

The Bundle Method in Combinatorial Optimization

Ilse Fischer* Gerald Gruber† Franz Rendl* Renata Sotirov*‡

September 17, 2003

Key words: Bundle Method, Semidefinite Programming, Max-Cut.

AMS Subject Classification: 90C22, 90C27.

Abstract

We propose a dynamic version of the bundle method to get approximate solutions to semidefinite programs with a nearly arbitrary number of linear inequalities. Our approach is based on Lagrangian duality, where the inequalities are dualized, and only a basic set of constraints is maintained explicitly. This leads to function evaluations requiring to solve a relatively simple semidefinite program.

Our approach provides accurate solutions to semidefinite relaxations of the Max-Cut and the Equipartition problem, which are not achievable by direct approaches based only on interior-point methods.

1 Introduction

Many combinatorial optimization problems have a natural formulation in 0-1 variables, leading to either linear or quadratic problems in binary variables. To get tractable relaxations, one can use a (partial) description of the convex hull of integer solutions (polyhedral approach), or more recently, semidefinite programming.

After having formulated some basic semidefinite relaxation, it is usually possible to tighten it by adding additional linear inequalities valid for all 0-1 solutions. These combinatorially derived 'cutting planes' pose a serious challenge even to state of the art software for semidefinite programming, because the potential number of these constraints can be prohibitively large.

The most reliable methods to solve semidefinite programs are based on the interior-point idea, which means applying the Newton method to some variant of parameterized optimality conditions. The primal-dual path-following variants seem to yield the most reliable computational results, see [27].

The computational effort for one interior-point iteration is determined by the following factors:

- Maintaining positive definite matrices of some given order n .
- Linear algebra operations with matrices of order n , such as matrix multiplication or forming the inverse explicitly.

*Universität Klagenfurt, Institut für Mathematik, Universitätsstraße 65-67, A-9020 Klagenfurt, Austria, (`{ifischer,rendl,rsotirov}@uni-klu.ac.at`)

†Carinthia Tech Institute, Department of Geoinformation, Europastrasse 4, A-9524 Villach, Austria, (`g.gruber@cti.ac.at`)

‡Support from the Austrian Science Foundation FWF Project P12660-MAT is gratefully acknowledged.

- Solving a linear system of order given by the number m of constraints of the problem.

The first two factors can be bounded by $O(n^3)$ operations, independent of the number of constraints. The last factor depends heavily on the number m of constraints, and is at least $O(m^3)$ plus the effort to determine the matrix of the linear system to be solved. The main conclusion is therefore that the number m of constraints has to be limited, if interior-point methods are used to solve the relaxation.

In this paper we investigate alternative algorithmic approaches, where we still assume that the order n of the matrix space is reasonably small ($n \leq 1000$), but the number of constraints can be significantly larger.

Recently, Barahona and Anhil [3] and Barahona and Ladanyi [4] revived the idea of working with the Lagrangian dual to handle constraints $Ax \leq b$ only indirectly. The Lagrangian dual functional is nonsmooth, so one can use the algorithmic machinery from nonsmooth optimization to deal with it. In [3, 4] a simple version of the subgradient iteration scheme of Polyak [23] is refined to carry out the iterations. The key point with these methods is that they rely on a subroutine, which carries out a function and subgradient evaluation of the dual functional. Using the subgradient information, a new trial point is determined, and the process is iterated.

Contrary to the setup in [3, 4], we have a more general framework in mind, where the function evaluation itself may be nontrivial. Therefore it is important to provide an algorithmic tool, which works well even if the number of function evaluations is limited (to a few dozen function evaluations). We will use the bundle method instead of the subgradient scheme, to make efficient use of subgradient information, gathered during the computational process. It was recently shown [2] that the approach from [3, 4] can in fact be viewed as a simple variant of the bundle method.

It is the main purpose of this paper to show that a dynamic version of the bundle method provides an efficient machinery to approximately solve semidefinite problems with a nearly arbitrary number of inequalities. More precisely, we will concentrate on the Max-Cut and the Equipartition problem. For both these problems there exists a rather 'simple' semidefinite relaxation which can be refined by including additional cutting planes ($O(n^3)$ triangle inequalities for Max-Cut or $O(n^2)$ sign constraints for Equipartition). We demonstrate that these problems can be solved rather accurately for problems of medium size (number of vertices of the graphs no more than 500). To the best of our knowledge, such results are currently not available using other methods. We explain the main features of our version of the bundle method, which consists of repeatedly identifying new interesting constraints and removing old unimportant constraints. We also describe other types of problems where we think that the present approach might work well. We finally also show the practical limits of our approach.

Notation: We use standard notation for semidefinite programming. $\text{tr}(A)$ denotes the trace of the square matrix A . The vector of all ones is e .

2 Large Semidefinite programs from combinatorial optimization

We develop our ideas on the following two graph optimization problems. Both are defined on an undirected edge-weighted graph G on vertex set $V := \{1, \dots, n\}$. The relevant data are conveniently expressed through the weighted adjacency matrix $A = (a_{ij})$, where

$$a_{ij} = a_{ji} = w_e \text{ for edge } e = [ij] \in E,$$

$$a_{ij} = 0 \text{ if } [ij] \notin E.$$

Given A , we also introduce the Laplacian matrix $L = L_A$, associated to A , which is defined as

$$l_{ii} = \sum_k a_{ik}, \quad \forall i, \quad l_{ij} = -a_{ij} \quad i \neq j.$$

The **Max-Cut Problem (MC)** asks to partition V into S and $V \setminus S$ in such a way that the total weight of the edges joining S and $V \setminus S$ is maximized. (If $S = \emptyset$ or $S = V$, the weight of the cut is defined to be 0, since no edges are cut in this case.)

