

**A unifying framework
for several cutting plane methods for
semidefinite programming**

John E. Mitchell

Department of Mathematical Sciences

Rensselaer Polytechnic Institute

Troy, NY 12180 USA

mitchj@rpi.edu

<http://www.rpi.edu/~mitchj>

Joint work with **Kartik Krishnan**.

Multiscale Optimization Methods and Applications

University of Florida

28 February 2004

Abstract

Cutting plane methods provide the means to solve large scale semidefinite programs (SDP) cheaply and quickly. We give a survey of various cutting plane approaches for SDP in this paper. These cutting plane approaches arise from various perspectives, and include techniques based on interior point cutting plane approaches, and an approach which mimics the simplex method for linear programming. We present an accessible introduction to various cutting plane approaches that have appeared in the literature. We place these methods in a unifying framework which illustrates how each approach arises as a natural enhancement of a primordial linear programming cutting plane scheme based on a semi-infinite formulation of the SDP.

Contents

1. Semidefinite programming
2. A semi-infinite linear program
3. Nonpolyhedral cutting plane models
4. Approximately solving the dual subproblem
5. Properties and a polar cone
6. Complexity of the algorithm
7. Conclusions

1 Semidefinite programming

Standard form:

$$\begin{aligned} \min \quad & C \bullet X \\ \text{subject to} \quad & \mathcal{A}(X) = b \quad (SDP) \\ & X \succeq 0, \end{aligned}$$

with dual

$$\begin{aligned} \max \quad & b^T y \\ \text{subject to} \quad & \mathcal{A}^T y + S = C \quad (SDD) \\ & S \succeq 0 \end{aligned}$$

- C , X and S are square symmetric matrices.
- Require X and S be **positive semidefinite** (psd).
- Inner products are the **Frobenius inner product**, so $C \bullet X = \text{trace}(CX)$.

- $\mathcal{A}X = \begin{bmatrix} A_1 \bullet X \\ \vdots \\ A_k \bullet X \end{bmatrix}$, $\mathcal{A}^T y = \sum_{i=1}^k y_i A_i$, where each A_j is a symmetric matrix.

- Denote

$$Y := \{y \in \mathbb{R}^m : C - \mathcal{A}^T y \succeq 0\}.$$

Applications

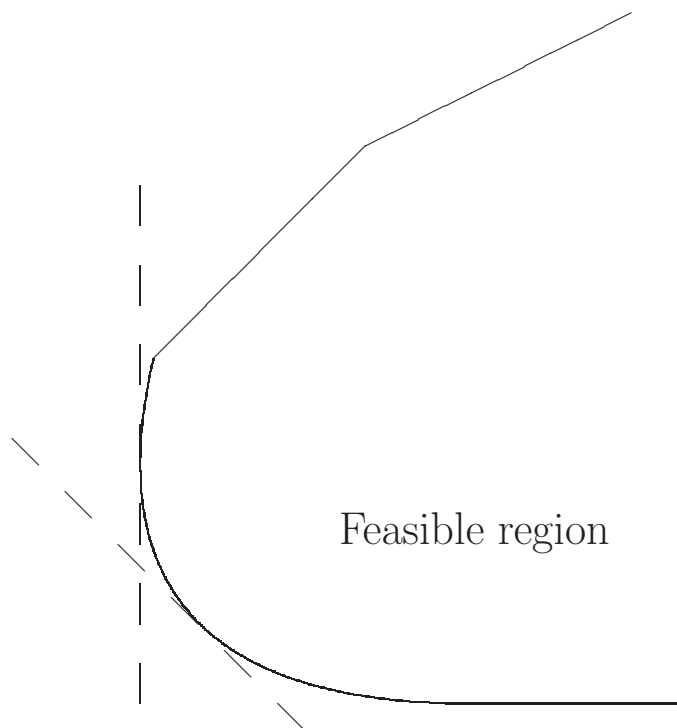
- control
- developing approximation algorithms for combinatorial optimization problems
- finance
- statistics
- ...

Methods for solving SDPs

- **Primal-dual interior point approaches** are limited in the size of problems they can solve.
- Alternatives include:
 - The **bundle method** of Helmberg, Rendl, and Kiwiel.
 - The **nonlinear programming approaches** of Burer, Monteiro, and Zhang and Benson and Vanderbei.
 - Methods based on **cutting plane models**: Krishnan and Mitchell, Goffin and Oskoorouchi.

