch our separation procedures. Section 3 provides the necessary background of the sparse semidefinite approach of [22] for improving (2) by cutting planes with support extension and describes the semidefinite rounding heuristics employed. Aspects concerning the incorporation of this semidefinite cutting-plane approach in a branching scheme are treated in Section 4. The LP-based and the semidefinite branch-and-cut algorithm were implemented within the same framework SCIP [1] and use the same separation machinery. Numerical results for the linear and semidefinite root node relaxations are given in Section 5. The branch-and-cut results are discussed in Section 6. Conclusions and some open problems will complete the paper in Section 7. This work is an outgrowth of the dissertations [4, 16], which contain and investigate several further approaches and ideas for efficient implementation of linear and semidefinite branch-and-cut approaches in great depth.

2 Valid inequalities for P_B and their separation

A large variety of valid inequalities for the cut polytope, the equipartition polytope, and the polytope associated with MNCGP is known: cycle inequalities [7] of the cut polytope; tree, star, and cycle inequalities [8, 9] as well as suspended tree and path block cycle inequalities [11, 31] for the equipartition polytope; tree, star, cycle with ear, cycle with tails, and knapsack tree inequalities [13] valid for the polytope associated with MNCGP. Since MB is a special case of MNCGP and $P_B \subseteq P_C$ the bisection cut polytope inherits most of the valid inequalities mentioned above.

Exemplarily we cite the cycle inequalities which we will later use in both models for MB. Let the subgraph $C = (V_C, E_C)$ be a cycle in G . Let D be a subset of E_C such that $|D|$ is odd. Then the *cycle inequality*

$$\sum_{e \in D} y_e - \sum_{e \in E_C \setminus D} y_e \leq |D| - 1 \quad (3)$$

is a valid inequality for the cut polytope P_C . For separation we employ the exact polynomial approach of [7] that searches for shortest paths in an auxiliary graph constructed from two copies of the original graph with a total of $4|E|$ edges.

The cut structure implies that whenever there is a walk between two nodes of the graph with an even number of edges in the cut, the two end-nodes of the walk have to be in the same cluster. So, given a special root node r , a walk P_{rv} in G to some node v , an edge subset $H_v \subseteq P_{rv}$ of even cardinality, and an incidence vector y of a cut, if the term

$$1 - \sum_{e \in P_{rv} \setminus H_v} y_e - \sum_{e \in H_v} (1 - y_e) \quad (4)$$

evaluates to one then r and v are on the same side of the cut; for nodes in opposite clusters, it is at most zero. In [6] this is used to set up an inequality linking the cut structure and the capacity constraint on the node weights.

Proposition 1 (bisection knapsack walk inequality [6]). *Let $\sum_{v \in V} a_v x_v \leq a_0$ be a valid inequality for the knapsack polytope P_K with $a_v \geq 0$ for all $v \in V$. For a subset $V' \subseteq V$, a fixed root node $r \in V'$, walks $P_{rv} \subseteq E$, and sets $H_v \subseteq P_{rv}$ with $|H_v|$ even, the bisection knapsack walk inequality*

$$\sum_{v \in V'} a_v \left(1 - \sum_{e \in P_{rv} \setminus H_v} y_e - \sum_{e \in H_v} (1 - y_e) \right) \leq a_0. \quad (5)$$

is valid for the polytope P_B .

Given a root node r and a vector $y \in [0, 1]^{|E|}$ the optimal walks P_{rv} and subsets H_v maximizing (4) can be found in polynomial time with an algorithm that follows the one for separating cycle inequalities, see [6].

The *knapsack tree inequalities* of [13] form a special case, where the walks P_{rv} are taken from a tree (T, E_T) of G rooted at r and $H_v = \emptyset$ for all $v \in V'$,

$$\sum_{v \in T} a_v \left(1 - \sum_{e \in P_{rv}} y_e \right) \leq a_0. \quad (6)$$

Following [13], one may strengthen the coefficients of (6) to

$$\sum_{e \in E_T} \min \left\{ \sum_{v: e \in P_{rv}} a_v, \sum_{v \in T} a_v - a_0 \right\} y_e \geq \sum_{v \in T} a_v - a_0, \quad (7)$$

we call this a *truncated knapsack tree inequality*. A less obvious strengthening exploits the dependence of the coefficients in (7) on the choice of the root node, which we express by the notation

$$\alpha_0 := \sum_{v \in T} a_v - a_0, \quad \alpha_e^r := \min \left\{ \sum_{v: e \in P_{rv}} a_v, \alpha_0 \right\}, \quad e \in E_T,$$

The strongest form is achieved if r enforces a sort of balance with respect to the cumulated node weights on the paths to r .

Theorem 2. [6] *Let (T, E_T) be a tree in G . The strongest truncated knapsack tree inequality, with respect to the knapsack inequality $\sum_{v \in V} a_v x_v \leq a_0$, defined on (T, E_T) is obtained for a root $r \in \mathcal{R} := \text{Argmin}_{v \in T} \sum_{e \in E_T} \alpha_e^v$. That is, if $r \in \mathcal{R}$ then $\sum_{e \in E_T} \alpha_e^s y_e \geq \sum_{e \in E_T} \alpha_e^r y_e \geq \alpha_0$ holds for all $s \in T$ and all $y \in P_B$. In particular, $\sum_{e \in E_T} \alpha_e^r y_e = \sum_{e \in E_T} \alpha_e^s y_e$ holds for all $r, s \in \mathcal{R}$ and all $y \in P_B$.*

The elements of the set \mathcal{R} are called *minimal roots* of a given tree (T, E_T) , and by Theorem 2 all minimal roots of (T, E_T) deliver the same strongest truncated knapsack tree inequality. Additional structural results allow to locate minimal roots algorithmically at almost no cost. This strengthening proved highly effective in our experiments. Note, if the inequality induces a facet then r is a minimal root by Theorem 2. In some cases the minimal root condition is also sufficient. In order to state this result, call a path in (T, E_T) *branch-less*, if its inner nodes are all of degree 2 in the tree.

Theorem 3. [6] *Assume that $G = (T, E_T)$ is a tree rooted at a node $r \in T$, $f_v = 1$ for all $v \in T$ and $\frac{|T|}{2} + 1 \leq F < |T|$. The truncated knapsack tree inequality $\sum_{e \in E} \min\{|\{v : e \in P_{rv}\}|, |V| - F\} \geq |V| - F$ is facet-defining for P_B if and only if one of the following conditions is satisfied:*

- (a) r is a minimal root and (T, E_T) satisfies the so-called branch-less path condition: each branch-less path with F nodes has one end-edge that is a leaf in (T, E_T) ,
- (b) $F = |T| - 1$.

To motivate a strengthening for general bisection knapsack walk inequalities consider the case of a disconnected graph with two components, one of them being a single edge $\{u, v\}$, the other connected one being $V' = V \setminus \{u, v\}$. If $y_{uv} = 1$ then u and v belong to different clusters and therefore the capacity remaining for the clusters in V' (e.g., the right-hand side of (5)) can be reduced to $F - \min\{f_u, f_v\}y_{uv}$. To generalize this idea we define for $\bar{G} \subseteq G$ with $\bar{V} \subseteq V$, $\bar{E} \subseteq E(\bar{V})$ and $a \in \mathbb{R}_+^{|\bar{V}|}$ a function $\beta_{\bar{G}} : \{0, 1\}^{|\bar{E}|} \rightarrow \mathbb{R} \cup \infty$ with

$$\beta_{\bar{G}}(y) = \inf \left\{ a(S), a(\bar{V} \setminus S) : S \subseteq \bar{V}, \max \{a(S), a(\bar{V} \setminus S)\} \leq a_0, y = \chi^{\delta_{\bar{G}}(S)} \right\}.$$

Now we look at the convex envelope $\check{\beta}_{\bar{G}} : \mathbb{R}^{|\bar{E}|} \rightarrow \mathbb{R} \cup \infty$ of $\beta_{\bar{G}}(y)$, i.e.,

$$\check{\beta}_{\bar{G}}(y) = \sup \left\{ \check{\beta}(y) : \check{\beta} : \mathbb{R}^{|\bar{E}|} \rightarrow \mathbb{R}, \check{\beta} \text{ convex}, \check{\beta}(z) \leq \beta_{\bar{G}}(z) \text{ for } z \in \{0, 1\}^{|\bar{E}|} \right\}.$$

Note that $\check{\beta}_{\bar{G}}$ is a piecewise linear function on its domain.