Let us encode $S \subseteq V$ by setting $x_i = 1 \quad \forall i \in S$ and $x_i = -1 \quad \forall i \notin S$. It is a well-known fact, easy to verify, that the weight of the cut, given by $S \subseteq V$, $cut(S, V \setminus S)$ is given by

$$cut(S, V \setminus S) = \frac{1}{4} x^T L x.$$

Hence Max-Cut is equivalent to

$$z_{mc} := \max\{x^T L x : x \in \{-1, 1\}^n\}. \quad (1)$$

In case of the **k -Equipartition Problem (k -EP)** we make the assumption that the number n of vertices is a multiple of a given (integer) parameter k , $n = mk$. The problem now consists in partitioning V into k sets of cardinality m , so as to minimize the total weight of edges joining different partition blocks. Now we represent k -partitions S_1, \dots, S_k by 0 – 1 matrices X of size $n \times k$ with $x_{ij} = 1$ if vertex $i \in S_j$, and $x_{ij} = 0$ otherwise. The problem now amounts to the following, see e.g. [6].

$$z_{ep} = \min\{\text{tr} X^T L X : X e = e, \quad X^T e = \frac{n}{k} e, \quad x_{ij} \in \{0, 1\}\}. \quad (2)$$

2.1 Semidefinite Relaxations

The formulation (1) of Max-Cut gives rise to the following relaxation, which has attracted a lot of attention due to the famous approximation analysis of Goemans and Williamson [8]. It is most easily derived through the trivial identity

$$x^T L x = \text{tr} L (x x^T).$$

We replace $x x^T$ by a new matrix variable X and observe that $x \in \{-1, 1\}^n$ implies in this case that $X \succeq 0$ and $\text{diag}(X) = e$. Hence we obtain the following basic relaxation of (1):

$$z_{mc\text{-basic}} = \max\{\text{tr} L X : \text{diag}(X) = e, \quad X \succeq 0\}. \quad (3)$$

This is a semidefinite program in the matrix variable X of order n , and n equality constraints. This relaxation was introduced by Delorme and Poljak [5] in a dual form. The primal version (3) can be found in [22]. In [8] it is shown that this relaxation has an error of no more than 13.82 %, i.e.

$$\frac{z_{mc\text{-basic}}}{z_{mc}} \leq 1.1382,$$

if all $a_{ij} \geq 0$. Nesterov [21] provides a weaker bound for a larger class of instances with $L \succeq 0$ by showing that

$$\frac{z_{mc\text{-basic}}}{z_{mc}} \leq 1.57.$$

It is not too difficult to see that the relaxation (3) can be strengthened by requiring X to satisfy the triangle inequalities, $X \in \text{MET}$, see [22, 12]. By definition, $X \in \text{MET}$ if and only if

$$\begin{aligned} x_{ij} + x_{ik} + x_{jk} &\geq -1, \\ x_{ij} - x_{ik} - x_{jk} &\geq -1 \\ -x_{ij} + x_{ik} - x_{jk} &\geq -1, \\ -x_{ij} - x_{ik} + x_{jk} &\geq -1 \quad \forall i < j < k. \end{aligned}$$

Hence we get

$$z_{mc-met} = \max\{\text{tr } LX : \text{diag}(X) = e, X \in \text{MET}, X \succeq 0\}. \quad (4)$$

This is again a semidefinite program, but it has $4\binom{n}{3}$ triangle inequalities in addition to the n equations fixing the main diagonal of X to e . The computational effort to solve this problem is nontrivial, even for small n like $n \approx 100$, see e.g. [11].

Turning to k -equipartition, we find a similar situation. Linearizing again using $\text{tr} X^T L X = \text{tr} L (X X^T)$ and introducing a new variable Y for $X X^T$, we observe, see e.g. [1, 13] that Y has to satisfy the constraints $Y \succeq 0$, $\text{diag}(Y) = e$, $Y e = \frac{n}{k} e$.

This gives rise to the following basic semidefinite relaxation for k -EP, introduced in [13].

$$z_{ep-basic} := \min\{\text{tr } LY : Y e = \frac{n}{k} e, \text{diag}(Y) = e, Y \succeq 0\}. \quad (5)$$

In [13] it is observed that the condition $Y e = \frac{n}{k} e$ can be eliminated by parameterizing Y as $Y = \frac{1}{k} e e^t + V R V^T$. Here V is an arbitrary $n \times (n-1)$ matrix with its columns spanning the orthogonal complement of e , and R is symmetric of order $n-1$. This representation insures that Y has the eigenvector e to the eigenvalue $\frac{n}{k}$ and all the other eigenvectors are orthogonal to e . Since $L e = 0$ by the definition of the Laplacian L , we obtain the following equivalent formulation, see [13]:

$$z_{ep-basic} = \min\{\text{tr } (V^T L V) R : \text{diag}(V R V^T) = \frac{k-1}{k} e, R \succeq 0\}.$$

This basic relaxation is again a problem in the matrix variable R of order $n-1$ with n equations.

In [13], several refinements of this model are discussed. The simplest one includes the additional constraint $Y \geq 0$ elementwise.

$$z_{ep-nonneg} := \min\{\text{tr } LY : Y e = \frac{n}{k} e, \text{diag}(Y) = e, Y \geq 0, Y \succeq 0\}. \quad (6)$$

Including all $\binom{n}{2}$ sign constraints $y_{ij} \geq 0$ directly is computationally prohibitive for problems of medium size, like $n \approx 200$.

In both cases we have the situation that the basic semidefinite relaxation can be solved with reasonable effort for problem sizes n rather large. A Matlab routine to compute $z_{mc-basic}$ is for instance contained in [12]. (It is also available on the following web-site. <http://www-math.uni-klu.ac.at/or/software/>) Typical computation times with this routine on randomly generated problems of various sizes are given in Table 1.

A Matlab routine that computes $z_{ep-basic}$ is given in [13]. The computation times are similar to Table 1. The situation is different, once we include additional constraints such as the triangle inequalities for Max-Cut or the sign constraints for k -EP. We conclude this section with a discussion of the limitations of standard interior-point methods to deal with semidefinite relaxations like (4) or (6).

n	seconds
200	1.84
400	8.92
600	24.10
800	51.45
1000	99.27
1500	314.99
2000	714.21

Table 1: Computation times (seconds) to solve the SDP relaxation (3) on a PC (Pentium 4, 1.7 Ghz).