2 A semi-infinite linear program

- (SDP) and (SDD) are **convex programming** problems.
- Only nonlinearity is the positive semidefiniteness (PSD) requirement.
- **Replace PSD constraint by linear constraints.**



Dual formulation:

$$\begin{aligned} \max \quad & b^T y && (LDD) \\ \text{subject to} \quad & d^T (C - \mathcal{A}^T y) d \geq 0 \quad \forall d \text{ with } \|d\|^2 = 1. \end{aligned}$$

- If (SDP) has m constraints, then (LDD) has m variables.
- Need to be **selective** with the vectors d included in the constraints. For example, use a **cutting plane approach** to select vectors d .
- Typically, we will need some vectors d that are **dense**, leading to a dense linear programming problem.
- Have far fewer variables than in a primal formulation, so get smaller linear programs (although they are still dense).

Regaining a primal solution

- We only use a finite number of vectors d , so we are solving a **relaxation** of the dual:

$$\begin{aligned} \max \quad & b^T y && (LDR) \\ \text{subject to} \quad & d_i^T (C - \mathcal{A}^T y) d_i \geq 0 \\ & \text{for vectors } d_i, i = 1, \dots, m. \end{aligned}$$

- The optimal value to (LDR) gives an upper bound on the optimal value of (SDP) .
- (LDR) is a relaxation of (SDD) . The dual linear program to (LDR) is a **constrained version** of (SDP) :

$$\begin{aligned} \min \quad & C \bullet (DMD^T) \\ \text{s.t.} \quad & A_j \bullet (DMD^T) = b_j \quad j = 1, \dots, m \\ & M \succeq 0 && (LPR) \\ & M \text{ diagonal} \end{aligned}$$

where the columns of D are the vectors d_i .

- Thus, any feasible solution to (LPR) gives a **feasible solution to (SDP) , $X = DMD^T$.**

3 Nonpolyhedral cutting plane models

Recall that (LPR) can be written as

$$\begin{aligned}
 \min \quad & C \bullet (DMD^T) \\
 \text{s.t.} \quad & A_j \bullet (DMD^T) = b_j \quad j = 1, \dots, m \\
 & M \succeq 0 \\
 & M \text{ diagonal}
 \end{aligned} \tag{1}$$

Here $M \in \mathcal{S}^m$, and $D \in \mathbb{R}^{n \times m}$ with j th column d_j .

If the columns of D contain **eigenbases for all the strictly positive eigenvalues of X^*** , then the solution to (1) is an exact solution to (SDP).

In other cases, a solution provides an **upper bound** on this objective value.

Drop requirement that M be diagonal

Consider a relaxation of (1) dropping the requirement that M be diagonal.

$$\begin{aligned} \min \quad & C \bullet (DM D^T) \\ \text{s.t.} \quad & A_j \bullet (DM D^T) = b_j \quad j = 1, \dots, m \\ & M \succeq 0 \end{aligned} \quad (2)$$

with $M \in \mathcal{S}^m$, and $D \in \mathbb{R}^{n \times m}$.

If the rank of D is n , then (2) is essentially (SDP).

In fact, **if $\text{Range}(D) \supset \text{Range}(X^*)$, then a solution to (2) is an exact solution to (SDP).**

This is a **less stringent** requirement than the polyhedral cutting plane model, where we require the exact eigenvectors of X^* .

But this is a semidefinite program, so harder to solve than the linear program (*LPR*).

Necessary to **limit the number of columns** in D to keep the algorithm practical.

This model is the basis for the **spectral bundle method** of Helmberg.

The dual problem

The dual to (2) is

$$\begin{aligned} & \max && b^T y \\ & \text{subject to} && D^T(C - \mathcal{A}^T y)D \succeq 0 \end{aligned} \tag{3}$$

If D has rank n then this is equivalent to the original problem (SDD).

The linear programming formulation (LDR) requires only that the **diagonal entries** of $D^T(C - \mathcal{A}^T y)D$ be non-negative.