Proposition 4 (capacity reduced bisection knapsack walk inequality [6]). *Let $\sum_{v \in V} a_v x_v \leq a_0$ with $a_v \geq 0$ for all $v \in V$ be a valid inequality for the knapsack polytope P_K . Let V_0 be a non-empty subset of V and $r \in V_0$. Select subgraphs $(V_l, E_l) = G_l \subset G$ with pairwise disjoint sets V_l , $V_l \cap V_0 = \emptyset$ and $E_l \subseteq E(V_l)$ for $l = 1, \dots, L$. Find for each l an affine minorant for the convex envelope $\check{\beta}_{G_l}$ such that*

$$c_0^l + \sum_{e \in E_l} c_e y_e \leq \check{\beta}_{G_l}(y) \quad (8)$$

holds for all y in P_B . Then the capacity reduced bisection knapsack walk inequality

$$\sum_{v \in V_0} a_v \left(1 - \sum_{e \in P_{rv} \setminus H_v} y_e - \sum_{e \in P_{rv} \cap H_v} (1 - y_e) \right) \leq a_0 - \sum_{l=1}^L (c_0^l + \sum_{e \in E_l} c_e y_e) \quad (9)$$

is valid for P_B .

In certain cases it is possible to establish a full description of $\check{\beta}_{\bar{G}}$ via a complete description of the cluster weight polytope defined as follows. Given a graph $G = (V, E)$ with non-negative node weights $a_v \in \mathbb{R}$ for all $v \in V$. For a set $S \subseteq V$ we define $h(S) := (a(S), (\chi^{\delta(S)})^T)^T \in \mathbb{R}^{|E|+1}$. With respect to a given $a_0 \in \mathbb{R}$ we call

$$P_{CW} = \text{conv} \{ h(S) : S \subseteq V, a(S) \leq a_0, a(V \setminus S) \leq a_0 \}$$

the *cluster weight polytope*.

In [6], a full description of $P_{CW}(\bar{G})$ is given for the special case that the subgraph $\bar{G} = (\bar{V}, \bar{E})$ is a star centered at some node $r \in \bar{V}$, $a \geq 0$, and a_0 satisfies $a(\bar{V}) \leq a_0$. In order to state the nontrivial facets of this case, assume $a(\bar{V} \setminus \{r\}) > a_r$ and call a triple (V_p, \bar{v}, V_n)

feasible if it fulfills $\bar{V} = \{r, \bar{v}\} \dot{\cup} V_p \dot{\cup} V_n$ and $a(V_p) \leq \frac{1}{2}a(\bar{V}) < a(V_p) + a_{\bar{v}}$. For all feasible triples (V_p, \bar{v}, V_n) the inequalities

$$y_0 - \sum_{v \in V_p} a_v y_{rv} - (a(\bar{V}) - 2a(V_p) - a_{\bar{v}}) y_{r\bar{v}} + \sum_{v \in V_n} a_v y_{rv} \geq 0 \quad (10)$$

are the facet-inducing inequalities for $P_{CW}(\bar{G})$. Thus, inequalities (10) form the best linear minorants (8) to be used in (9) in this case.

In our experiments this rather involved strengthening technique proved far less effective than the simple root strengthening for knapsack tree inequalities, but this may be due to the dominating non-negative cost structure in our test instances.

3 A sparse semidefinite cutting-plane approach

We employ the semidefinite cutting-plane approach of [22] which is based on the spectral bundle method [26, 24]. For a brief sketch of the essentials, let \mathcal{S}^n and \mathcal{S}_+^n denote the space of symmetric matrices and symmetric positive semidefinite matrices of order n , let $C = -\frac{1}{4}L$ denote the cost matrix of (2) and collect all constraint matrices of (2) together with additional cutting planes in a linear operator $\mathcal{A} : \mathcal{S}^n \rightarrow \mathbb{R}^m$ with $\mathcal{A}X = [\langle A_i, X \rangle]_{i \in \{1, \dots, m\}}$. It is convenient and computationally useful to require $\|A_i\| := \sqrt{\langle A_i, A_i \rangle} = 1$ for $i \in \{1, \dots, m\}$. In order to keep notation simple we assume that all constraints are given as inequalities with a right-hand side vector $b \in \mathbb{R}^m$, then the negative of (2) reads

$$\max\{\langle C, X \rangle : \mathcal{A}X + \eta = b, X \in \mathcal{S}_+^n, \eta \in \mathbb{R}_+^m\}. \quad (11)$$

The dual eigenvalue formulation.

Because of the constraints $\text{diag}(X) = \mathbf{1}$, the trace of any feasible X is equal to n , *i.e.*, defining

$$\mathcal{W} := \{X \in \mathcal{S}_+^n : \langle I, X \rangle = 1\} = \text{conv}\{vv^T : \|v\| = 1, v \in \mathbb{R}^n\}$$

we have $\{X \in \mathcal{S}_+^n : \mathcal{A}X \leq b\} \subseteq n\mathcal{W}$. Introducing Lagrange multipliers $\mu \in \mathbb{R}^m$, a Slater point or compactness argument shows that there is no duality gap for the Lagrangian dual,

$$\sup_{X \in n\mathcal{W}, \eta \in \mathbb{R}_+^m} \inf_{\mu \in \mathbb{R}^n} [\langle C, X \rangle + (b - \mathcal{A}X - \eta)^T \mu] = \inf_{\mu \in \mathbb{R}^n} \sup_{X \in n\mathcal{W}, \eta \in \mathbb{R}_+^m} [\langle C - \mathcal{A}^T \mu, X \rangle + (b - \eta)^T \mu].$$

The Rayleigh-Ritz characterization of the maximum eigenvalue yields

$$\lambda_{\max}(C - \mathcal{A}^T \mu) = \max_{\|v\|=1} v^T (C - \mathcal{A}^T \mu) v = \max_{\|v\|=1} \langle C - \mathcal{A}^T \mu, vv^T \rangle = \max_{W \in \mathcal{W}} \langle C - \mathcal{A}^T \mu, W \rangle. \quad (12)$$

Therefore the dual is equivalent to the eigenvalue optimization problem

$$\min_{\mu \in \mathbb{R}^m} \varphi(\mu) := n\lambda_{\max}(C - \mathcal{A}^T \mu) + \sup_{\eta \in \mathbb{R}_+^m} (b - \eta)^T \mu.$$

A semidefinite model of λ_{\max} .

The nonsmooth convex function φ can be evaluated at some given $\mu \geq 0$ by computing the maximum eigenvalue λ_{\max} of the symmetric matrix $C - \sum_{i=1}^m A_i \mu_i$. If this is done by an iterative method (the ConicBundle package [23] uses a Lanczos method), structural properties like low rank representations and sparsity can be exploited conveniently. An eigenvector v to λ_{\max} with $\|v\| = 1$ generates a subgradient via $b - \mathcal{A}(nvv^T)$, but any other normalized vector v also generates a linear minorant to the dual function via $W = vv^T$ as in (12). The key idea of the spectral bundle method is to form a manageable minorant of the objective by accumulating (approximate) eigenvector information in a subset of \mathcal{W} of the form

$$\widehat{\mathcal{W}} := \{PUP^T + \alpha\overline{W} : U \in \mathcal{S}_+^r, \alpha \geq 0, \langle I, U \rangle + \alpha = 1\} \subseteq \mathcal{W},$$

where $P \in \mathbb{R}^{n \times r}$ consists of r orthonormal columns ($P^T P = I_r$) and $\overline{W} \in \mathcal{W}$. The columns of P are meant to form a basis of the r most important directions spanned by previous eigenvectors while \overline{W} collects hopefully less important residual contributions; both are updated at every iteration of the bundle method.

The bundle subproblem and the primal aggregate.

