European Journal of Operational Research 216 (2012) 509–520

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Invited Review

Copositive optimization – Recent developments and applications

Immanuel M. Bomze *

Department of Statistics and Operations Research, University of Vienna, Austria

ARTICLE INFO

Article history:
Received 26 January 2011
Accepted 26 April 2011
Available online 30 April 2011

Keywords:
Clique number
Completely positive matrix
Convexity gap
Crossing number
Robust optimization
Standard quadratic optimization

ABSTRACT

Due to its versatility, copositive optimization receives increasing interest in the Operational Research community, and is a rapidly expanding and fertile field of research. It is a special case of conic optimization, which consists of minimizing a linear function over a cone subject to linear constraints. The diversity of copositive formulations in different domains of optimization is impressive, since problem classes both in the continuous and discrete world, as well as both deterministic and stochastic models are covered. Copositivity appears in local and global optimality conditions for quadratic optimization, but can also yield tighter bounds for NP-hard combinatorial optimization problems. Here some of the recent success stories are told, along with principles, algorithms and applications.

1. Introduction

1.1. Motivation, notation and basic ideas

Copositive optimization (or copositive programming, coined in [19]) is a special case of conic optimization, which consists of minimizing a linear function over a (convex) cone subject to additional (inhomogeneous) linear (inequality or equality) constraints. This problem class has a close connection to that of quadratic optimization, which represents the simplest class of hard problems in continuous optimization [102] – to minimize a (possibly indefinite) quadratic form over a polyhedron given in standard form:

$$\min\{\mathbf{x}^\top Q\mathbf{x} : A\mathbf{x} = \mathbf{b}, \mathbf{x} \in \mathbb{R}^n_+\}. \quad (1)$$

Here we denote by bold-faced letters vectors in n-dimensional Euclidean space $\mathbb{R}^n$, the positive orthant therein by $\mathbb{R}^n_+$ (we write $\mathbf{a} \geqslant \mathbf{b}$ for $\mathbf{a} - \mathbf{b} \in \mathbb{R}^n_+$), and by $^\top$ transposition. I_n is the $n \times n$ identity matrix (sometimes with subscript suppressed if the order of I_n is clear from the context), $\mathbf{o}$ and O stand for zero vectors, and matrices, respectively, of appropriate orders. For two integers m and n with $m \leqslant n$ we abbreviate $[m : n]$ for the integer interval $\{m, m+1, \ldots, n\}$.

The basic *lifting* idea (see, e.g. [92]) is to linearize the quadratic form

$$\mathbf{x}^\top Q\mathbf{x} = \mathrm{trace}(\mathbf{x}^\top Q\mathbf{x}) = \mathrm{trace}(Q\mathbf{x}\mathbf{x}^\top) = \langle Q, \mathbf{x}\mathbf{x}^\top\rangle$$

by introducing the new symmetric matrix variable $X = \mathbf{x}\mathbf{x}^\top$ and Frobenius duality $\langle X, Y\rangle = \mathrm{trace}(XY)$. If $A\mathbf{x} \in \mathbb{R}^m_+$ for all $\mathbf{x} \in \mathbb{R}^n_+$ and $\mathbf{b} \in \mathbb{R}^m_+$, then the linear constraints in (1) can be squared, to arrive in a similar way at linear constraints of the form $\langle A_i, X\rangle = b_i^2$, where $A_i = \mathbf{a}_i\mathbf{a}_i^\top$ and $\mathbf{a}_i^\top$ is the ith row of A.

Now the set of all these $X = \mathbf{x}\mathbf{x}^\top$ generated by feasible $\mathbf{x}$ is nonconvex since rank $(\mathbf{x}\mathbf{x}^\top) = 1$. The convex hull

$$\mathcal{C} = \mathrm{conv}\{\mathbf{x}\mathbf{x}^\top : \mathbf{x} \in \mathbb{R}^n_+\},$$

results in a convex matrix cone called the cone of *completely positive matrices* since [71]; for a text see [7]. Note that a similar construction dropping nonnegativity constraints leads to

$$\mathcal{P} = \mathrm{conv}\{\mathbf{x}\mathbf{x}^\top : \mathbf{x} \in \mathbb{R}^n\},$$

the cone of positive-semidefinite matrices, the basic set in *Semidefinite Optimization* (*SDP*), wherefrom above lifting idea was borrowed.