	$n = 100$	$n = 200$	$n = 300$
$m = 500$	21 (19)	34 (20)	49 (19)
$m = 1000$	103 (21)	136 (22)	164 (21)
$m = 1500$	304 (24)	358 (24)	422 (24)
$m = 2000$	643 (25)	763 (26)	816 (24)
$m = 2500$	1090 (24)	1313 (26)	1360 (24)

Table 2: Computation times (seconds) on a PC (Pentium 4, 1.7 GHz) to compute the semidefinite relaxation of Max-Cut for a graph with n nodes and m triangle inequalities. The number of interior point iterations is given in parentheses.

2.2 Limitations of Interior-Point methods

Let us consider solving (4) directly by interior point methods, as proposed e.g. in [9, 11]. The total number of inequalities in (4) is roughly $\frac{2}{3}n^3$, hence impossible to include directly for problems of interest ($n \geq 100$). On the other hand, there are only $\binom{n}{2}$ variables in (4), so the 'right' choice of active inequalities should not contain more than $\binom{n}{2}$ of the triangle inequalities. Unfortunately, this number is still too big. To illustrate how the number of triangle inequalities influences the computation times, we provide in Table 2 computation times to solve (4) with $n \in \{100, 200, 300\}$ vertices and $m \in \{500, 1000, 1500, 2000, 2500\}$ triangle constraints by interior point methods. The computation times clearly indicate that problems with significantly more than 2000 triangle constraints are computationally expensive to solve.

Another reason for potential difficulties with a direct approach lies in the problem of quickly identifying important inequalities, that is to say those which are active at the optimum. In [9] it was proposed to take the constraints most violated by the optimal solution X of (3), and solve the relaxation (4) with those constraints included, and iterate this process. This idea should work well, if a reasonable proportion of constraints violated by X would appear in the list of constraints active at the optimum. Unfortunately, our computational experiments do not indicate that this is likely to happen.

To elaborate on this point, we report results about the following experiment. We first solved the relaxation (4) for the problems g1d, g1s, spin5 and g2d given in the appendix.

Then we took the optimal solution X_{sdp} of the basic relaxation (3) and sorted all the triangle constraints according to violation with respect to X_{sdp} . Thus the first constraint is the most

graph	n	active for (4)	500	1000	2000	5000	inactive
g1d	100	1053	66	120	205	406	20
g1s	100	1371	33	63	118	253	199
spin5	125	1924	21	29	51	116	420
g2d	200	2108	56	94	158	302	24

Table 3: The active constraints of (4) are difficult to identify by looking at the most violated constraints of the basic relaxation. The last column gives the number of active constraints of (4), which are inactive at the basic relaxation.

violated one, and so on. In Table 3 we provide the number of active constraints for these problems at the optimal solution of (4) in column 3. We first observe that this number is much smaller than $\binom{n}{2}$.

We then checked, how many of the most violated constraints of (3) are among the active constraints at the optimal solution to (4). We looked separately at the 500, 1000, 2000 and 5000 most violated constraints (columns 4,5,6 and 7). Thus we see that among the 500 most violated constraints of X_{sdp} for g1d, only 66 are among the final active constraints. The situation is even worse for spin5, where among the 5000 most violated constraints of (3), only 116 are active at (4).

Finally, we also checked, how many of the active constraints of (4) are inactive at the optimal solution of (3) (column 8). Ideally, this number should be quite small, but unfortunately it is not in all our examples.

From these results it should not come as a surprise that solving (4) iteratively by including the most violated constraints of (3) takes a nontrivial computational effort.

As a possible alternative we therefore follow the philosophy of Lagrangian duality and try to approximate the Lagrangian dual using techniques from non-smooth optimization. We recall the necessary mathematical machinery and also the relevant technicalities in the following section.

3 Lagrangian Duality and the Bundle Method

3.1 Lagrangian Duality

To describe the semidefinite problems from before in a general setting, we first reformulate them as follows. It will be irrelevant that our problems are formulated in the space of symmetric matrices, hence we prefer expressing them simply as finite-dimensional (nonlinear) optimization problems. Using $x = \text{vec}(X)$, $c = \text{vec}(L)$, we introduce the set $\mathcal{X} := \{x : x = \text{vec}(X), \text{diag}(X) = e, X \succeq 0\}$. In this case, (3) can equivalently be written as

$$\max\{c^T x : x \in \mathcal{X}\}. \tag{7}$$

A similar definition would give the basic relaxation of k -EP. Including the triangle constraints in (7) amounts to including the constraints

$$Ax \leq b$$

where A is a matrix of order $4\binom{n}{3} \times n^2$, with each row of A representing one of the triangle inequalities. Thus we end up with a problem of the following general form:

$$z = \max\{c^T x : x \in \mathcal{X}, Ax \leq b\}. \tag{8}$$

We assume that the problem without the inequalities is 'easy' to solve, but the inclusion of $Ax \leq b$ makes it computationally expensive, see the discussion in the previous section.

We introduce the Lagrangian

$$L(x, \gamma) := c^T x + \gamma^T (b - Ax) \quad (9)$$

and the dual functional

$$f(\gamma) := \max_{x \in \mathcal{X}} L(x, \gamma) = b^T \gamma + \max_{x \in \mathcal{X}} (c^T - A^T \gamma)x. \quad (10)$$

The following notation will be useful. We call a pair (γ, x) a **matching pair** for f , if $f(\gamma) = L(x, \gamma)$. The Lagrangian dual now is $z = \min_{\gamma \geq 0} f(\gamma)$. Note that we assume \mathcal{X} to be nice in the sense that $\max_{x \in \mathcal{X}} \min_{\gamma \geq 0} L(x, \gamma) = \min_{\gamma \geq 0} f(\gamma)$ holds.

The problem now consists in finding $x^* \in \mathcal{X}$ and $\gamma^* \geq 0$ such that

$$|f(\gamma^*) - c^T x^*| \approx 0 \text{ and } \|\max\{0, Ax^* - b\}\| \approx 0.$$

Prosaically speaking, we would like to find x^* which is nearly feasible and nearly optimal. The dual variables γ^* serve to estimate z from above. We note that evaluating f for some γ amounts to solving a problem of type (7). Our idea is now to get a good approximation of $\min f(\gamma)$ with only a limited number of function evaluations. The bundle method is a good way to achieve this goal.