In iteration k of the bundle method a next multiplier candidate μ^{k+1} is determined in the vicinity of a current *stability center* $\hat{\mu}^k$ by approximating

$$\operatorname{argmin}_{\mu \in \mathbb{R}^n} \varphi_{\widehat{\mathcal{W}}^k}(\mu) + \frac{1}{2} \|\mu - \hat{\mu}^k\|^2, \quad \text{where} \quad \varphi_{\widehat{\mathcal{W}}^k}(\mu) := \max_{W \in \widehat{\mathcal{W}}^k} n \langle C - \mathcal{A}^T \mu, W \rangle + \sup_{\eta \in \mathbb{R}_+^m} (b - \eta)^T \mu. \quad (13)$$

If the actual progress $\varphi(\hat{\mu}^k) - \varphi(\mu^{k+1})$ is sufficiently large in comparison to the progress predicted by $\varphi(\hat{\mu}^k) - \varphi_{\widehat{\mathcal{W}}^k}(\mu^{k+1})$, a *descent step* is made by setting $\hat{\mu}^{k+1} := \mu^{k+1}$. Otherwise, in a *null step*, the stability center is left unchanged, $\hat{\mu}^{k+1} := \hat{\mu}^k$, but the subgradient information obtained by evaluating $\varphi(\mu^{k+1})$ is used to improve the new model $\widehat{\mathcal{W}}^{k+1}$. For computing the candidate μ^{k+1} the spectral bundle method solves – for some slack vector estimate $\hat{\eta}^k$ – a (small) convex quadratic semidefinite program of the form

$$\begin{aligned} & \text{maximize} && -\frac{1}{2} \|b - \hat{\eta}^k - \mathcal{A}(nW)\|^2 + \langle C - \mathcal{A}^T \hat{\mu}^k, nW \rangle + (b - \hat{\eta}^k)^T \hat{\mu}^k \\ & \text{subject to} && W = P_k U P_k^T + \alpha \overline{W}_k, \\ & && \langle I, U \rangle + \alpha = 1, \\ & && U \in \mathcal{S}_+^r, \alpha \geq 0. \end{aligned} \quad (14)$$

Let $W^k = P^k U^k (P^k)^T + \alpha_k \overline{W}^k$ be the computed optimal solution and v^{k+1} a normalized eigenvector to $\lambda_{\max}(C - \mathcal{A}^T \mu^{k+1})$, then the next P^{k+1} is set to hold a basis of v^{k+1} and the eigenvectors of $P^k U^k (P^k)^T$ belonging to large eigenvalues of U^k (this only requires an eigenvector factorization of the $r \times r$ matrix $U^k = Q^k \Lambda_{U^k} (Q^k)^T$), while \overline{W}^{k+1} collects that part of W^k that is not spanned by P^{k+1} . In order to guarantee convergence, the choice $P^{k+1} = v^{k+1}$ and $\overline{W}^{k+1} = W^k$ would be sufficient, but for purposes explained below we prefer a rather rich P and require $r \leq 20$ only for efficiency reasons. It is important to note that for the spectral bundle algorithm there is no need to store \overline{W}^k , only $\langle C, \overline{W}^k \rangle$ and $\mathcal{A} \overline{W}^k$ are needed. For other purposes, however, it will be useful to store at least part of it, *e.g.*, its elements \overline{W}_{ij} with $ij \in E$ or even some elements on an enlarged support outside E . The

computed solution $\bar{X}^k := nW^k \succeq 0$ of (14) is of central importance to the cutting-plane algorithm. We will refer to \bar{X}^k as the *(primal) aggregate*. The iteration index k is of little relevance in the considerations to follow and will be dropped most of the time.

Separation.

Cluster points of appropriate subsequences of the primal aggregates are optimal solutions of (11), see [21]. Therefore the vector $\bar{y} \in \mathbb{R}^E$ with $\bar{y}_{ij} = \frac{1}{2}(1 - \bar{X}_{ij})$ for $ij \in E$ should eventually be landed in P_B , so our separation routines work with respect to the aggregate \bar{X} . Unfortunately, \bar{X} is only a rough approximation that is typically infeasible and does not satisfy even the current set of constraints $AX \leq b$. One consequence is that \bar{y} may not even satisfy the bound constraints, *i.e.*, in general $\bar{y} \notin [0, 1]^E$ (the bounds hold whenever $\text{diag}(\bar{X}) = \mathbf{1}$). Thus, \bar{y} is truncated or computed with respect to some scaled \bar{X} so as to lie in $[0, 1]^E$ before separation routines are called. A second consequence is that separation routines may find the same inequalities repeatedly and some care has to be taken that the algorithm does not stall due to this property. It was shown in [22] that convergence of appropriate subsequences to optimal solutions can be guaranteed for such a cutting plane approach within the bundle framework if a maximum violation oracle is used for separation, *i.e.*, for a given finite set of inequalities and a given input point y the oracle always returns an inequality whose violation with respect to y is maximal. The result relies on the fact that violation of each individual inequality will be reduced sufficiently over time so that eventually all violated inequalities will be detected. In theory and in practice repeated separation of the same inequalities is typical and requires some additional work in order to eliminate duplicates. After the addition of newly separated cutting planes there is no difficulty in restarting the bundle method from the old Lagrange multiplier solution by setting the new multipliers to zero. Extending the old subgradients to the new coordinates (*i.e.*, setting up the bundle subproblem (14)) can be done easily, if the support of the new inequalities is restricted to the support of \bar{W} . This way there is no need to rebuild the bundle model in spite of the changes in dimension. After every descent step that satisfies the bundle method's termination criterion for a relative precision of 0.2 we delete cuts with sufficiently large slack and call the separation routines for the current \bar{y} , see [4] for a detailed description. Fortunately, no dramatic scaling problems seem to arise after the separation process, maybe because in our current scheme violation of the inequalities and changes in the multipliers seem to converge to zero at a common speed.

Support extension.

Because the semidefiniteness constraint $X \succeq 0$ and the capacity constraint $\langle ff^T, X \rangle \leq (2F - f(V))^2$ couple all entries X_{ij} , the bound can sometimes be improved considerably by enforcing cycle inequalities on entries ij with $ij \notin E$. Within the max-cut setting, the approaches of [25, 35] separate triangle inequalities on the complete graph, but this is computationally too expensive for graphs having more than a few hundred nodes. It was shown in [22] that adding just a few edges per node may already lead to strong improvements while maintaining sparsity. Good guidelines for choosing the additional edges, however, are not yet known. For the unit weight case $f = \mathbf{1}$ it seems likely that there exists some favorable sparse graph structure that works reasonably well when added to any graph. Some experiments in this direction are reported in [4] but so far the dynamic support extension approach of [22] seems to work better. In order to keep the paper self contained we give a brief sketch of it. Given the current

support E^k and the exact values of the aggregate \bar{X} on E^k (this is achieved by storing all values \bar{W}_{ij} for $ij \in E^k$), let $\bar{E}^k := \binom{V}{2} \setminus E^k$ denote the set of edges not in E^k . The idea is to add for each node $\bar{i} \in V$ at most two edges from \bar{E}^k that are hopefully part of a violated cycle inequality (3) in \bar{X} . For edges $\{i, j\} \in \bar{E}^k$ the correct value of \bar{X}_{ij} is not available, but $\tilde{X}_{ij} := n[P^k U^k (P^k)^T]_{ij}$ may be used as a rough approximation. For each root node $\bar{i} \in V$ a shortest path tree is computed within (V, E^k) with respect to edge weights $1 - |\bar{X}_{ij}/\sqrt{\bar{X}_{ii}\bar{X}_{jj}}|$ for $ij \in E^k$. With respect to this tree, each edge $\bar{ij} \in \bar{E}^k$ induces a cycle $\mathcal{C}_{\bar{j}}$. For the cycle inequality (3) on this cycle we use the values $\tilde{y}_{ij} := \frac{1}{2}(1 - \bar{X}_{ij}/\sqrt{\bar{X}_{ii}\bar{X}_{jj}})$ for $ij \in \mathcal{C}_{\bar{j}} \cap E^k$ and $\tilde{y}_{\bar{ij}} = \frac{1}{2}(1 - \tilde{X}_{ij})$ to determine a violation measure

$$v_{\bar{i}}(\bar{j}) := \max\{1 + \sum_{e \in D} (\tilde{y}_e - 1) - \sum_{e \in \mathcal{C}_{\bar{j}} \setminus D} \tilde{y}_e : D \subseteq \mathcal{C}_{\bar{j}}, |D| \text{ odd}\}.$$

For each \bar{i} we now add a “most violated” edge \bar{ij} with $\hat{j} \in \text{Argmax}\{v_{\bar{i}}(j) : \bar{ij} \in \bar{E}^k\}$ and a most undecided edge \bar{ij} with $\tilde{j} \in \text{Argmin}\{|\tilde{X}_{\bar{ij}}| : \bar{ij} \in \bar{E}^k\}$. This support extension routine is called directly after the first separation step and after each tenth separation step. At this point we also eliminate all edges $e \in E^k \setminus E$ that are not in the support of separated inequalities. After every modification of E^k the matrix \bar{W}^k has to be rebuilt from scratch.

