14 Relaxations

In this section we approximate difficult optimization problems by a semi-definite program. Usually the solution of the SDP furnishes a lower bound on the optimal value of the original (minimization) problem. In some cases we can give performance guarantees and recover a suboptimal solution of the original problem from the optimal solution of the SDP.

14.1 Max-Cut

Consider again the Max-Cut problem. We are given a graph G with edge weights w_{ij}, $i,j = 1,\dots,n$. The problem consists in partitioning the vertices $1,\dots,n$ of G into two disjoint subsets S,T such that the weight $\sum_{i\in S,j\in T} w_{ij}$ of the resulting cut is maximized. In contrast to the Min-Cut problem with nonnegative weights this problem is NP-complete.

We represent the partition into subsets S,T by a vector $x \in \{-1,+1\}^n$, for any partition the weight of the cut is then given by the matrix scalar product $\langle A, W\rangle$, where $W = (W_{ij})$ is the matrix of the edge weights and $A = \frac{1}{4}(\mathbf{1} - xx^T) = \frac{1}{4}(\mathbf{1} - X)$. The problem can then be written as

$$\max_{X\in\mathcal{S}_+^n} \frac{1}{4}\langle W, \mathbf{1} - X\rangle \ : \quad \mathrm{diag}(X) = \mathbf{1},\ \mathrm{rk}\ X = 1. \tag{1}$$

This would be a semi-definite program, if there were not the rank constraint. By dropping the rank constraint we obtain the *semi-definite relaxation*

$$\max_{X\in\mathcal{S}_+^n} \frac{1}{4}\langle W, \mathbf{1} - X\rangle \ : \quad \mathrm{diag}(X) = \mathbf{1}. \tag{2}$$

Clearly the optimal value c_{MC}^{opt} of (1) is upper bounded by the optimal value c_{SR}^{opt} of (2), because the feasible set of the latter is overbounding the former.

We shall now employ a *randomized rounding* procedure to generate suboptimal solutions of the original problem from the solution of the relaxed problem, due to Goemans and Williamson [4]. Let X_{SR}^{opt} be the maximizer in (2), and let k be its rank. Then X is the Gramian of n unit norm vectors $f_1,\dots,f_n \in \mathbb{R}^k$, which can be computed as rows of the $n\times k$ matrix F in the factorization $X_{SR}^{opt} = FF^T$. Let now $\xi \in \mathbb{R}^k$ be a uniformly distributed random unit norm vector. For each such ξ, we generate a cut by setting

$$S_\xi = \{i \,|\, \langle \xi, f_i\rangle \geq 0\}, \quad T_\xi = \{i \,|\, \langle \xi, f_i\rangle < 0\}. \tag{3}$$

In other words, the hyperplane which is normal to ξ separates the vectors f_i into two subsets, which will serve as the partition defining the cut.

Let us compute the expectation of the weight c_ξ^{subopt} of this cut. The probability that ξ separates two vectors f_i, f_j depends only on the angle $\theta_{ij} = \arccos\langle f_i, f_j\rangle = \arccos(X_{SR}^{opt})_{ij}$ between f_i and f_j and increases linearly from 0 for $\theta_{ij} = 0$ to 1 for $\theta_{ij} = \pi$. The expectation then equals

$$\mathbb{E}c_\xi^{subopt} = \frac{1}{2\pi}\sum_{i,j=1}^n w_{ij}\theta_{ij} = \frac{1}{2\pi}\sum_{i,j=1}^n w_{ij}\arccos(X_{SR}^{opt})_{ij}.$$

Let us now assume that the weights w_{ij} are nonnegative. We have the inequality

$$\alpha\cdot\frac{1}{4}(1 - (X_{SR}^{opt})_{ij}) \leq \frac{1}{2\pi}\arccos(X_{SR}^{opt})_{ij},$$

where $\alpha = \frac{2}{\pi}\min_{y\in[-1,1)}\frac{\arccos y}{1-y} = \frac{2}{\pi}\min_{\theta\in(0,\pi]}\frac{\theta}{1-\cos\theta} \approx 0.87856$. Multiplying by w_{ij} and summing over i,j, we obtain

$$\alpha\cdot c_{SR}^{opt} \leq \mathbb{E}c_\xi^{subopt} \leq c_{MC}^{opt} \leq c_{SR}^{opt}.$$

Thus the upper bound c_{SR}^{opt} is not larger than $\alpha^{-1} \approx 1.138217$ times the optimal value c_{MC}^{opt} of the original problem. On the other hand, by repeatedly generating random cuts we obtain sub-optimal objective values which are not smaller than $\mathbb{E}c_\xi^{subopt} - \varepsilon$ for ε as small as desired. The expectation is in turn not smaller than α times the optimal objective value.