A block-diagonal model

Finally, another relaxation of (1) which is more restrictive than (2) is the following.

$$\begin{aligned}
 \min \quad & C \bullet (DM D^T) \\
 \text{s.t.} \quad & A_j \bullet (DM D^T) = b_j \quad j = 1, \dots, m \\
 & M \succeq 0 \\
 & M \text{ block diagonal}
 \end{aligned} \tag{4}$$

Dual:

$$\begin{aligned}
 \max \quad & b^T y \\
 \text{subject to} \quad & D_i^T (C - \mathcal{A}^T y) D_i \succeq 0
 \end{aligned} \tag{5}$$

where each D_i corresponds to a diagonal block of M .

This is the model of Goffin and Oskoorouchi.

Generic representation of the dual

Write the dual relaxation as

$$\begin{aligned} \max \quad & b^T y \\ \text{s.t.} \quad & D_i^T (C - \mathcal{A}^T y) D_i \succeq 0 \quad i \in \mathcal{D} \end{aligned} \tag{6}$$

where \mathcal{D} represents our current collection of constraints.

- If we add linear constraints, each D_i is a column vector.
- If we use the block diagonal formulation, each D_i is a matrix with a small number of columns.
- If we drop the restrictions on M , \mathcal{D} consists of just one constraint. Then D_i is a matrix with approximately \sqrt{k} columns, typically.

4 Approximately solving the dual subproblem

- Need to find good cutting planes.
- Useful to search for cutting planes before solving the current dual problem to optimality.
- This should lead to deeper cuts, with more of the feasible region removed.

Methods for approximate solution

- Solve the relaxation **approximately using an interior point method**.

- Use a **bundle method**:

have a proximity term $\|y - \hat{y}\|^2$ in the objective to keep iterates close to the old iterate \hat{y} .

Solve subproblem

$$\begin{aligned} \max \quad & b^T y - u \|y - \hat{y}\|^2 \\ \text{s.t.} \quad & D_i^T (C - \mathcal{A}^T y) D_i \succeq 0 \quad i \in \mathcal{D} \end{aligned}$$

for a dynamically updated parameter u .

Generic algorithm

1. Choose an initial finite set $\mathcal{D} = \{D_i\}$.
2. Approximately solve the current dual relaxation to get \bar{y} .
3. If $\bar{S} = (C - \mathcal{A}^T \bar{y}) \succeq 0$:
 - If duality gap small enough, STOP.
 - Else, solve the current dual relaxation more accurately.
4. Else find $D \in \mathbb{R}^{n \times r}$ such that $D^T(C - \mathcal{A}^T \bar{y})D \not\succeq 0$.
5. Either add D to \mathcal{D} , or aggregate D into \mathcal{D} and return to Step 2.

Algorithm	Model	\mathbf{u}	Form of \mathcal{D}	Add or aggregate?
1	Polyhedral	zero	$D \in \mathbb{R}^{n \times 1}$ only	add, letting $ \mathcal{D} $ grow
2	Polyhedral bundle	positive	$D \in \mathbb{R}^{n \times 1}$ only	add, letting $ \mathcal{D} $ grow
3	Nonlinear Block-Diag	zero	$D \in \mathbb{R}^{n \times r}$ with $r \leq \sqrt{2k}$	add, letting $ \mathcal{D} $ grow
4	Spectral bundle	positive	$D \in \mathbb{R}^{n \times r}$	aggregate, keeping $ \mathcal{D} = 1$
5	Primal active set	zero	$D \in \mathbb{R}^{n \times r}$	aggregate, keeping $ \mathcal{D} = 1$

5 Properties and a polar cone

Assume:

1. The constraints $\mathcal{A}(X) = b$ imply $\mathbf{trace}(X) = a$ for some fixed a .
2. Given \bar{y} and $\bar{S} = C - \mathcal{A}^T \bar{y}$, the column(s) of the cutting matrix D are eigenvectors of \bar{S} corresponding to the **most negative eigenvalue** λ_{\min} of \bar{S} .

Then:

There exists \hat{y} feasible in (SDD) satisfying the constraint $D^T(C - \mathcal{A}^T \hat{y})D = 0$.