Rounding heuristics.

The primal aggregate $X^k = n(P^k U^k (P^k)^T + \alpha_k \bar{W}^k)$ offers several possibilities for designing rounding heuristics. In all approaches we first try to find a partition vector $\bar{x} \in \{-1, 1\}^n$ so that $\bar{x}\bar{x}^T$ is close to X^k and then apply local exchange heuristics in order to enforce feasibility of the bisection constraint $|f^T x| \leq 2F - f(V)$. One way to generate \bar{x} is the random hyperplane rounding technique of Goemans and Williamson [19]. For this, we first compute an eigenvalue factorization of $U^k = Q^k \Lambda_{U^k} (Q^k)^T$ with $\lambda_{\max}(U^k) = [\Lambda_{U^k}]_{11} \geq \dots \geq [\Lambda_{U^k}]_{rr} = \lambda_{\min}(U^k) \geq 0$ and approximate \bar{X} by $\bar{P}\bar{P}^T$ with

$$\bar{P} = nP^k Q^k \sqrt{\Lambda_{U^k}}. \quad (15)$$

Random hyperplane rounding is then applied to the n normalized row vectors of \bar{P} . In fact, for random $h \in \mathbb{R}^k$ we sort the values $\bar{P}_{i,\bullet} h$ ($i \in V$) nonincreasingly and fill up one side of the partition in this order until the capacity restrictions are met or exceeded yielding a corresponding feasible or infeasible \bar{x} . A slightly simplified version of this seemed to work even better. Indeed, the hope is that \bar{X} is actually close to some rank one matrix xx^T with $x \in \{-1, 1\}$ and the sign structure of an eigenvector to the maximum eigenvalue of \bar{X} should already give a reasonable approximation. The first column of \bar{P} should be a reasonable approximation of this eigenvector and we simply try $\bar{x}_i = \text{sign } \bar{P}_{1i}$ as starting partition. If $[\Lambda_{U^k}]_{11}$ is not that much bigger than $[\Lambda_{U^k}]_{22}$ or if there are some values \bar{P}_{1i} close to zero it might make sense to include a bit information from the second eigenvector as well. Therefore we generate two more starting vectors based on $p^+ = \bar{P}_{\bullet,1} + \bar{P}_{\bullet,2}$ and $p^- = \bar{P}_{\bullet,1} - \bar{P}_{\bullet,2}$ using again the sorted values of these vectors to generate initial partitions. This heuristic, called *SumPi* in [4], actually performed best in the experiments of [4]. Some more classical rounding approaches based on the values $\bar{y}_{ij} = \frac{1}{2}(1 - \bar{X}_{ij})$ for $ij \in E$ were also quite successful in [4], but on average SumPi was observed to be slightly superior. In our experiments reported in Section 5 we called SumPi after every descent step.

4 Spectral Bundle Based Branching

In this section we will concentrate on two topics, namely the possibilities offered for finding good branching decisions and the aspect of warmstart after branching. While semidefinite relaxations are quite powerful, they also require a lot of time to compute, so typically any semidefinite branch-and-cut code will only be able to visit very few nodes of the branch-and-cut tree within reasonable time. We use the standard approach of setting one edge variable X_{ij} to either -1 (ij is in the cut) or $+1$ (i and j belong to the same set in the bipartition), but explore new possibilities to exploit the information offered by the primal aggregate \bar{X} . The branching decision is then implemented by adding the constraint $X_{ij} = \pm 1$ and induces a big change on the values of relatively few Lagrange multipliers while many others need almost no corrections. This causes scaling problems and is a notoriously difficult situation for first order methods. Therefore we investigate some ideas to improve warm start properties by scaling heuristics.

Variable selection for branching. First note that there is no need to require $ij \in E$ when determining an X_{ij} to be set to ± 1 in semidefinite relaxations, because ij may always be added to E with a zero cost coefficient. In [25] the branching strategy resulting in the smallest branch-and-cut tree was rule *R3*, it selects an edge $ij \in \text{Argmin}\{|\bar{X}_{ij}| : ij \in E\}$ that is “most undecided”. In our new approaches we would like to improve this criterion by taking into account the global information offered in \bar{P}^k of (15). Using the vector labeling interpretation of [33, 19] each row $\bar{P}_{i,\bullet}$ ($i \in V$) represents node i by a vector in \mathbb{R}^r and ideally they should all lie in a common one-dimensional subspace so that \bar{X} is of rank one. We try to identify a common dominating direction of the vectors $\bar{P}_{i,\bullet}$ by computing a singular value decomposition of $\bar{P} = \bar{U}\Sigma\bar{V}$, with $\Sigma_{11} \geq \dots \geq \Sigma_{rr} \geq 0$. As dominating direction we use $\bar{v} := (\bar{V}_{1,\bullet})^T \in \mathbb{R}^r$. A possible branching rule is to determine the nodes

$$\bar{i} \in \underset{i \in V}{\text{Argmax}} \left| \frac{\bar{P}_{i,\bullet}}{\|\bar{P}_{i,\bullet}\|} \bar{v} \right| \quad \text{and} \quad \bar{j} \in \underset{i \in V}{\text{Argmin}} \left| \frac{\bar{P}_{i,\bullet}}{\|\bar{P}_{i,\bullet}\|} \bar{v} \right|$$

that are most and least articulate with respect to the dominating direction \bar{v} with the intention to clarify the relative position of \bar{j} with respect to the majority of the nodes by setting $X_{\bar{i}\bar{j}}$. This rule is called *Most Orthogonal (MO)* in [4]. In general, one would hope that nodes of almost parallel vectors induce a subgraph that is well decided with respect to a favorable partition of its nodes. Fixing the decision between two pronounced representatives of different one-dimensional subspaces should therefore encourage a decision between all nodes of the corresponding two subgraphs in one step. The branching strategy *Most Orthogonal Elaborate (MOEL)* tries to identify such favorable sets as follows. Given a parameter $\gamma \in (0, 1)$, say $\gamma = 0.2$, the set of nodes considered “orthogonal” to the dominating direction \bar{v} is

$$O := \{i \in V : \left| \frac{\bar{P}_{i,\bullet}}{\|\bar{P}_{i,\bullet}\|} \bar{v} \right| \leq \gamma\}.$$