1.2. Terminology, duality and attainability

Duality theory for conic optimization problems requires the dual cone $\mathcal{C}^*$ of $\mathcal{C}$ w.r.t. the Frobenius inner product which is

$$\mathcal{C}^* = \{S \text{ a symmetric } n \times n \text{ matrix} : \langle S, X\rangle \geqslant 0 \text{ for all } X \in \mathcal{C}\}.$$

Here it can easily be shown that $\mathcal{C}^*$ coincides with the cone of *copositive matrices*, namely

$$\mathcal{C}^* = \{S \text{ a symmetric } n \times n \text{ matrix} : \mathbf{x}^\top S\mathbf{x} \geqslant 0 \text{ for all } \mathbf{x} \in \mathbb{R}^n_+\}.$$

This observation justifies terminology of our problem class. The term was coined by Motzkin (the usually cited source [106] however provides no evidence of this) who called a matrix S copositive (apparently abbreviating "conditionally *positive*-semidefinite"), if S

* Tel.: +43 1427738652.
E-mail address: bomze.ejor@univie.ac.at

0377-2217/$ - see front matter © 2011 Published by Elsevier B.V.
doi:10.1016/j.ejor.2011.04.026

generates a quadratic form $\mathbf{x}^\top S\mathbf{x}$ taking no negative values over the positive orthant. More generally, let $\Gamma \subseteq \mathbb{R}^n$ be a closed convex cone and consider the class

$$\mathcal{C}^*_\Gamma = \{S \text{ a symmetric } n \times n \text{ matrix} : \mathbf{x}^\top S\mathbf{x} \geqslant 0 \text{ for all } \mathbf{x} \in \Gamma\}$$

of all *Γ-copositive matrices*. This is the dual cone of

$$\mathcal{C}_\Gamma = \text{conv}\{\mathbf{x}\mathbf{x}^\top : \mathbf{x} \in \Gamma\}.$$

The first accounts on copositive optimization can be found in [116,19], where a copositive representation of a subclass of particular interest is established, namely for *Standard Quadratic Optimization Problems (StQPs)*. Here the feasible polyhedron is the standard simplex $\Delta = \{\mathbf{x} \in \mathbb{R}^n_+ : \mathbf{e}^\top\mathbf{x} = 1\}$, where $\mathbf{e} = [1,\ldots,1]^\top \in \mathbb{R}^n$: this subclass is also NP-hard (there can be up to $\sim 2^n/(1.25\sqrt{n})$ local non-global solutions [14]). Now, with $E = \mathbf{e}\mathbf{e}^\top$ the $n \times n$ all-ones matrix, we have

$$\min\{\mathbf{x}^\top Q\mathbf{x} : \mathbf{x} \in \Delta\} = \min\{\langle Q, X\rangle : \langle E, X\rangle = 1,\ X \in \mathcal{C}\}. \quad (2)$$

Note that the problem on the right-hand side is convex, so there are no more local, non-global solutions. In addition, the objective function is now linear, and there is just one linear equality constraint. The complexity has been completely pushed into the feasibility condition $X \in \mathcal{C}$, which also shows that there are indeed convex minimization problems which cannot be solved easily: while the most prominent conic optimization problems, namely SDPs, second-order cone optimization, and linear optimization problems (LPs), can be solved to arbitrary accuracy in polynomial time, copositive problems are NP-hard. The dual of problem (2) over $\mathcal{C}$ is then

$$\max\{y : S = Q - yE \in \mathcal{C}^*\}, \quad (3)$$

a linear objective in just one variable y with the innocent-looking feasibility constraint $S \in \mathcal{C}^*$. This shows that checking membership of $\mathcal{C}^*$ is NP-hard, which has been observed already by [102]. More generally, a typical primal-dual pair in copositive optimization is of the following form:

$$\inf\{\langle C, X\rangle : \langle A_i, X\rangle = b_i,\ i \in [1:m],\ X \in \mathcal{C}\}$$
$$\geqslant \sup\left\{\mathbf{b}^\top\mathbf{y} : \mathbf{y} \in \mathbb{R}^m,\ S = C - \sum_i y_iA_i \in \mathcal{C}^*\right\}.$$