Relaxation (2) hence has a *guaranteed performance*.

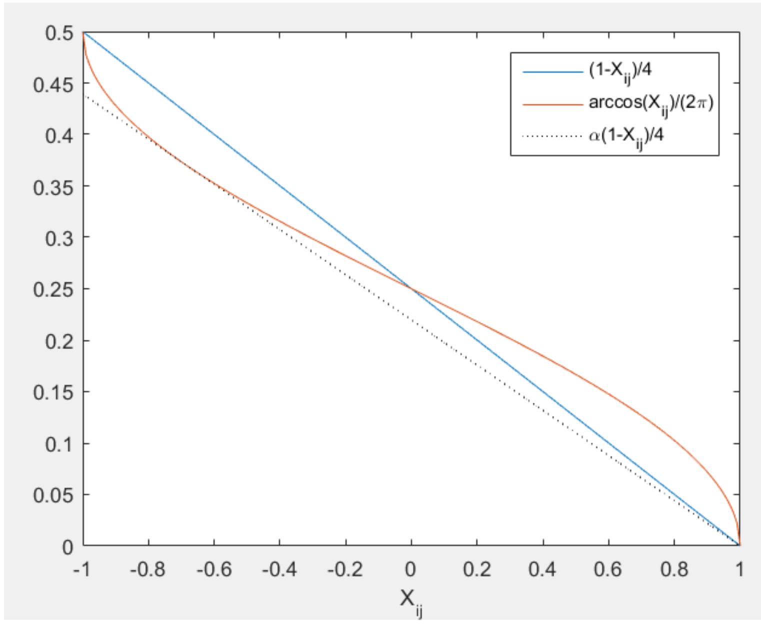

Figure 1: Element-wise functions involved in the Goemans-Williamson procedure

14.2 Nesterov's $\frac{\pi}{2}$ theorem

Let us have another look at the Max-Cut problem. The difference between this problem and its relaxation is that in (1) the matrix variable X takes values at the vertices of the *Max-Cut polytope* $\mathcal{MC}$, defined as the convex hull of the set $\{xx^T \mid x \in \{-1,+1\}^n\}$, and in (2) it takes values in the set

$$\mathcal{SR} = \{X \in \mathcal{S}_+^n \mid \text{diag}(X) = \mathbf{1}\}.$$

In (1) it is actually not relevant whether we maximize over the extreme points of $\mathcal{MC}$ or over the set $\mathcal{MC}$ itself, because the objective function is linear and will assume its maximum at an extreme point anyway.

We shall discard the factor $\frac{1}{4}$ and the constant term $\langle \mathbf{1}, W\rangle$ in the cost function, and consider the two problems

$$\max_{X \in \mathcal{MC}} \langle W, X\rangle \tag{4}$$

and

$$\max_{X \in \mathcal{SR}} \langle W, X\rangle. \tag{5}$$

Problem (5) is a semi-definite relaxation of problem (4), and the optimal value c_{SR}^{opt} of the former is an upper bound for the optimal value c_{MC}^{opt} of the latter, because we have the inclusion $\mathcal{MC} \subset \mathcal{SR}$.

For every $X \in \mathcal{SR}$ we now consider the matrix $\tilde{X} = \frac{2}{\pi} \arcsin X$, where the arcsin function is applied element-wise.