If O is empty, use rule (MO). Otherwise, form set O_h for $h = 1, 2, \dots, \bar{h}$ and suitable $\bar{h} \in \mathbb{N}$ by finding the next direction

$$\bar{v}_h = \left(\frac{\bar{P}_{i_h,\bullet}}{\|\bar{P}_{i_h,\bullet}\|} \right)^T \quad \text{for some} \quad i_h \in \underset{i \in O \setminus \bigcup_{h'=1}^{h-1} O_{h'}}{\text{Argmin}} \left\{ \left| \frac{\bar{P}_{i,\bullet}}{\|\bar{P}_{i,\bullet}\|} \bar{v} \right| : i \in O \setminus \bigcup_{h'=1}^{h-1} O_{h'} \right\}$$

and by collecting almost parallel vectors to v_h in

$$O_h := \{i \in O \setminus \bigcup_{h'=1}^{h-1} O_{h'} : |\frac{\bar{P}_{i,\bullet}}{\|\bar{P}_{i,\bullet}\|} \bar{v}_h| \geq 1 - \gamma\}$$

so that $O \setminus \bigcup_{h=1}^{\bar{h}} O_h$ becomes empty. As an aside, attempts to identify the directions via singular value decomposition of the rows belonging to O failed so far, maybe due to the rather large diversity of the row vectors belonging to O . Node i_h is used as a representative of the set O_h and it remains to decide which of the edges \bar{v}_h should be branched on. To this end we define for each $h = 1, \dots, \bar{h}$ two trial solutions

$$X_h^+ := P_h^+(P_h^+)^T + \bar{W} \quad \text{and} \quad X_h^- := P_h^-(P_h^-)^T + \bar{W}$$

where

$$[P_h^+]_{i,\bullet} := \begin{cases} \bar{P}_{i,\bullet} & \text{for } i \in V \setminus O_h \\ \text{sign}(\bar{P}_{i,\bullet} \bar{v}_h) \bar{P}_{i,\bullet} & \text{for } i \in O_h \end{cases}$$

and

$$[P_h^-]_{i,\bullet} := \begin{cases} \bar{P}_{i,\bullet} & \text{for } i \in V \setminus O_h \\ -\text{sign}(\bar{P}_{i,\bullet} \bar{v}_h) \bar{P}_{i,\bullet} & \text{for } i \in O_h. \end{cases}$$

They correspond to setting all vectors of nodes in O_h to being positively or negatively aligned with the vector of the representative \bar{i} of the dominating direction \bar{v} . Now we branch on an edge \bar{v}_h with

$$\hat{h} \in \underset{h=1, \dots, \bar{h}}{\text{Argmin}} (\langle C, X_h^+ \rangle + \langle C, X_h^- \rangle)$$

in the hope that this gives the best average improvement of the bound on both branches.

Scaling for warm start after branching. In principle there is no difficulty to start a bundle method from any given stability center $\hat{\mu}$, but subsequent progress of the method depends heavily on the scaling of the variables. If just a few multipliers have to change a lot while most others should roughly stay the same, first order methods typically exhibit very slow convergence. This is exactly the situation that is encountered whenever a branching constraint $X_{ij} = \pm 1$ is added to the converged Lagrangian relaxation of the preceding subproblem. The quadratic term $\frac{1}{2} \|\mu - \hat{\mu}^k\|^2$ in (13) penalizes movement in all directions equally while one would prefer wider steps in directions that are affected by the new constraint. One can try to improve progress by adjusting the quadratic term via diagonal scaling to $\frac{1}{2} \|\mu - \hat{\mu}^k\|_H^2 = \frac{1}{2} (\mu - \hat{\mu}^k)^T H (\mu - \hat{\mu}^k)$ for some appropriate diagonal matrix $H \succ 0$. Several attempts to set up H are described in [4], here we only report on the approach by *spectral analysis at the current point (SACP)* that seemed to perform best among these. To fix ideas, suppose $X_{\bar{i}\bar{j}} = -1$ is enforced by adding the constraint

$$\langle \bar{A}, X \rangle \leq -1 \quad \text{with} \quad \bar{A}_{ij} = \begin{cases} 1, & ij \in \{\bar{i}\bar{j}, \bar{j}\bar{i}\}, \\ 0, & \text{otherwise.} \end{cases}$$

Using Lagrange multiplier $\nu \geq 0$ for this constraint we would like to minimize

$$n\lambda_{\max}(C - \mathcal{A}^T \mu - \nu \bar{A}) + (b - \eta)^T \mu - \nu.$$

The current values are $\mu = \hat{\mu}$ and $\nu = 0$. We would like to increase ν to improve the bound. In order to study its influence on λ_{\max} observe that

$$\bar{A} = \bar{v}\bar{v}^T - \bar{w}\bar{w}^T \quad \text{with} \quad \bar{v}_i = \begin{cases} 1/\sqrt{2}, & i = \bar{i}, \\ 1/\sqrt{2}, & i = \bar{j}, \\ 0, & \text{otherwise,} \end{cases} \quad \bar{w}_i = \begin{cases} 1/\sqrt{2}, & i = \bar{i}, \\ -1/\sqrt{2}, & i = \bar{j}, \\ 0, & \text{otherwise.} \end{cases}$$

Thus increasing ν adds the positive semidefinite part $\nu\bar{w}\bar{w}^T$ to $C - \mathcal{A}^T\hat{\mu}$ while it subtracts a positive semidefinite part $\nu\bar{v}\bar{v}$. If \bar{w} is almost an eigenvector to λ_{\max} , *i.e.*, $\bar{w}^T(C - \mathcal{A}^T\hat{\mu})\bar{w} \geq \lambda_{\max}(C - \mathcal{A}^T\hat{\mu}) - \epsilon$ for some small $\epsilon > 0$, then an increase in ν will increase λ_{\max} by the same amount unless this can be compensated by changing some other variables of $\hat{\mu}$. Among these, only those indices i with rather large values $|\bar{w}^T A_i \bar{w}|$ can be expected to have a strong effect on the subspace spanned by \bar{w} . Because their interaction is hard to predict we allow larger movement for all of them. On the other hand, subtracting $\nu\bar{v}\bar{v}$ can do no harm but might open some space to move for variables μ_i with $|\bar{v}^T A_i \bar{v}|$ large. Motivated by this and putting $A_{m+1} := \frac{1}{\sqrt{2}}\bar{A}$, $\hat{\mu}_{m+1} := \nu$, we define a diagonal scaling matrix H via $H_{ii} := 1/s_i$ for $i = 1, \dots, m+1$ with

$$s_i := 1 + \sigma \begin{cases} \max\{(\bar{v}^T A_i \bar{v})^2, (\bar{w}^T A_i \bar{w})^2\}, & \text{if } \bar{w}^T(C - \mathcal{A}^T\hat{\mu})\bar{w} \geq \lambda_{\max}(C - \mathcal{A}^T\hat{\mu}) - \epsilon, \\ (\bar{v}^T A_i \bar{v})^2, & \text{otherwise.} \end{cases}$$

Based on the experiments in [4] we choose $\sigma = 20$ and $\epsilon = 0.1$ in our actual computations. The scaling is reset to $H = I$ after the next descent step.

5 Computational results at the root node

In this section we first introduce the common software framework and the data sets in Section 5.1. Next, in Section 5.2, we discuss the results of the cutting plane algorithms when separating the cycle, the knapsack tree and the bisection knapsack walk inequalities within the LP and the SDP model. Then, in Section 5.3, we analyze the performance of the heuristics METIS, SumPi and Goemans-Williamson when working with the semidefinite model. Section 5.4 illustrates the effect of separating cycle inequalities on a slightly extended support of the graph. Finally, in Section 5.5, we show the evolution of the lower bound over time for the LP and the SDP model. All computations of this section are carried out only in the root node of the branch-and-bound tree.

Note that in plotting the results in the following figures we commit the sin of linearly connecting the evaluations corresponding to identical algorithmic choices for otherwise independent instances. We do so on purpose because we found it helpful in visually linking the results to the actual parameter settings.

5.1 Computational environment and instances

For solving the minimum graph bisection problem we use the state of the art branch-and-cut framework SCIP (Solving Constraint Integer Programs [1]). In combination with CPLEX as LP-solver SCIP is one of the fastest MIP-solvers currently available. Because SCIP can handle non-linear constraints, it is possible to replace the LP-solver with a semidefinite solver or to use both within one run. SCIP follows the common branch-and-cut strategy which

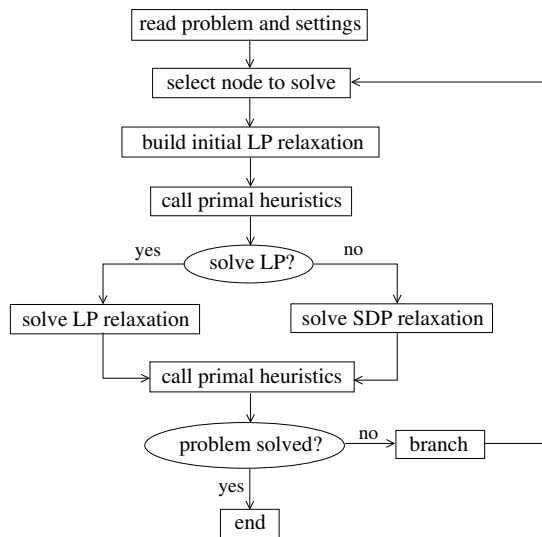


Figure 1: Main solving loop in SCIP.

