Incorporating Inequality Constraints in the Spectral Bundle Method

Christoph Helmberg[1], Krzysztof C. Kiwiel[2], and Franz Rendl[3]

[1] Konrad-Zuse-Zentrum für Informationstechnik Berlin
Takustraße 7, 14195 Berlin, Germany
helmberg@@zib.de, http://www.zib.de/helmberg

[2] Systems Research Institute, Polish Academy of Sciences
Newelska 6, 01-447 Warsaw, Poland
kiwiel@@ibspan.waw.pl

[3] Technische Universität Graz, Institut für Mathematik
Steyrergasse 30, A-8010 Graz, Austria
rendl@@opt.math.tu-graz.ac.at

Abstract. Semidefinite relaxations of quadratic 0-1 programming or graph partitioning problems are well known to be of high quality. However, solving them by primal-dual interior point methods can take much time even for problems of moderate size. Recently we proposed a spectral bundle method that allows to compute, within reasonable time, approximate solutions to structured, large equality constrained semidefinite programs if the trace of the primal matrix variable is fixed. The latter property holds for the aforementioned applications. We extend the spectral bundle method so that it can handle inequality constraints without seriously increasing computation time. This makes it possible to apply cutting plane algorithms to semidefinite relaxations of real world sized instances. We illustrate the efficacy of the approach by giving some preliminary computational results.

1 Introduction

Since the landmark papers [9,4,10,2] it is well known that semidefinite programming allows to design powerful relaxations for constrained quadratic 0-1 programming and graph partitioning problems. The most commonly used algorithms for solving these relaxations, primal-dual interior point methods, offer little possibilities to exploit problem structure. Typically their runtime is governed by the factorization of a dense symmetric positive definite matrix in the number of constraints and by the line search that ensures the positive definiteness of the matrix variables. Computation times for problems with more than 3000 constraints or matrix variables of order 500, say, are prohibitive. Very recently, a pure dual approach has been proposed in [1] that is able to exploit the sparsity of the cost matrix in the case of the max-cut relaxation with diagonal

R. E. Bixby, E. A. Boyd, and R. Z. Ríos-Mercado (Eds.): IPCO VI
LNCS 1412, pp. 423–435, 1998.

constraints. It is not yet clear whether these results extend to problems with a huge number of less structured constraints.

The spectral bundle method [6] works on a reformulation of semidefinite relaxations as eigenvalue optimization problems and was developed to provide approximate solutions to structured problems fast. In contrast to standard bundle methods [7,11], a non-polyhedral semidefinite cutting plane model is constructed from the subgradients. Reinterpreted in terms of the original semidefinite program, the semidefinite model ensures positive semidefiniteness of the dual matrix variable on a subspace only. Subgradients correspond to eigenvectors to negative eigenvalues and are used to correct the subspace. By means of an aggregate subgradient the dimension of the subspace can be kept small, thus ensuring efficient solvability of the subproblems. Lanczos methods (see e.g. [3]) allow to compute a few extremal eigenvalues and their eigenvectors efficiently by a series of matrix-vector multiplications which do not require the matrix in explicit form. Thus structural properties of cost and coefficient matrices can be exploited.

Like most first order methods, the spectral bundle method exhibits fast progress in the beginning, but shows a strong tailing off effect as the optimal solution is approached. Fortunately many semidefinite relaxations do not have to be solved exactly. Rather an approximate solution is used to improve, e.g., by cutting planes, the current relaxation, which is then resolved.

In its original form the spectral bundle method is designed for equality constraints only because sign constraints on the dual variables may increase computation times for the semidefinite subproblem significantly. In this paper we employ Lagrangian relaxation to approximate the solution of a sign constrained semidefinite subproblem. Surprisingly just one update of Lagrange multipliers per function evaluation suffices to ensure convergence. The semidefinite subproblem can be solved as efficiently as in the unconstrained case, thus rendering this method an attractive choice for large scale semidefinite cutting plane algorithms.

Section 2 introduces some notation and explains the connection to semidefinite programming. This is followed by a very brief review of important properties of the maximal eigenvalue function. Section 4 explains the extension of the spectral bundle method to inequality constraints. In Section 5 we discuss efficiency aspects of the subproblem solution. Section 6 gives computational results.

2 Semidefinite Programs

Let $\mathcal{S}_n$ denote the space of symmetric matrices of order n. The inner product in this space is the usual matrix inner product, $\langle A, B\rangle = \operatorname{tr}(B^T A)$ for $A, B \in \mathbb{R}^{m\times n}$. Let $\mathcal{S}_n^+$ denote the set of symmetric positive semidefinite matrices. $\mathcal{S}_n^+$ is a pointed closed convex cone. Except for its apex $\{0\}$, a face F of this cone can be described as

$$F = \left\{PVP^T : V \in \mathcal{S}_r^+\right\},$$

where $P \in \mathbb{R}^{n\times r}$ is some fixed matrix with orthonormal columns (w.l.o.g.). The dimension of such a face F is $\binom{r+1}{2}$. For $A, B \in \mathcal{S}_n$, $A \succeq B$ refers to the Löwner partial order induced by the cone $\mathcal{S}_n^+$ ($A \succeq B \iff A - B \in \mathcal{S}_n^+$).