Lemma 14.1. *For every $X \in \mathcal{SR}$ the matrix $\tilde{X}$ is an element of $\mathcal{MC}$.*

Proof. As in the previous section, let k be the rank of X, and let $F \in \mathbb{R}^{n \times k}$ such that $X = FF^T$. Let $f_1, \dots, f_n \in \mathbb{R}^k$ be the rows of F and let $\xi \in \mathbb{R}^k$ be a uniformly distributed random unit length vector. Let $x_\xi \in \{-1,+1\}$ be the vector defined by partition (3), and set $X_\xi = x_\xi x_\xi^T$. Then X_ξ is a random vertex of the Max-Cut polytope $\mathcal{MC}$. Moreover, we have $(X_\xi)_{ij} = -1$ with probability $\frac{\arccos X_{ij}}{\pi}$, and $(X_\xi)_{ij} = 1$ with probability $1 - \frac{\arccos X_{ij}}{\pi}$. Therefore

$$\mathbb{E}(X_\xi)_{ij} = 1 - \frac{2}{\pi} \arccos X_{ij} = \frac{2}{\pi} \arcsin X_{ij}.$$

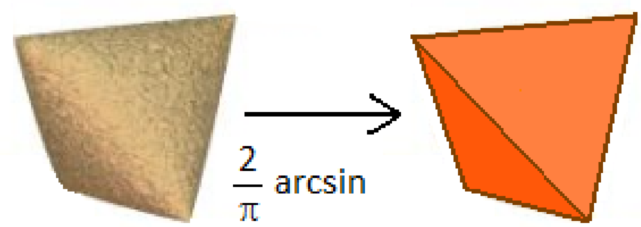

Figure 2: The sets $\mathcal{SR}$ and $\mathcal{MC}$ for $n = 3$

Hence $\mathbb{E}X_\xi = \tilde{X}$, and $\tilde{X}$ is a convex combination of the vertices of $\mathcal{MC}$. The proof is concluded by the convexity of $\mathcal{MC}$. $\square$

The set

$$\mathcal{TA} = \{\frac{2}{\pi} \arcsin X \mid X \in \mathcal{SR}\}$$

is hence an *inner* approximation of the Max-Cut polytope, the so-called *trigonometric approximation* [5]. For $n \leq 3$ this relaxation is exact, i.e., $\mathcal{TA} = \mathcal{MC}$.

The function arcsin has the Taylor series $\arcsin z = \sum_{k=0}^{\infty} \frac{(2k-1)!!\, z^{2k+1}}{(2k)!!\,(2k+1)}$, where the double factorial means a product of only odd or only even integers. It follows that the difference $\frac{2}{\pi} \arcsin z - \frac{2}{\pi} z$ has only nonnegative Taylor coefficients and converges for all $|z| < 1$. As a consequence, for every matrix $X \in \mathcal{SR}$ we have $\tilde{X} = \frac{2}{\pi} \arcsin X \succeq \frac{2}{\pi} X$.

Suppose now that the weighting matrix W is itself positive semi-definite. Setting X equal to the optimal solution X_{SR}^{opt} of the relaxation (5) and taking the matrix scalar product of the preceding inequality with W, we obtain that

$$\mathbb{E}\langle W, X_\xi \rangle = \langle W, \tilde{X} \rangle \geq \frac{2}{\pi} \langle W, X_{SR}^{opt} \rangle = \frac{2}{\pi} c_{SR}^{opt}.$$

We obtain the following performance guarantee of the semi-definite relaxation [8].

Theorem 14.2. *Consider problems* (4),(5) *with a positive semi-definite weighting matrix* W *and let* c_{MC}^{opt}, c_{SR}^{opt} *be their optimal values, respectively. Then*

$$\frac{2}{\pi} c_{SR}^{opt} \leq c_{MC}^{opt} \leq c_{SR}^{opt}.$$

Suboptimal solutions which have at least $(\frac{2}{\pi} - \varepsilon)$ times the optimal value can be obtained by the described randomized rounding procedure.

14.3 Smooth relaxation of MaxCut polytope

Recall that the MaxCut polytope $\mathcal{MC}$ is given by the convex hull of the discrete set

$$X: \qquad X \succeq 0, \quad \operatorname{diag} X = \mathbf{1}, \quad \operatorname{rk} X = 1.$$

It has 2^{n-1} vertices, where n is the matrix size, and many more facets. Optimization over $\mathcal{MC}$ is hence a difficult task.

In addition, due to the discrete nature of the above set continuous optimization methods cannot be applied in a straightforward manner. We therefore consider the following connected analytic relaxation of this set [1].