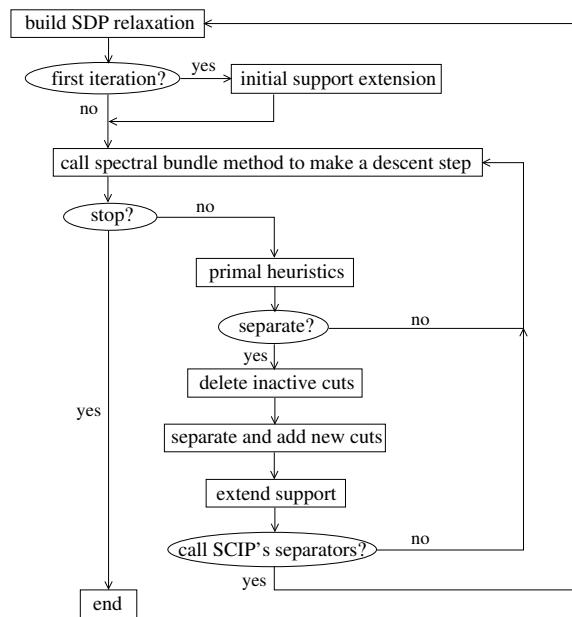


Figure 2: Main solving loop of the SDP-solver.

we outline in Figure 1. For full implementation details we refer to [1]. The semidefinite relaxation is solved by the ConicBundle code [23] which is an implementation of the spectral bundle method [26]. The integration of the SDP-solver into SCIP is explained in detail in [4]. The main solving loop of the SDP-solver is drafted in Figure 2. For our computations we used SCIP version 1.1.0 [2] with ILOG CPLEX 12.10 [27] as LP-solver and Conic Bundle version 0.3.5 as SDP-solver. The computations were executed on Intel(R) Core(TM) i7 CPU 920 machines with 8 MB cache and 12 GB RAM under openSUSE Linux 11.1 (x86_64) in single thread mode.

The first set of graph instances that we used in our computational experiments is based on nets from VLSI design discussed in [29]. The authors consider the placement problem in the layout design of electronic circuits, which consists of finding a non-overlapping assignment of rectangular cells to positions on the chip so that wireability is guaranteed and certain technical constraints are met. Using heuristic algorithms for the optimization problem involved, a partial placement is obtained for the chips *alue*, *alut*, *diw*, *dmxa*, *gap* and *taq*. We use this as input for the minimum graph bisection problem and call the instances *vlsi graphs*. The second set of graphs originates from nested dissection approaches for solving sparse symmetric linear KKT-systems and were communicated to us in [22] by Sharon Filipowski from Boeing. Their names start with *kkt*.

5.2 Separation at the root node

It is important to mention at this point that the SCIP-framework employs the same separation algorithms for both linear (1) and semidefinite (2) programs by transforming any \bar{X} to \bar{y} via $\bar{y}_{ij} = \frac{1-\bar{X}_{ij}}{2}$ for all $ij \in E$. Separated inequalities $\sum_{ij \in E} k_{ij} y_{ij} \leq k_l$ are translated into constraints $\langle K, X \rangle \leq k_s$ for the primal semidefinite relaxation by $K_{ij} = -\frac{1}{2}k_{ij}$ for all $ij \in E$ and $k_s = 2k_l - \sum_{ij \in E} k_{ij}$.

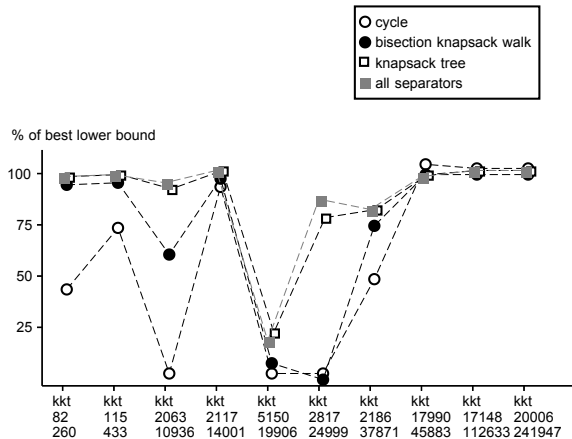


Figure 3: Performance of separation routines with LP-relaxation applied to *kkt*-graphs.

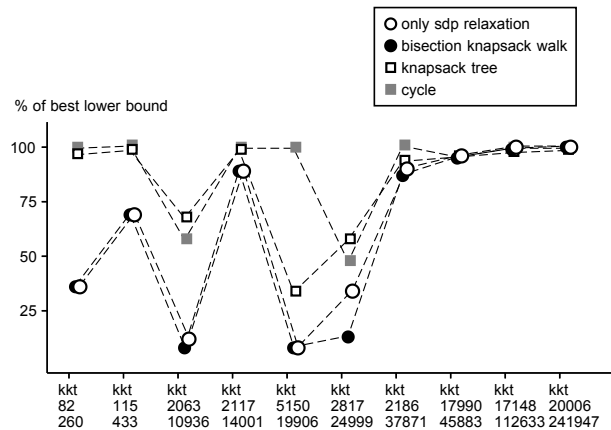


Figure 4: Performance of separation routines with SDP-relaxation applied to *kkt*-graphs.

In Figures 3–6 we present a comparison of the performance of the cutting plane algorithms based on the knapsack tree inequalities with minimal roots (Theorem 2), bisection knapsack walk inequalities (9) and cycle inequalities (3) in combination with the linear and the semidefinite relaxation. Note that the cycle inequalities are part of the integer linear model (1), they are separated in each setting of the LP relaxation. These tests consider the root node without branching and with a time limit of 2 hours. We added inequalities of each class separately as long as violated cuts were found. Within the linear relaxation the knapsack tree inequalities have the biggest impact on the improvement of the lower bound. This may seem surprising in view of the fact that the knapsack walk inequalities subsume the knapsack tree inequalities; the reason is that, for speed, both separators start from a few seed nodes and then the minimal root strengthening of Theorem 2 boosts the performance of knapsack trees. In the LP model the best results are achieved when applying all separators, as one can see in Figures 3 and 5. The cycle separator alone, however, achieves very poor lower bounds, thus studying the bisection cut polytope P_B pays off when considering the linear relaxation of MB. The situation is stunningly different in the semidefinite case. Here, the pure semidefinite relaxation yields already reasonable good lower bounds. For very large instances like *alut2292.494500* the separation routines only slow down the solution process and thus lead to worse bounds when computing time is a major limiting factor. The best results are obtained when the cycle inequalities are separated exclusively (on an extended support, see Section 5.4), additional knowledge about the bisection knapsack polytope does not seem to help for these instances. From the bisection specific inequalities, *i.e.*, the knapsack tree and bisection knapsack walk inequalities, only the knapsack tree inequalities improve the basic SDP bound significantly, while the bisection knapsack walk inequalities seem not to help at all. In comparison to cycle inequalities, however, the bounds are worse and computation times are higher.

5.3 Primal heuristics

A major advantage of semidefinite relaxations over LP relaxations are the powerful possibilities they offer for rounding heuristics. Indeed, while LP rounding methods need a long time to produce reasonable results and have to rely on separation heuristics like MeTiS [30] ini-




Figure 7: The performance of the primal heuristics applied to the semidefinite program.

tially, the random hyperplane rounding approach of Goemans and Williamson [19] produces very good solutions after the first few iterations of the semidefinite solver. In comparison, the slight variation SumPi suggested in Section 3 is sometimes slightly worse but seems to be a bit more stable than pure random hyperplane rounding. This is illustrated in Figure 7 which collects all instances where differences are visible.