Example

Take $Y \subseteq \mathbb{R}^2$ to be the vectors y that make the following matrix S positive semidefinite:

$$S = \begin{bmatrix} y_1 & y_2 & 0 \\ y_2 & y_1 - 3 & 0 \\ 0 & 0 & y_1 - 4 \end{bmatrix}.$$

Let $\bar{y} = [-1, 2]^T \in Y$.


The minimum eigenvalue of \bar{S} is -5, with multiplicity 2.

Basis for eigenspace of \bar{S} : $[1, -2, 0]^T$, $[0, 0, 1]^T$.

These vectors give the constraints

$$\begin{aligned} y_1 &\geq 4 \\ 5y_1 - 4y_2 &\geq 12 \end{aligned}$$

The cone of tangents and its polar cone



Note that $\mathcal{A}^T \begin{bmatrix} 1 \\ 0 \end{bmatrix} = I$, and so we take $\hat{y} = \bar{y} + \begin{bmatrix} 5 \\ 0 \end{bmatrix}$.

Extreme rays of polar cone C^* give the constraints

$$\begin{aligned} y_1 &\geq 4 \\ 5y_1 - 4y_2 &\geq 12 \end{aligned}$$

which bound the cone of tangents.

More on the polar cone

- Other vectors in the nullspace of \bar{S} will not give such strong constraints.
- In higher dimensions, the **polar cone may not be polyhedral**, and the number of extreme rays may be uncountable.
- If the columns of the matrix D form an eigenbasis for the minimum eigenvalue, then the constraint $D^T(C - \mathcal{A}^T y)D \succeq 0$ **represents exactly the cone of tangents** at \hat{y} .

The form of the of the polar cone

- In general, the polar cone C^* at a boundary point y is equal to the closure of the set of vectors of the form

$$\tilde{C}^* := \{\mathcal{A}(PZP^T) : Z \succeq 0\}$$

where the columns of P give a basis for the nullspace of S .

- The set \tilde{C}^* may not be closed.

Example (Pataki):

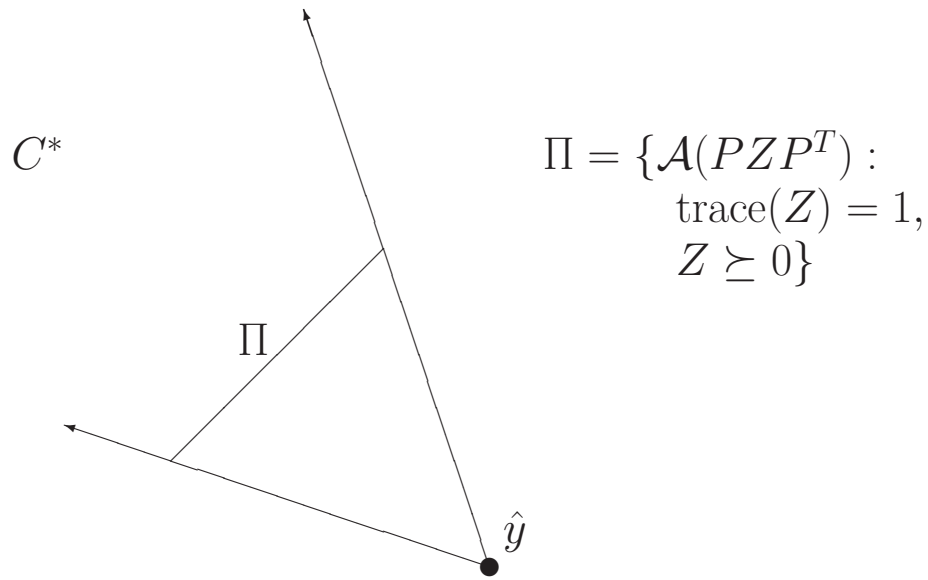
$$P^T A_1 P = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \quad P^T A_2 P = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

The ray $(0, 1)^T$ is not in the cone $\{\mathcal{A}(PZP^T) : Z \succeq 0\}$, but it is in its closure.

(Consider $Z = \begin{bmatrix} \epsilon & 0.5 \\ 0.5 & 1/\epsilon \end{bmatrix}$.)

- Under our first assumption, the set \tilde{C}^* is closed.

Finding the extreme rays of the polar cone



- The extreme rays of C^* correspond to the extreme points of the compact set Π .
- Find these by minimizing a linear objective function over Π . This is equivalent to finding an eigenvector of minimum eigenvalue of a particular matrix.