Instead of dropping the rank constraint, as in the semi-definite relaxation, we weaken it slightly, to obtain the set

$$X: \qquad X \succeq 0, \quad \operatorname{diag} X = \mathbf{1}, \quad \operatorname{rk} X \leq 2. \tag{6}$$

A matrix X of this form can be factorized as $X = FF^T$ with $F \in \mathbb{R}^{n \times 2}$. Since the rows of F are of unit length, they must be of the form $(\cos \varphi_i, \sin \varphi_i)$ for some angles $\varphi_1, \dots, \varphi_n$.

The points of this set are hence parameterized by coordinates living on an n-dimensional torus $\mathbb{T}^n$. Since the torus is a discrete quotient of the space $\mathbb{R}^n$, this makes application of continuous optimization algorithms particularly simple.

14.4 Metric cuts

The basic semi-definite relaxation $\mathcal{SR}$ is an over-bounding approximation of the MaxCut polytope $\mathcal{MC}$. Since $\mathcal{MC}$ is compact, for every point $X \in \mathcal{SR} \setminus \mathcal{MC}$ there exists a hyperplane which separates X strongly from $\mathcal{MC}$. One may then *strengthen* the relaxation by adding linear constraints generated by such separating hyperplanes.

A family of hyperplanes (or linear constraints) which are particularly simple are the so-called *metric inequalities*. Recall that the polytope $\mathcal{MC}$ is defined for every positive integer matrix size n. The different polytopes are related to each other. In particular, if one projects $\mathcal{MC}$ on the subset of entries in a given $k \times k$ principal sub-matrix, the result will be the polytope $\mathcal{MC}$ for matrix size k. For small k the polytopes $\mathcal{MC}$ can be explicitly computed, and are delineated by a reasonable, finite number of hyperplanes (cuts).

In particular, the polytope $\mathcal{MC}$ for matrix size $k = 3$ is given by the 4 linear constraints

$$\begin{pmatrix} 1 & 1 & 1 \\ 1 & -1 & -1 \\ -1 & 1 & -1 \\ -1 & -1 & 1 \end{pmatrix} \begin{pmatrix} X_{12} \\ X_{13} \\ X_{23} \end{pmatrix} \geq \begin{pmatrix} -1 \\ -1 \\ -1 \\ -1 \end{pmatrix}.$$

One may then *impose*, in addition to the conditions defining $\mathcal{SR}$, that the projection on every such sub-space is contained in the corresponding smaller-dimensional $\mathcal{MC}$. In other words, for every triple $\{i, j, k\}$ of mutually distinct indices we add the linear constraints

$$X_{ij} + X_{jk} + X_{ki} \geq -1, \quad X_{ij} - X_{jk} - X_{ki} \geq -1, \quad -X_{ij} + X_{jk} - X_{ki} \geq -1, \quad -X_{ij} - X_{jk} + X_{ki} \geq -1.$$

These are $4 \times \frac{n(n-1)(n-2)}{6}$ conditions, which considerably complicates the relaxation already for moderate values of n.

We may, however, curb the growth in complexity by imposing only subsets of the linear constraints. In particular, we may impose only those constraints which are violated by the solution of the relaxation (2). We may also design an iterative scheme, keeping those constraints at the next step which have been active or violated in the current one.

14.5 Quadratically constrained quadratic programs

We shall now consider a very general class of NP-hard optimization problems, the quadratically constrained quadratic programs (QCQP):

$$\min_{x \in \mathbb{R}^n} (x^T A_0 x + 2b_0^T x + c_0) \ : \quad x^T A_i x + 2b_i^T x + c_i = 0, \ x^T A'_j x + 2{b'}_j^T x + c'_j \leq 0, \ Cx = d. \tag{7}$$

Thus the objective function of the problem is quadratic, and the constraints are linear and quadratic.

Let us introduce the vector $y = (1, x^T)^T \in \mathbb{R}^{n+1}$, whose entries we index from 0 to n, the rank 1 matrix $Y = yy^T$, and the matrices

$$\mathbf{A}_i = \begin{pmatrix} c_i & b_i^T \\ b_i & A_i \end{pmatrix}, \quad \mathbf{A}'_j = \begin{pmatrix} c'_j & {b'}_j^T \\ b'_j & A'_j \end{pmatrix}, \ \mathbf{C} = (-d, C).$$

Then our problem can be rewritten as

$$\min_{Y \in \mathcal{S}_+^{n+1}} \langle \mathbf{A}_0, Y \rangle \ : \quad \langle \mathbf{A}_i, Y \rangle = 0, \ \langle \mathbf{A}'_j, Y \rangle \leq 0, \ \mathbf{C}Y = 0, \ Y_{00} = 1, \ \mathrm{rk}\, Y = 1. \tag{8}$$

By dropping the rank constraint we obtain a semi-definite relaxation of the QCQP. Its optimal value lower bounds the optimal value of the original QCQP.