3.2 The bundle method: basic ideas

The bundle method goes back at least to the 1970's, see e.g. [15, 14, 25]. It was originally developed to minimize a nonsmooth convex function $f(\gamma)$ over $\gamma \in \mathbb{R}^n$. The function f is assumed to be given by an oracle, which, for some input γ returns the function value $f(\gamma)$ and a vector g contained in the subdifferential of f at γ , $g \in \partial f(\gamma)$. Imposing the sign constraints $\gamma \geq 0$, as required in our situation, does not make the problem much harder. We will use the bundle method tailored for our problem, hence it is useful to briefly recall its basic ideas and practical issues.

If $f(\gamma)$ is given as before, $f(\gamma) = b^T \gamma + \max_{x \in \mathcal{X}} (c - A^T \gamma)^T x$, the bundle method is most conveniently explained as follows. For given $\gamma \geq 0$ we solve (10) returning $f(\gamma)$ and x , such that (γ, x) is a matching pair. A subgradient $g \in \partial f(\gamma)$ is then implicitly given by $g = b - Ax$, as can easily be verified.

We start with some initial $\gamma_{start} \geq 0$, for instance $\gamma_{start} = 0$ and evaluate f . Thus initially we have a matching pair $(\gamma_{start}, x_{start})$ with $f(\gamma_{start}) = L(x_{start}, \gamma_{start})$.

The algorithm is iterative and maintains in each iteration a currently best approximation $\hat{\gamma}$ to the minimizer of f and a sequence $X = (x_1, \dots, x_k)$ where each $x_i \in \mathcal{X}$ and $(\hat{\gamma}, x_k)$ is a matching pair.

To describe a general step, we assume to have $X = (x_1, \dots, x_k)$ and $\hat{\gamma}$, with $(\hat{\gamma}, x_k)$ a matching pair. Given X , we compute the following information:

$$g_i := b - Ax_i, \quad G = (g_1, \dots, g_k), \quad \phi_i := c^T x_i \text{ and } F = (\phi_1, \dots, \phi_k)^T. \quad (11)$$

The bundle method now combines two ideas to determine a new trial point γ_{test} . First, the function $f(\gamma)$ is approximated by

$$f_{appr}(\gamma) := \max\{L(x, \gamma) : x \in \text{conv}(x_1, \dots, x_k)\}.$$

Since $x \in \text{conv}(x_1, \dots, x_k)$ if and only if $\exists \lambda \in \Lambda := \{\lambda \in \mathbb{R}^k : \lambda \geq 0, e^T \lambda = 1\}$ with $x = X\lambda$, we can write out the definition of f_{appr} to get

$$f_{appr}(\gamma) = \max_{\lambda \in \Lambda} b^T \gamma + (c - A^T \gamma)^T (X\lambda) = \max_{\lambda \in \Lambda} F^T \lambda + \gamma^T G \lambda,$$

using the definition of F and G , and noting that $G\lambda = b - AX\lambda$.

The second ingredient of the bundle method consists in the proximal point idea, which penalizes displacements from the currently best point $\hat{\gamma}$ with a term proportional to $\|\gamma - \hat{\gamma}\|^2$.

In summary therefore, the bundle method asks to find a new trial point $\gamma_{test} \geq 0$ by minimizing, for some prescribed parameter $t > 0$ the function

$$f_{appr}(\gamma) + \frac{1}{2t} \|\gamma - \hat{\gamma}\|^2 \tag{12}$$

over the set $\gamma \geq 0$. We obtain the following dual problem to be solved.

$$\begin{aligned} \min_{\gamma \geq 0} \max_{\lambda \in \Lambda} F^T \lambda + (G\lambda)^T \gamma + \frac{1}{2t} \|\gamma - \hat{\gamma}\|^2 \\ = \min_{\gamma} \max_{\lambda \in \Lambda, \eta \geq 0} F^T \lambda + (G\lambda)^T \gamma + \frac{1}{2t} \|\gamma - \hat{\gamma}\|^2 - \gamma^T \eta \\ = \max_{\lambda \in \Lambda, \eta \geq 0} \min_{\gamma} F^T \lambda + (G\lambda)^T \gamma + \frac{1}{2t} \|\gamma - \hat{\gamma}\|^2 - \gamma^T \eta \end{aligned}$$

The inner minimization problem is convex quadratic in γ , hence we can replace it by insuring $\frac{\partial}{\partial \gamma}(\cdot) = 0$ which is equivalent to $\gamma = \hat{\gamma} + t(\eta - G\lambda)$. Therefore minimizing $f_{appr}(\gamma) + \frac{1}{2t} \|\gamma - \hat{\gamma}\|^2$ is dual to the following maximization

$$\max_{\lambda \in \Lambda, \eta \geq 0} (F + G^T \hat{\gamma})^T \lambda - \frac{t}{2} \|\eta - G\lambda\|^2 - \hat{\gamma}^T \eta, \tag{13}$$

yielding λ and η . The minimizing γ of (12) is then given by $\gamma = \hat{\gamma} + t(\eta - G\lambda)$.

The final problem in λ and η is easy, if one set of the variables is kept constant: keeping η constant results in a convex quadratic problem over the set Λ . Keeping λ constant allows to solve for η coordinatewise, see for instance [10]. Thus we start with $\eta = 0$, solve for λ which we then keep constant to solve for η and iterate this process several times to get (approximate) solutions λ, η of (13). Using these estimates λ and η we get the new trial point

$$\gamma_{test} = \hat{\gamma} + t(\eta - G\lambda).$$

As a byproduct of minimizing (12) we also get a 'primal' point

$$\hat{x} := X\lambda \tag{14}$$

in the convex hull of the current bundle X , which is least infeasible among all points in the convex hull of X .

Note also that $\lambda_i = 0$ implies that the column x_i of X has no influence on the maximization. We use this observation and remove any columns of X with corresponding $\lambda_i = 0$.

The bundle method now asks to evaluate f at γ_{test} , producing the function value $f(\gamma_{test})$ and the matching pair $(\gamma_{test}, x_{test})$. Using some standard criteria, we decide whether or not γ_{test} becomes the currently best point, we update our information and start a new iteration.

One iteration of the bundle method therefore has the following two significant steps:

- Solve (12) approximately by first solving a sequence of convex quadratic problems in the k variables λ . This gives the new trial point γ_{test} .
- Evaluate f at γ_{test} .