6 Complexity of the polyhedral cutting plane algorithm

- Theoretically, could use a **volumetric barrier cutting plane algorithm** to solve (SDD).
- Want to get within ϵ of optimality.
- Requires $O(m \log(\frac{1}{\epsilon}))$ **calls to the quadratic subproblem** and a similar number of Newton steps.
- Each **solution of the subproblem** requires $O(n^3)$ work if the QR algorithm is used to find the solutions to the subproblems. It also requires $O(mn^2)$ work to calculate \bar{S} .
- Each **Newton step** in (LDR) requires $O(m^3)$ work.
- Thus, for $m \geq n$, **total number of arithmetic operations** is of the order of $O(m^4 \log(\frac{1}{\epsilon}))$.
- A **primal dual SDP method** requires $O(mn^3 + m^2n^2 + m^3)$ arithmetic operations at each iteration and $O(\sqrt{n} \log(\frac{1}{\epsilon}))$ iterations. So overall complexity is $O(m^2n^{2.5} \log(\frac{1}{\epsilon}))$ for $n \leq m \leq n^2$. (Complexity can be reduced slightly if constraint matrices have special structure.)

7 Conclusions

1. The general framework allows the design of **many different algorithms**.
2. Cutting plane approaches to the solution of semidefinite programs have attractive **theoretical complexity**.
3. High-dimensional polyhedral cutting planes can be found efficiently from the extreme rays of the polar cone.
4. The **cone of tangents** can be captured exactly through the use of a nonpolyhedral cutting plane.
5. Assumption 1 is necessary to allow the straightforward formulation of an eigenvalue optimization problem equivalent to (SDD) , and to allow the translation of any infeasible point to a point on the boundary of Y while preserving the relevance of the eigenbasis for the minimum eigenvalue. The assumption is equivalent to assuming the feasible region of (SDP) is bounded.

References

- [1] H. Y. Benson and R. J. Vanderbei. Solving problems with semidefinite and related constraints using interior-point methods for nonlinear programming. *Mathematical Programming*, 95(2):279–302, 2003.
- [2] S. Burer and R. D. C. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Mathematical Programming*, 95(2):329–357, 2003.
- [3] S. Burer, R. D. C. Monteiro, and Y. Zhang. Solving a class of semidefinite programs via nonlinear programming. *Mathematical Programming*, 93(1):97–122, 2002.
- [4] C. Helmberg. Semidefinite programming for combinatorial optimization. Technical Report ZR-00-34, TU Berlin, Konrad-Zuse-Zentrum, Berlin, October 2000. Habilitationsschrift.
- [5] C. Helmberg and K. Kiwiel. A spectral bundle method with bounds. *Mathematical Programming*, 93(2):173–194, 2002.
- [6] C. Helmberg and F. Rendl. A spectral bundle method for semidefinite programming. *SIAM Journal on Optimization*, 10(3):673–696, 2000.
- [7] K. Krishnan and J. E. Mitchell. A unifying framework for several cutting plane methods for semidefinite programming. Technical report, Mathematical Sciences, Rensselaer Polytechnic Institute, Troy, NY 12180, November 2002. Revised: December 15, 2003.
- [8] K. Krishnan and J. E. Mitchell. Properties of a cutting plane method for semidefinite programming. Technical report, Mathematical Sciences, Rensselaer Polytechnic Institute, Troy, NY 12180, May 2003.
- [9] K. Krishnan and J. E. Mitchell. Semi-infinite linear programming approaches to semidefinite programming (SDP) problems. In P. M. Pardalos and H. Wolkowicz, editors, *Novel approaches to hard discrete optimization problems*, volume 37 of *Fields Institute Communications Series*, pages 123–142. AMS, 2003.
- [10] M. R. Oskoorouchi and J.-L. Goffin. The analytic center cutting plane method with semidefinite cuts. *SIAM Journal on Optimization*, 13(4):1029–1053, 2003.
- [11] G. Pataki. On the closedness of the linear image of a closed convex cone. Part I: necessary and sufficient conditions for nice cones, and core separation. Technical report, Department of Operations Research, University of North Carolina, Chapel Hill, 2003.