In general nothing can be said about the performance of the semi-definite relaxation.

14.6 Dines' theorem

We consider the following special cases of QCQPs:

$$\min_{x\in\mathbb{R}^n} x^TAx \ :\quad x^TBx = b, \tag{9}$$

$$\min_{x\in\mathbb{R}^n} x^TAx \ :\quad x^TBx \leq b. \tag{10}$$

Setting $X = xx^T$, we can rewrite these problems as

$$\min_{X\in\mathcal{S}^n_+} \langle A,X\rangle \ :\quad \langle B,X\rangle = b,\ \text{rk}\ X = 1, \tag{11}$$

$$\min_{X\in\mathcal{S}^n_+} \langle A,X\rangle \ :\quad \langle B,X\rangle \leq b,\ \text{rk}\ X = 1. \tag{12}$$

We have the following result.

Lemma 14.3. *The semi-definite relaxations of problems* (9),(10) *which are obtained by dropping the rank constraint from* (11),(12)*, respectively, are exact.*

The proof of the lemma relies on the following theorem by Dines [3].

Theorem 14.4. *Let $A,B\in\mathcal{S}^n$ be real symmetric matrices. Then the set $\{(x^TAx,x^TBx)\in\mathbb{R}^2\,|\,x\in\mathbb{R}^n\}$ is convex.*

The set $\{(x^TAx,x^TBx)\in\mathbb{R}^2\,|\,x\in\mathbb{R}^n\}$ is called the *numerical range* of the pair (A,B). The theorem then says that the numerical range is a convex cone.

Proof. (of Lemma 14.3) If the semi-definite relaxations are infeasible, then the original problems are also infeasible. We shall hence assume that the semi-definite relaxations are feasible.

Let $\hat X\in\mathcal{S}^n_+$ be a feasible point, and set $\hat b = \langle B,\hat X\rangle$, $\hat a = \langle A,\hat X\rangle$. Let $k = \text{rk}\ \hat X$ and factor $\hat X = FF^T$, where the columns of the matrix $F\in\mathbb{R}^{n\times k}$ are denoted by $f_1,\dots,f_k$. Then $\hat X = \sum_{i=1}^k f_if_i^T$. The points $(f_i^TAf_i, f_i^TBf_i)$ are in the numerical range of (A,B). Since the numerical range is a convex cone, the conic combination

$$\sum_{i=1}^k (f_i^TAf_i, f_i^TBf_i) = (\langle A,\hat X\rangle, \langle B,\hat X\rangle) = (\hat a,\hat b)$$

is also in the numerical range. Hence there exists a vector $\hat x\in\mathbb{R}^n$ such that $(\hat x^TA\hat x, \hat x^TB\hat x) = (\hat a,\hat b)$. Therefore problems (11),(12) are equivalent to their counterparts without the rank constraint. □

14.7 Stability number

In this section we consider another graph-related optimization problem.

Definition 14.5. A *clique* of a graph G is a subset S of vertices of G such that every two vertices from S are connected by an edge. A *maximal* clique is a clique which cannot be enlarged by adding another vertex. The *clique number* $\alpha(G)$ of a graph G is the cardinality of its largest clique.

Complementary to this definition is the following.

Definition 14.6. A *stable set* of a graph G is a subset S of vertices of G such that no two vertices from S are connected by an edge. A *maximal* stable set is a stable set which cannot be enlarged by adding another vertex. The *stability number* of a graph G is the cardinality of its largest stable set.

Thus a stable set of G is a clique of the complementary graph $\bar G$ and vice versa. In particular, computing the clique number and computing the stability number of graphs are problems of the same complexity.