In our applications, the last step is the dominating operation in each iteration. At termination, the bundle method returns $\hat{\gamma}$ as an approximate dual solution, and \hat{x} , see (14) as an approximate primal solution. We note, that these do not form a matching pair, but the general convergence theory for the bundle method, see e.g. [16, 25, 7], insures that under appropriate stopping conditions these vectors converge towards an optimal primal-dual solution pair of (8).

3.3 A dynamic version of the bundle method

Contrary to the “classical” use of the bundle method, where the function $f(\gamma)$ to be minimized is considered to be fixed, we have a more flexible version in mind, where f will change in the course of the algorithm. To get the initial function $f(\gamma)$ we solve

$$\max\{c^T x : x \in \mathcal{X}\}$$

yielding an initial maximizer x^* . In principle we could now define $f(\gamma)$ by dualizing all the constraints $Ax \leq b$. Recall however, that the number m of constraints may be substantially larger than the dimension of the problem. To maintain efficiency, we are interested only in those constraints, which are likely to be active at the optimum. Therefore we look at $r^* := b - Ax^*$ and set $r_{\min} := \min\{r_i^*\}$. In the unlikely event that $r_{\min} \geq 0$, x^* is optimal for (8) and we are done. Otherwise $r_{\min} < 0$ and we consider now only those constraints from $Ax \leq b$ which are ‘badly’ violated by x^* to define f .

Specifically, let $0 < \alpha < 1$ and set

$$I := \{i : r_i^* \leq \alpha r_{\min}\}.$$

We denote by A_I the submatrix of A with rows indexed by I , a similar definition holds for b_I . For notational convenience we define

$$f_I(\gamma_I) := \max_{x \in \mathcal{X}} c^T x + (b_I - A_I x)^T \gamma_I.$$

Thus we start by minimizing f_I using the bundle method described above. In order to maintain computational efficiency, we include as an additional stopping condition an upper bound on the number of iterations. Therefore, when the method stops, returning some $\hat{\gamma}_I$ and some \hat{x} , it is likely that $\hat{\gamma}_I$ is still far from the true minimizer of f_I . To continue, we need to decide on the following two issues:

- Which of the constraints in I should be kept?
- Are there additional constraints that should be added to I ?

We use $\hat{\gamma}_I$ to answer to the first question. A large value $\hat{\gamma}_i$ indicates that the constraint i is binding and hence should not be removed. $\hat{\gamma}_i = 0$ indicates that constraint i may be inactive, and therefore could be removed. In summary, we use $\hat{\gamma}_I$ to purge I by removing constraints i with a value $\hat{\gamma}_i$ smaller than a prescribed fraction of $\max(\hat{\gamma}_I)$.

It is less obvious to decide which new constraints should be added to I . In our experiments we found that using a convex combination of \hat{x} and the previous point \hat{x}_{old} to identify violated constraints gives the most satisfactory behaviour of the algorithm.

4 Solving the relaxations with the bundle method

We now turn to computational experiments with our version of the bundle method. We consider Max-Cut and Equipartition problems in the following two subsections separately.

4.1 Solving the semidefinite relaxation (4) of Max-Cut

To make our experiments reproduceable and accessible to other researchers, we generate our test instances using the graph generator `rudy`, introduced by Giovanni Rinaldi. It is available under <http://www-user.tu-chemnitz.de/~helmsberg/rudy.tar.gz>.

We generate three types of instances, all with edge weights c_{ij} randomly chosen from $\{1, -1\}$: dense graphs, sparse graphs with edge density of 10 %, and three-dimensional grid graphs. Finally we take some graphs from the literature, see [11].

The calls to `rudy` to generate these graphs are given in the appendix. We first give some statistics about these graphs in Table 4, including the number of vertices (n), the number of edges ($|E|$) and the value of the best cut known to us. We are mostly interested in difficult instances, where neither the approximation results of Goemans and Williamson ($a_{ij} \geq 0$) nor the theory of Nesterov ($L \succeq 0$) applies. We will show that even for this difficult class of Max-Cut instances, the relaxation (4) provides a good (a-posteriori) approximation to the true value of the Max-Cut problem, and in particular constitutes a nontrivial improvement of the basic relaxation (3).

We apply the bundle method with the modifications previously described to these graphs. For problems with $n \leq 729$ we have set a limit of 100 function evaluations, for the bigger problems ($800 \leq n \leq 1000$) we allowed only 50 function evaluations. We use the graphs on 2000 vertices to indicate also the limits of the present approach. For these graphs we allow only 30 function evaluations to estimate the optimum of (4).

In Table 4 we summarize our computational results. The table identifies the graphs (in column 1), then gives the initial bound (optimal value of (3)) and the final bound found after the maximum number of function evaluations was reached. We also include the gap (in %) of the bounds relative to the value of the best known cut. The last two columns contain the number of $\gamma_i > 0$, i.e. the number of inequalities actually active at the stop, and the total computation times in minutes on a PC (Pentium 4, 1.7 Ghz).

Accuracy: The first question is to estimate, how good our bounds approximate the correct optimal value of (4). To estimate the quality of our bounds, we use the final set of active constraints of the bundle method, to solve the relaxation with these constraints by interior-point methods. We kept adding violated constraints, until the largest violation of all triangle constraints was less than 0.02. It is interesting to observe that the constraints identified by our bundle approach serve as a very good starting point to solve (4). It took typically only very few rounds (at most 4) of adding additional violated constraints until the largest violation fell below the desired threshold. We collect the results in Table 5 for some of the smaller problems. We note that the bundle method provides very good approximations to the true optimal value of (4). The difference is in all cases except `spin5` significantly smaller than 1%, hence we may safely assume that the results of Table 4 give tight approximations to (4). We emphasize that solving (4) without any knowledge of a good starting set of triangle constraints would take much longer than the times given in the last column of Table 5. We recall the discussion in subsection 2.2 and the experience reported in [11].