The inequality above is just standard weak duality, but observe we have to use inf and sup since – as in general conic optimization – there may be problems with attainability of either or both problems above, and likewise there could be a (finite or infinite) positive duality gap without any further conditions like strict feasibility (*Slater's condition*). For the above representation of StQPs, this is not the case:

$$min\{\langle Q, X\rangle : \langle E, X\rangle = 1, X \in \mathcal{C}\} = \max\{y : S = Q - yE \in \mathcal{C}^*\}.$$

But for a similar class arising in many applications, the *Multi-Standard Quadratic Optimization Problems* [25], dual attainability is not guaranteed while the duality gap is zero – an intermediate form between weak and strong duality [117]. A complete picture of possible attainability/duality gap constellations in primal-dual pairs of copositive optimization problems is provided in [26], which also lists some elementary algebraic properties and counterexamples illustrating the difference between the semidefinite cone $\mathcal{P}$ and the copositive/completely positive cone $\mathcal{C}^*/\mathcal{C}$. This is important for many copositivity detection procedures, and as we saw in (3), the feasibility constraint incorporates most of the hardness in copositive optimization.

1.3. Surveys, reviews, entries, book chapters

Copositive optimization receives increasing interest in the Operational Research community, and is a rapidly expanding and fertile field of research. While the time may not yet be ripe for writing up the final standard text book in this domain, several authors nonetheless bravely took the challenge of providing an overview, thereby aiming at a rapidly moving target. A recent survey on copositive optimization is offered by [57], while [77,74] provide reviews on copositivity with less emphasis on optimization. Bomze [16] and Busygin [37] provided entries in the most recent edition of the Encyclopedia of Optimization. Recent book chapters with some character of a survey on copositivity from an optimization viewpoint are [17, Section 1.4] and [34]. Finally, [26] offers a rough literature review by clustering a considerable part of copositivity-related publications.

1.4. Organization of this paper

We start in Section 2 by demonstrating the diversity of copositive formulations in different domains of optimization: continuous and discrete, deterministic and stochastic. Section 3 briefly sketches the ideas of approximation hierarchies, a field with many contacts to (semi-)algebraic geometry and positive polynomials, therefore closely related to the Positivstellensatz [120,118,114], an extension of Hilbert's famous Nullstellensatz. Also some complexity issues are discussed here. We turn to the core of Operational Research in discussing the role of copositivity for local and global optimality conditions in Section 4. In the world of quadratic optimization, it turns out that checking global optimality requires an effort which differs from that of checking local optimality only by a factor smaller than the number of constraints. This may be somewhat surprising at first thought. On the other hand, elementary geometric intuition also suggests that the gap between global and local optimization opens more widely when curvature of the objective is no longer constant. In Section 5, we give a short account on some algorithmic approaches to checking copositivity, and to solve copositive optimization problems. Finally, in Section 6, some success stories are reported: how to obtain tractable yet tight bounds for NP-complete combinatorial problems like the Maximum-Clique problem, how to find the best known asymptotic bound for crossing numbers, and, in the continuous domain, how to construct tight convex underestimators by means of copositive optimization, or Lyapunov functions for switched dynamical systems in optimal control.

2. Copositive reformulations of NP-hard problems

2.1. The standard quadratic case

This case was already addressed above as a motivating (and historically first) example for copositive reformulation, see [116,19]. The Maximum-Clique problem provides a thoroughly studied example of application, see Section 6.3. So consider the StQP

$$\alpha_Q = \min\{\mathbf{x}^\top Q\mathbf{x} : \mathbf{x} \in \Delta\}, \quad (4)$$

and its copositive formulation

$$\alpha_Q = \min\{\langle Q, X\rangle : \langle E, X\rangle = 1, X \in \mathcal{C}\}.$$

Not only the optimal values are equal, we also know there is always a rank-one solution X^* to the latter problem over $\mathcal{C}$, which encodes an optimal solution $\mathbf{x}^*$ to the StQP (4) by way of $X^* = \mathbf{x}^*(\mathbf{x}^*)^\top$. However, if there are multiple optimal solutions to the former (or the latter), we only know that any optimal solution $\overline{X}$ (which may be returned by an – ideal – copositive optimization procedure) is a convex combination of rank-one solutions of the type X^* above. Therefore a rounding procedure is required to retrieve the solution of the StQP in general, unless the above addressed procedure 'automatically' delivers a rank-one solution. In any case, a strength of