There exists a semi-definite relaxation for computing the clique number of a graph. Its value is called the *Lovasz ϑ-function*, computed by solving the SDP [6]

$$\max_{X\succeq 0} \langle X,\mathbf{1}\rangle :\ X\bullet A_{\bar G} = 0,\ \text{tr}\ X = 1$$

or its dual

$$\min \lambda_{\max}(Y + \mathbf{1}) : \ Y \bullet A_G = 0, \ \operatorname{diag} Y = 0.$$

Here A_G is the incidence matrix of the graph, i.e., $A_{ij} = 0$ if vertices i, j are not linked by an edge, and $A_{ij} = 1$ if they are linked by an edge. The operation $\bullet$ is the Hadamard (element-wise) multiplication. The diagonal of A_G is always assumed to be zero.

Let us show that $\alpha(G) \leq \vartheta(G)$. Let $S \subset V$ be the largest clique of the graph G, and $k = \alpha(G)$ its cardinality. Define the matrix $X = (X_{ij})$ by

$$X_{ij} = \begin{cases} \frac{1}{k}, & i, j \in S, \\ 0, & \{i, j\} \not\subset S. \end{cases}$$

Then we have

$$\operatorname{tr} X = 1, \quad X \succeq 0, \quad X \bullet A_{\bar{G}} = 0, \quad \langle X, \mathbf{1} \rangle = k$$

and consequently $k \leq \vartheta(G)$.

The problem of computing the clique (or stability) number can actually be directly formulated as a conic program, albeit over a cone which is computationally difficult to access.

Definition 14.7. A matrix $A \in \mathcal{S}^n$ is called *copositive* if $x^T A x \geq 0$ for all $x \in \mathbb{R}^n_+$. The set of copositive matrices forms the *copositive cone* $\mathcal{COP}^n$.

The clique number of a graph G can be represented as the optimal value of the conic program [7]

$$\min_{Z \in \mathcal{COP}^n} \alpha : \quad Z = \alpha(I + A_{\bar{G}}) - \mathbf{1} = (\alpha - 1)\mathbf{1} - \alpha A_G \tag{13}$$

over the copositive cone. Such programs are called *copositive programs*. Here, as usual, $\mathbf{1}$ is the $n \times n$ matrix with all elements equal to 1.

Now note that if a real symmetric matrix can be represented as a sum $A = P + N$, where $P \in \mathcal{S}^n_+$ and $N \geq 0$, then $A \in \mathcal{COP}^n$. Hence the semi-definite representable cone

$$\mathcal{SPN}^n = \mathcal{S}^n_+ + \mathcal{N}^n,$$

where $\mathcal{N}^n$ is the cone of element-wise nonnegative matrices, is an inner approximation of the copositive cone.

We have the following result [2]:

Theorem 14.8. *For $n \leq 4$ equality $\mathcal{COP}^n = \mathcal{SPN}^n$ holds. For $n \geq 5$ the cone $\mathcal{SPN}$ is strictly contained in $\mathcal{COP}^n$.*

Hence $\alpha(G)$ can be bounded from above by the value of the semi-definite relaxation

$$\min_{Z \succeq 0} \lambda : \quad Z \leq \lambda(I + A_{\bar{G}}) - \mathbf{1} = (\lambda - 1)\mathbf{1} - \lambda A_G.$$

It can be obtained from the exact reformulation (13) by replacing $\mathcal{COP}^n$ by $\mathcal{S}^n_+ + \mathcal{N}^n$.

This relaxation is more complex, but also stronger than the Lovasz ϑ-function.

References

[1] Samuel Burer. Rank-two relaxation heuristics for max-cut and other binary quadratic programs. *SIAM J. Optim.*, 12(2):503–521, 2002.

[2] P. H. Diananda. On nonnegative forms in real variables some or all of which are nonnegative. *Proc. Cambridge Philos. Soc.*, 58:17–25, 1962.

[3] Lloyd L. Dines. On the mapping of quadratic forms. *Bull. Amer. Math. Soc.*, 47(6):494–498, 1941.

[4] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42(6):1115–1145, Nov 1995.

[5] Bernd Hirschfeld. *Approximative Lösungen des Max-Cut Problems mit semi-definiten Programmen.* PhD thesis, Universität Düsseldorf, 2004.

[6] L. Lovasz. On the Shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25(1):1–7, 1979.

[7] T. Motzkin and E. Straus. Maxima for graphs and a new proof of a theorem of turán. *Canadian J. Math.*, 17:533–540, 1965.

[8] Yurii Nesterov. Semidefinite relaxation and nonconvex quadratic optimization. *Optim. Methods Softw.*, 9(1–3):141–160, 1998.