Bound versus Bundle Iterations: We now take a closer look at the development of

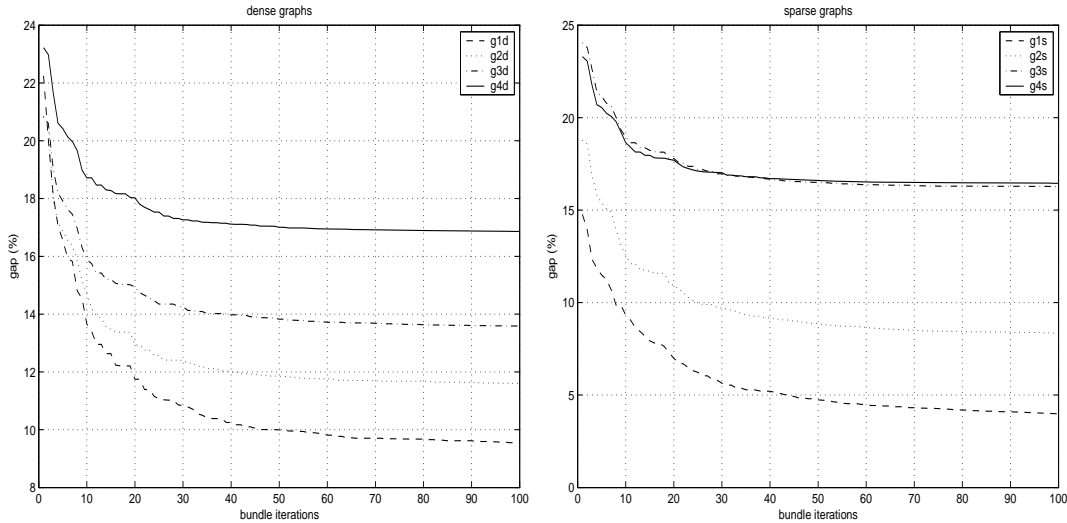


Figure 1: Reduction of gap (in %) for the dense and sparse graphs

the bounds during the iterations of the bundle method. In Figure 1 we provide a graphical development of the results for the dense graphs $g1d, \dots, g4d$ and also for the sparse graphs $g1s, \dots, g4s$. All plots show a similar qualitative behaviour: there is quick improvement of the bound, and a very flat tail.

In Figure 2 we also give the development of the bounds for the three-dimensional grid graphs, and here one would expect still some improvement of the bounds, if more iterations were carried out. This may explain the bigger error for $spin5$ in Table 4 as compared to Table 5. The figures also show that the biggest improvement is obtained in the first iterations. This is an attractive feature, because it allows the user to trade quality against computation time.

Big graphs: Looking at the graphs on 2000 nodes, we also see the limits of the present approach. The function evaluation itself is already quite expensive, see also Table 1, leading to computation times of several hours. We note however that even for these large problems, the inclusion of triangle inequalities provides a significant improvement of the initial basic semidefinite programming bound.

LP versus SDP: Finally, we compare solving (4) with and without the semidefiniteness constraint on X . Optimizing over the triangle inequalities alone amounts to solving a Linear Program, so one might think that this is a straight-forward task. Unfortunately, this is not the case. In Table 6 we summarize the relevant results for all the instances where we managed to compute the optimal solution of the triangle relaxation. We note that the linear relaxation is quite strong in case of the grid graphs, but tends to be very weak for all other instances. We find it surprising that solving (at least approximately) the relaxation (4) seems easier than solving the relaxation with the semidefiniteness condition omitted.

4.2 Computational results for Equipartitioning

We now turn to the equipartition problem. Here we use the data from [17]. These are graphs with edge density of 80 % and edge weights drawn randomly from $\{1, \dots, 100\}$ and are available from <http://www.math.uni-klu.ac.at/or/software>. We have scaled the cost elements by

problem	n	$ E $	cut	initial bd	gap (%)	final bd	gap (%)	m	time
g1d	100	4901	324	396.1	22.25	353.6	9.14	1299	0.44
g2d	200	19701	1050	1268.9	20.85	1170.4	11.47	2427	2.49
g3d	300	44402	1953	2359.6	20.82	2218.3	13.58	2996	6.63
g4d	400	79002	2979	3670.8	23.22	3481.1	16.86	3644	14.20
g5d	500	123503	3911	4937.3	26.24	4728.6	20.91	4431	25.50
g1s	100	495	126	144.6	14.77	130.8	3.81	1816	0.42
g2s	200	1990	318	377.7	18.77	344.2	8.24	2934	2.36
g3s	300	4485	547	678.5	24.04	635.7	16.22	3158	6.25
g4s	400	7980	907	1118.3	23.30	1055.7	16.39	4433	13.20
g5s	500	12475	1234	1557.6	26.22	1487.6	20.55	4321	24.57
spin5	125	375	108	125.3	15.99	111.1	2.87	3714	0.73
spin6	216	648	182	211.8	16.37	186.9	2.69	6223	3.16
spin7	343	1029	304	347.9	14.44	308.4	1.45	11323	10.55
spin8	512	1536	454	524.3	15.48	463.1	2.00	18375	33.44
spin9	729	2187	630	744.7	15.27	666.7	5.83	18265	42.75
spin10	1000	3000	874	1032.6	18.15	921.4	5.42	27830	106.1
G1	800	19176	11575	12083.2	4.39	12005.4	3.72	7372	51.76
G6	800	19176	2133	2656.2	24.53	2566.2	20.31	6983	43.11
G11	800	1600	564	629.2	11.56	572.7	1.54	15946	60.20
G14	800	4694	3032	3139.6	5.26	3140.7	3.59	8973	59.68
G18	800	4694	965	1166.0	20.83	1063.4	10.20	17635	69.19
G22	2000	19990	13220	14135.9	6.93	14045.8	6.25	18325	278.06
G27	2000	19990	3186	4141.7	30.00	4048.4	27.05	15178	406.66
G39	2000	11779	2314	2877.7	24.36	2672.7	15.50	26471	533.36

Table 4: Results for Max-Cut. The initial bound is the solution of (3). The final bound is obtained after 100 function evaluations with the bundle method in case $n \leq 729$, after 50 iterations in case $800 \leq n \leq 1000$ and after 30 iterations for $n = 2000$. m denotes the number of active constraints at the stop of the bundle method.

multiplying with 10^{-3} .