5.4 Support extension

The surprising strength of the cycle inequalities in combination with the SDP relaxation observed in Section 5.2 depends crucially on the fact that the inequalities are separated on an extended support. Figure 8 displays the effect of the dynamic support extension technique of [22] sketched in Section 3 in comparison to separating the cycle inequalities on the original support only. Indeed, for some instances without support extension almost no progress is achieved with respect to the pure SDP relaxation because none of the relevant edges are contained in cycles within the original graph.

5.5 Development of the lower bound over time

In Figures 9 - 12 we compare the development of the lower bound of the LP versus the SDP relaxation over time for a few characteristic instances. The results are given for root node computations with the best respective setting for each relaxation. In particular, LP uses all separators and SDP employs only the cycle separator together with dynamic support extension. The time limit is two hours. The plot of Figure 9 gives the typical development of the objective function value over time. The semidefinite relaxation delivers a good bound rather quickly and then exhibits a strong tailing off effect often observed for bundle methods and cutting plane approaches. In contrast, the lower bound of the linear relaxation grows much slower, looking linear on a logarithmic time scale. Figures 10 and 11 give examples of some bigger instances where progress is very slow for SDP, LP performance being abysmal




Figure 8: Semidefinite relaxation in the root node with and without support extension.

in these cases. An exception is the instance *kkt2817_24999* of Figure 12, where the linear relaxation outperforms the semidefinite one, but the shape of both curves is again typical.




Figure 9: dmx1755.3686




Figure 10: alut2292.494500

6 Computational results for branch-and-cut

The second part of our computational experiments focuses on our SCIP based semidefinite branch-and-cut approach sketched in Section 5.1. Motivated by the results of the previous section the parameters were set as follows. For the LP-relaxation the knapsack tree separator is given the highest priority and separation frequency, followed by the cycle and the bisection knapsack walk separators. The LP branching decision is determined by the standard setting provided by SCIP. For the semidefinite relaxation the cycle separator is the only separator. The SDP branching strategy is derived from the experiments presented in this section. We limited computation time to 4 hours for each instance. In Section 6.1 we show the performance




Figure 11: kkt5150.19906

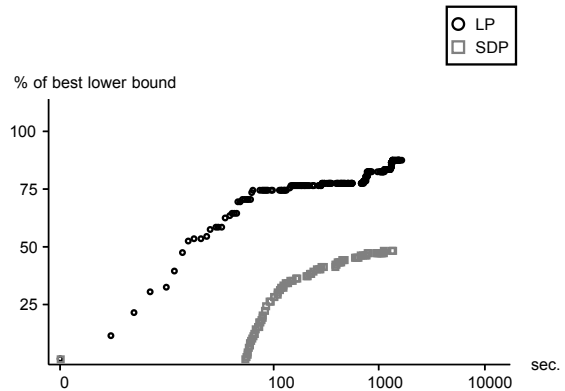


Figure 12: kkt2817.24999

of the four branching rules implemented for SDP. The tests presented in Section 6.2 concern the decision when to start to branch in order to reduce tailing off effects. In Section 6.3 we show the impact of the scaling and warm start methods. Finally, in Section 6.4, we compare the performance of the LP and SDP based branch-and-cut algorithm.

Before embarking on the actual comparisons please note that all results of this section need to be interpreted with some care. Indeed, a general difficulty in judging the effectiveness of branching rules for SDP is the fact that due to rather long computation times in the root node the actual number of branching decisions observed within a computation time of 4 hours is quite limited, in particular if the emphasis is on medium to large instances. An extensive study on smaller and in part random instances is given in [4] and results there point into the same directions observed here.

6.1 Branching rules for SDP

Figure 13 displays the performance of the branching strategies *most orthogonal* (MO) and *most orthogonal elaborate* (MOEL) suggested in Section 4 in comparison to the standard rules *most infeasible*, which picks the edge ij within the original support E that minimizes $|X_{ij}|$, and *random infeasible*, which picks a random edge $ij \in E$ among all edges with $|X_{ij}| \neq 1$. Results are only given for instances with at least three branch-and-bound nodes. While MOEL is not always the best choice, it produces good results quite reliably and so this is selected as the default rule in all further experiments.

Note that all newer branching rules (pseudocost branching, reliable branching, etc., see for instance [3]) for LP based branch-and-cut have also been taken into account with limited success. The reason is the same as mentioned above. These ideas are especially successful if the branch-and-bound trees are huge, because they gather information throughout the enumeration process about the variables to branch on. Since the SDP trees are much smaller these techniques tend to be much less effective.

6.2 Early versus late branching

As illustrated in Section 5.5, the joint tailing off effects of the spectral bundle method and the cutting plane approach result in very slow progress after strong initial improvements of the bound. Once this slow down is observed, the hope is to obtain faster progress by branching




Figure 13: The performance of the branching rules.

and computing bounds for the subproblems of each branch. In lack of a mathematically well founded indicator for identifying a good point for switching from the separation process to branching, we experimented with a decision rule based on the inner mechanics of the bundle method. For bundle methods, the number of functions evaluations that do not lead to descent steps is a good indicator for the progress of the method. In Figure 14 we consider three scenarios. In all of them we stop computation in a node for the usual combination of criteria regarding the progress of the bound, the violation of inequalities and the relative precision obtained (see [4] for details), giving the basic setting of *late branching*. In *early branching* we stop even before these criteria are met if the number of descent steps is at least ten and the total number of evaluations exceeds twenty times the number of descent steps. In *less early branching* the latter factor is set to 200. For small to medium size instances the *early branching* strategy seems to work best on the instances tested. Because branching is not really attractive for large scale instances we choose *early branching* as our default strategy in the sequel.

6.3 Warmstart and scaling

As discussed in Section 4, any nonnegative dual point may be chosen as starting point for the spectral bundle method, but because of the importance of scaling issues for first order methods it would be desirable to start at points, where all dual variables have to change by roughly the same amount. When considering the initialization of a new branch-and-cut node, it is therefore worth to investigate the effect of *warmstarts* the method from the best dual point computed in its parent node. In Section 4 we also introduced a scaling heuristic intended to mitigate the effects of scaling difficulties caused by the addition of a branching constraint $X_{ij} = \pm 1$. The impact of warmstarting and scaling options is displayed in Figure 15. On these instances, warmstarting has a positive effect while the influence of




Figure 14: Early versus late branching.

the scaling heuristic is somewhat ambiguous. Additional experiences collected in [4] indicate some advantage for scaling on small to medium scale instances and so we decided in favor of warmstarting in combination with the scaling heuristic as our default strategy. It should be noted, however, that we do not consider the current approach satisfactory and more needs to be done. Yet, even in the case of significant progress in warmstarting it seems unlikely that the same efficiencies can be reached that are observed in restarting the simplex method within linear branch-and-cut methods.

6.4 LP versus SDP

Figure 16 displays the final results of the branch-and-cut codes for LP and SDP relaxations for the default settings obtained within 4 hours of computation time for each instance and compares these to the lower bounds achieved by the SDP-relaxation in the root within 2 hours of computation time as described in Section 5.5. The advantage of the SDP over the LP approach matches the performance of the bounds as discussed in Section 5.5. The rather ambiguous outcome of the comparison between SDP branch-and-cut bounds and the SDP root node bound may be seen as a disappointment at first, but is actually in line with the previous observations and folklore experience in branch-and-cut. For small to medium size instances branch-and-cut may actually help to speed up the bounding process, and the results indicate that our approach is a reliable choice for instances with up to 15000 edges and a few thousands of nodes. For larger problems, the time limit is best exploited by putting all efforts into the root node and branching is a waste of time.

7 Conclusion

While SDP relaxations are widely believed to be of use only for small dense instances, our work demonstrates by direct comparison within the same branch-and-cut framework that they are also an attractive and often superior choice in the classical linear programming domain




Figure 15: The impact of warmstart and scaling.

of large and sparse graph partitioning instances. For bisection instances, the semidefinite cutting plane approach based on the spectral bundle method is certainly competitive to linear programming cutting plane methods. Problem specific cutting planes studied in [6], in particular knapsack tree inequalities centered at minimal roots, boost the quality of the LP bound but do not seem to improve the performance of the SDP bounds. Indeed, it is one of the major surprises that for the bisection instances considered here the best results were obtained for the basic semidefinite relaxation improved by exclusively separating cycle inequalities on a slightly extended support. This raises hopes for obtaining semidefinite bounds of similar quality for general constrained quadratic 0-1 programming problems without using problem specific cutting planes. The proposed extension of the semidefinite cutting plane method to a semidefinite branch-and-cut approach works well on instances with up to 15000 edges and a few thousand nodes. For larger instances more work is needed to increase the efficiency of warmstarting the spectral bundle method after the addition of branching constraints.

In order to improve the results it would be valuable to better understand structural properties that are required in the selection of a slightly extended support that is favorable for the separation of general valid inequalities of the cut polytope. This might also give further insight on why the current problem specific inequalities show so little effect on the SDP bound. On the algorithmic side, any progress on the scaling difficulties caused by new primal cuts within the spectral bundle method should be helpful for the semidefinite cutting plane method as well as for the semidefinite branch-and-cut approach.

Acknowledgements. This work was supported by grants HE 3524/1-2 and MA1324/1-2 of Deutsche Forschungsgemeinschaft (DFG) and mostly carried out while Michael Armbruster was at Chemnitz University of Technology and Marzena Fügenschuh and Alexander Martin were at Darmstadt University of Technology.




Figure 16: Branch-and-cut based on the linear versus the semidefinite relaxation. (early branch)

References

- [1] T. Achterberg. *Constraint Integer Programming*. PhD-thesis, Technische Universität Berlin, München, 2008.
- [2] T. Achterberg. *SCIP 1.1.0*. Konrad-Zuse-Zentrum für Informationstechnik Berlin, Takustr. 7, D-14195 Berlin-Dahlem, Germany, 2009. URL: <http://scip.zib.de>.
- [3] T. Achterberg, T. Koch, and A. Martin. Branching rules revisited. *Operations Research Letters*, 33:42 – 54, 2005.
- [4] M. Armbruster. *Branch-and-Cut for a Semidefinite Relaxation of Large-scale Minimum Bisection Problems*. PhD-Thesis, Technische Universität Chemnitz, Chemnitz, Germany, 2007.
- [5] M. Armbruster, M. Fügenschuh, C. Helmberg, and A. Martin. A comparative study of linear and semidefinite branch-and-cut methods for solving the minimum graph bisection problem. In A. Lodi, A. Panconesi, and G. Rinaldi, editors, *Integer Programming and Combinatorial Optimization*, volume 5035 of *Lecture Notes in Computer Science*, pages 112–124, Berlin, 2008. IPCO 2008, Bertinoro, Italy, Springer.
- [6] M. Armbruster, M. Fügenschuh, C. Helmberg, and A. Martin. On the graph bisection cut polytope. *SIAM J. Discrete Math.*, 22(3):1073–1098, 2008.
- [7] F. Barahona and A. R. Mahjoub. On the cut polytope. *Math. Programming*, 36:157–173, 1986.
- [8] M. Conforti, M. Rao, and A. Sassano. The equipartition polytope I. *Math. Programming*, 49:49–70, 1990.

- [9] M. Conforti, M. Rao, and A. Sassano. The equipartition polytope II. *Math. Programming*, 49:71–90, 1990.
- [10] C. C. de Souza. *The graph equipartition problem: Optimal solutions, extensions and applications*. PhD-Thesis, Université Catholique de Louvain, Louvain-la-Neuve, Belgium, 1993.
- [11] M. Deza and M. Laurent. *Geometry of Cuts and Metrics*, volume 15 of *Algorithms and Combinatorics*. Springer, 1997.
- [12] A. Eisenblätter. *Frequency Assignment in GSM Networks*. PhD-Thesis, Technische Universität Berlin, Berlin, 2001. ISBN 3-89873-213-4, <ftp://ftp.zib.de/pub/zib-publications/books/PhD.eisenblaetter.ps.Z>.
- [13] C. E. Ferreira, A. Martin, C. C. de Souza, R. Weismantel, and L. A. Wolsey. Formulations and valid inequalities for the node capacitated graph partitioning problem. *Math. Programming*, 74:247–267, 1996.
- [14] C. E. Ferreira, A. Martin, C. C. de Souza, R. Weismantel, and L. A. Wolsey. The node capacitated graph partitioning problem: a computational study. *Math. Programming*, 81(2):229–256, 1998.
- [15] A. Frieze and M. Jerrum. Improved approximation algorithms for MAX k -CUT and MAX BISECTION. *Algorithmica*, 18(1):67–81, 1997.
- [16] M. Fügenschuh. *Relaxations and Solutions for the Minimum Graph Bisection Problem*. Phd-thesis, Darmstadt University of Technology, Darmstadt, Germany, 2007.
- [17] M. R. Garey and D. S. Johnson. *Computers and Intractability*. W. H. Freeman and Company, New York, 1979.
- [18] J. R. Gilbert and R. E. Tarjan. The analysis of a nested dissection algorithm. *Numer. Math.*, 50:377–404, 1987.
- [19] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42:1115–1145, 1995.
- [20] C. Helmberg. Fixing variables in semidefinite relaxations. *SIAM J. Matrix Anal. Appl.*, 21(3):952–969, 2000.
- [21] C. Helmberg. Semidefinite programming for combinatorial optimization. Habilitationsschrift TU Berlin, Jan. 2000; ZIB-Report ZR 00-34, Konrad-Zuse-Zentrum für Informationstechnik Berlin, Takustraße 7, 14195 Berlin, Germany, Oct. 2000.
- [22] C. Helmberg. A cutting plane algorithm for large scale semidefinite relaxations. In M. Grötschel, editor, *The Sharpest Cut*, MPS-SIAM Series on Optimization, pages 233–256. SIAM/MPS, 2004.
- [23] C. Helmberg. *ConicBundle 0.3*. Fakultät für Mathematik, Technische Universität Chemnitz, 2009. <http://www.tu-chemnitz.de/~helmberg/ConicBundle>.

- [24] C. Helmberg and K. C. Kiwiel. A spectral bundle method with bounds. *Math. Programming*, 93(2):173–194, 2002.
- [25] C. Helmberg and F. Rendl. Solving quadratic (0,1)-problems by semidefinite programs and cutting planes. *Math. Programming*, 82(3):291–315, Aug. 1998.
- [26] C. Helmberg and F. Rendl. A spectral bundle method for semidefinite programming. *SIAM J. Optim.*, 10(3):673–696, 2000.
- [27] ILOG S.A., 9 Rue de Verdun, 94253 Gentilly Cedex, France. *ILOG AMPL CPLEX System, Version 12.1, User’s Guide*, 2008. Information available at <http://www.ilog.com>.
- [28] E. Johnson, A. Mehrotra, and G. Nemhauser. Min-cut clustering. *Math. Programming*, 62:133–152, 1993.
- [29] M. Jünger, A. Martin, G. Reinelt, and R. Weismantel. Quadratic 0/1 optimization and a decomposition approach for the placement of electronic circuits. *Math. Programming*, 63(3):257–279, 1994.
- [30] G. Karypis and V. Kumar. MeTiS: A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices, version 4.0. Technical report, Department of Computer Science, University of Minnesota, 1998. URL (13.03.2007): <http://glaros.dtc.umn.edu/gkhome/views/metis>.
- [31] M. Laurent and C. C. de Souza. Some new classes of facets for the equicut polytope. *Discrete Applied Mathematics*, 62:167–191, 1995.
- [32] T. Lengauer. *Combinatorial Algorithms for Integrated Circuit Layout*. John Wiley and Sons Ltd, Chichester, 1990.
- [33] L. Lovász. On the Shannon capacity of a graph. *IEEE Transactions on Information Theory*, IT-25(1):1–7, Jan. 1979.
- [34] S. Poljak and F. Rendl. Nonpolyhedral relaxations of graph-bisection problems. *SIAM J. Optim.*, 5(3):467–487, 1995.
- [35] F. Rendl, G. Rinaldi, and A. Wiegele. Solving max-cut to optimality by intersecting semidefinite and polyhedral relaxations. Technical report, Alpen-Adria-Universität Klagenfurt, 2007.
- [36] R. Weismantel. On the 0/1 knapsack polytope. *Math. Programming*, 77:49–68, 1997.