The function evaluation of the bundle method amounts now to solving the basic relaxation (5). We dualize the sign constraints, and summarize the results in Table 7. Again, the bundle method provides a strong method to deal with many inequalities efficiently. Looking at the number of active constraints (column 7), it is clear that a direct approach to solve any of the relaxations using interior-point methods is hopeless. We also note that including the sign constraints constitutes a significant improvement of the initial basic relaxation. We provide a graphical comparison of the improvement by comparing the initial gap with the final gap in Figure 3. We point out that the improvement using the bundle method as compared to a direct approach based entirely on interior-point methods, see [17], is impressive.

In summary, the results provided here should give enough computational evidence that the bundle method in combination with simple semidefinite programs is a good way to approximate more complicated semidefinite programs.

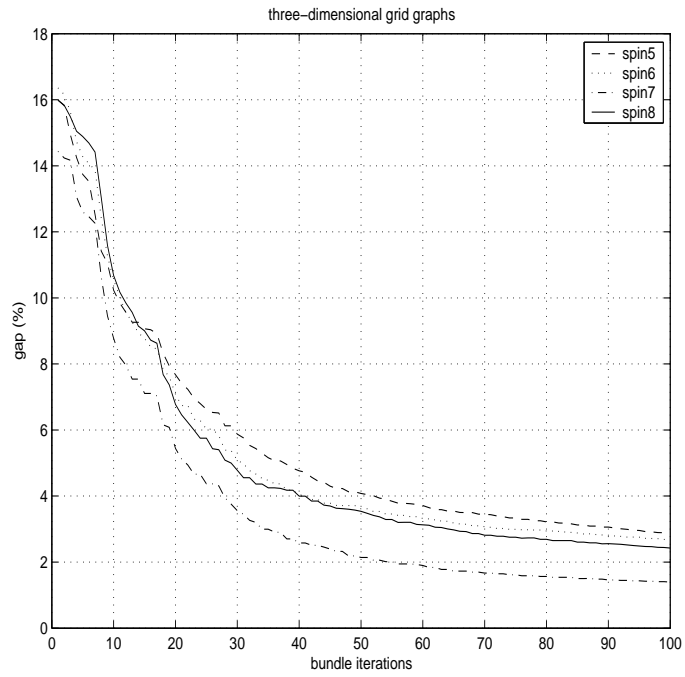


Figure 2: Reduction of gap (in %) for the grid graphs




Figure 3: Reduction of gap (in %) for equipartition

problem	n	z_{mc-met}	active constr.	100 bundle it.	gap (%)	time
g1d	100	352.375	1053	353.6	0.35	9.1
g2d	200	1167.992	2108	1170.4	0.20	37.7
g3d	300	2215.257	3053	2218.3	0.14	112.6
g1s	100	130.008	1371	130.8	0.60	16.0
g2s	200	343.156	2376	344.2	0.30	50.9
spin5	125	109.344	2616	110.8	1.60	180.0

Table 5: Exact solution of (4) for smaller instances, compared to 100 iterations of the bundle method. The gap refers to the relative error of the bundle bound compared to the true optimum.

problem	LP	gap	SDP	gap
spin5	110.3	2.13	111.1	2.87
spin6	187.1	2.80	186.9	2.69
spin7	307.5	1.15	308.4	1.45
spin8	460.9	1.52	463.1	2.00
spin9	658.8	4.57	666.7	5.83
spin10	909.4	4.05	921.4	5.42
g1s	137.5	9.13	130.8	3.81
g2s	410.0	28.93	344.2	8.24
g3s	854.4	56.20	635.7	16.22
g1d	786.3	142.69	353.6	9.14

Table 6: Comparing (4) (SDP) with the LP-based bound where the semidefiniteness condition on X is omitted.

5 Extensions and Conclusion

We briefly sketch some other combinatorial optimization problems, where our approach might provide an attractive alternative to pure interior-point based methods.

We start out with the **Quadratic Assignment Problem (QAP)**, which can be expressed as follows:

$$\min \operatorname{tr}(AXB + C)X^T$$

over the set of permutation matrices X . In the recent dissertation [26], an approach similar to the one presented here is applied to obtain bounds for QAP. The resulting bounds are currently by far the best ones available. Since the derivation of the relaxations includes many technicalities that would go beyond the limits of an introductory paper like this one, we refer to [26, 24] for further details and computational results. The technical report [24] is available at <http://www.math.uni-klu.ac.at/or/Forschung/>.

Another interesting source of problems of the type considered in this paper is given by the ϑ -function of a graph. The number $\vartheta(G)$ associated to a graph G separates the clique number $\omega(G)$ from the chromatic number $\chi(G)$,

$$\omega(G) \leq \vartheta(G) \leq \chi(G),$$

see the seminal paper [18]. $\vartheta(G)$ is the optimal value of a semidefinite program of rather simple

problem	feas sol.	initial bd	gap %	final bd	(gap %)	m	total time
r100	168.712	155.1	8.78	165.3	2.06	2238	0.52
r200	680.633	641.4	6.11	665.3	2.30	6204	3.35
r300	1558.767	1485.9	4.91	1525.3	2.20	10938	9.08
r400	2768.879	2659.6	4.11	2712.9	2.06	16083	19.12
r500	4354.869	4195.7	3.79	4267.7	2.04	22421	34.51
r200	765.728	698.4	9.64	757.5	1.08	15489	4.32
r300	1739.303	1617.9	7.50	1719.1	1.18	32642	11.04
r400	3078.133	2895.9	6.30	3039.9	1.26	54996	22.45
r500	4825.907	4568.7	5.63	4762.9	1.32	83030	40.01

Table 7: Results for equipartition, $k=10$ and $k=50$.

structure. This semidefinite program can now be strengthened to give better approximations to $\omega(G)$, see e.g. [19], or to get tighter approximations of the chromatic number, see [20]. These tighter models again are computationally intractable by standard interior-point methods, hence they are another class of problems, where our bundle approach might be a practical alternative.

To summarize, we draw the following conclusions from our computational experiments presented in this paper:

- The bundle method in combination with interior-point methods is a good way to approximate semidefinite programs with a huge number of combinatorial cutting planes.
- The present approach provides the currently strongest bounds at reasonable computational cost for several NP-hard optimization problems like Max-Cut, Equipartitioning and also the Quadratic Assignment Problem.
- The number of function evaluations to reach good approximations is surprisingly small.
- It is also surprising that solving (4) seems less difficult than solving (4) with the semidefiniteness condition omitted.

6 Appendix: rudy-calls to generate the test graphs

```

%% rudy call
% 3d-springlass random 1,-1 weights
rudy -springlass3pm 5 5 5 50 5551 > spin5
rudy -springlass3pm 6 6 6 50 6661 > spin6
rudy -springlass3pm 7 7 7 50 7771 > spin7
rudy -springlass3pm 8 8 8 50 8881 > spin8
rudy -springlass3pm 9 9 9 50 9991 > spin9
rudy -springlass3pm 10 10 10 50 1010101 > spin10

% dense graphs: random 1,-1 weights
rudy -rnd_graph 100 99 1001 -random 0 1 1001 -times 2 -plus -1 > g1d
rudy -rnd_graph 200 99 2001 -random 0 1 2001 -times 2 -plus -1 > g2d
rudy -rnd_graph 300 99 3001 -random 0 1 3001 -times 2 -plus -1 > g3d
rudy -rnd_graph 400 99 4001 -random 0 1 4001 -times 2 -plus -1 > g4d
rudy -rnd_graph 500 99 5001 -random 0 1 5001 -times 2 -plus -1 > g5d

% sparse (10% density) graphs: random 1,-1 weights
rudy -rnd_graph 100 10 1001 -random 0 1 1001 -times 2 -plus -1 > g1s
rudy -rnd_graph 200 10 2001 -random 0 1 2001 -times 2 -plus -1 > g2s
rudy -rnd_graph 300 10 3001 -random 0 1 3001 -times 2 -plus -1 > g3s
rudy -rnd_graph 400 10 4001 -random 0 1 4001 -times 2 -plus -1 > g4s
rudy -rnd_graph 500 10 5001 -random 0 1 5001 -times 2 -plus -1 > g5s

```

References

- [1] F. ALIZADEH. Interior Point methods in semidefinite programming with applications to combinatorial optimization. *SIAM Journal on Optimization*, 5:13–51, 1995.
- [2] L. BAHIENSE, N. MACULAN, and C. SAGASTIZABAL. The volume algorithm revisited: relation with bundle methods. *Mathematical Programming*, 94:41–69, 2002.
- [3] F. BARAHONA and R. ANBIL. The Volume algorithm: producing primal solutions with a subgradient method. *Mathematical Programming*, 87:385–399, 2000.
- [4] F. BARAHONA and L. LADANYI. Branch and Cut based on the Volume Algorithm: Steiner trees in graphs and Max-Cut. 2001.
- [5] C. DELORME and S. POLJAK. Laplacian eigenvalues and the maximum cut problem. *Mathematical Programming*, 62:557–574, 1993.
- [6] W.E. DONATH and A.J. HOFFMAN. Lower bounds for the partitioning of graphs. *IBM J. of Research and Development*, 17:420–425, 1973.
- [7] A. FRANGIONI. Generalized bundle methods. *SIAM Journal on Optimization*, 13:117–156, 2002.
- [8] M.X. GOEMANS and D.P. WILLIAMSON. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42:1115–1145, 1995.
- [9] C. HELMBERG. *An interior point method for semidefinite programming and max-cut bounds*. PhD thesis, Graz University of Technology, Austria, 1994.
- [10] C. HELMBERG, K.C. KIWIEL, and F. RENDL. Incorporating inequality constraints in the spectral bundle method. In E.A. Boyd R.E. Bixby and R.Z. Rios-Mercado, editors, *Integer Programming and combinatorial optimization*, pages 423–435. Springer Lecture Notes 1412, 1998.
- [11] C. HELMBERG and F. RENDL. Solving quadratic (0,1)-problems by semidefinite programming and cutting planes. *Mathematical Programming*, 82:291–315, 1998.
- [12] C. HELMBERG, F. RENDL, R. VANDERBEI, and H. WOLKOWICZ. An interior-point method for semidefinite programming. *SIAM Journal on Optimization*, 6:342–361, 1996.
- [13] S.E. KARISCH and F. RENDL. Semidefinite Programming and Graph Equipartition. *Fields Institute Communications*, 18:77–95, 1998.
- [14] K.C. KIWIEL. *Methods of descent for nondifferentiable optimization*. Springer, Berlin, 1985.
- [15] C. LEMARECHAL. Nonsmooth optimization and descent methods. Technical report, International Institute for Applied Systems Analysis, 1978.
- [16] C. LEMARECHAL, A. NEMIROVSKII, and Y. NESTEROV. New variants of bundle methods. *Mathematical Programming*, 69:111–147, 1995.

- [17] A. LISSER and F. RENDL. Graph partitioning using Linear and Semidefinite Programming. *Mathematical Programming (B)*, 95:91–101, 2002.
- [18] L. LOVÁSZ. On the shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25:1–7, 1979.
- [19] L. LOVÁSZ and A. SCHRIJVER. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1:166–190, 1991.
- [20] P. MEURDESOLF. Strengthening the lovasz bound for graph coloring. *Mathematical Programming (forthcoming)*.
- [21] Y. NESTEROV. Quality of semidefinite relaxation for nonconvex quadratic optimization. Technical report, 1997.
- [22] S. POLJAK and F. RENDL. Nonpolyhedral relaxations of graph bisection problems. *SIAM Journal on Optimization*, 5:467–487, 1995.
- [23] B.T. POLYAK. A general method for solving extremum problems. *Soviet Mathematics*, 8:593–597.
- [24] F. RENDL and R. SOTIROV. Bounds for the quadratic assignment problem using the bundle method. Technical report, University of Klagenfurt, Austria, 2003.
- [25] H. SCHRAMM and J. ZOWE. A version of the bundle idea for minimizing a nonsmooth function: Conceptual idea, convergence analysis, numerical results. *SIAM J. Optimization*, 2:121–152, 1992.
- [26] R. SOTIROV. *The Bundle Method in combinatorial optimization*. PhD thesis, University of Klagenfurt, Austria, 2003.
- [27] M.J. TODD. A study of search directions in primal-dual interior-point methods for semidefinite programming. *Optimization Methods and Software*, 11:1–46, 1999.