

Using quadratic cuts to iteratively strengthen convexifications of box quadratic programs

Amélie Lambert¹ and Daniel Porumbel¹

¹Cedric-Cnam, 292, rue Saint-Martin 75141 Paris Cedex 03 France.

Contributing authors: amelie.lambert@cnam.fr;
daniel.porumbel@cnam.fr;

Abstract

We seek the global minimum of a quadratic function f with box constrained variables. For this goal, we underestimate f by a convex piecewise-quadratic function defined as the maximum of $p \geq 1$ convex quadratic functions (p underestimators). We show that when $p \rightarrow \infty$ the optimal solution of this relaxation converges to an optimal solution of the strong “Shor plus RLT” semi-definite relaxation of the initial problem. To compute the new relaxation, we introduce an iterative algorithm that adds convex quadratic cuts (or cutting-quadratics) one by one in a cutting plane fashion. The resulting convexification is tighter than the one produced by previous related methods that use $p = 1$, *i.e.*, using multiple underestimators leads to a stronger convexification than using a unique one (as in past work). Its integration into a spatial branch-and-bound algorithm brings a second advantage: compared to previous work, we can refine the lower bound at each node of the branching tree. This is because we are able to compute underestimators that act specifically on any particular node of the branching tree. Numerical results show that even a small value of $p \in \{2, 3\}$ can often be enough to reduce the branching tree size by half compared to sticking to $p = 1$. The resulting algorithm is also competitive in terms of CPU time compared to well-established solvers that rely on other techniques.

Keywords: Quadratic Programming, piecewise-quadratic underestimator, cutting-quadratics algorithm

1 Introduction and literature review

Our goal is to find the exact solution of the following non-convex Quadratic Box-constrained Program (QBP):

$$(QBP) \begin{cases} \min f(x) \equiv \langle Q, xx^\top \rangle + c^\top x \\ \ell_i \leq x_i \leq u_i \end{cases} \quad \forall i \in \mathcal{I} \quad (1)$$

where $\langle A, B \rangle = \sum_{i=1}^n \sum_{j=1}^n a_{ij} b_{ij}$, $\mathcal{I} = \{1, \dots, n\}$, $(Q, c, \ell, u) \in \mathcal{S}_n \times \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^n$, and \mathcal{S}_n is the set of real symmetric matrices of order n . Without loss of generality, we assume that the box constraints (1) take the form $x_i \in [0, 1]$ and that the feasible domain of (QBP) is non-empty.

(QBP) is a fundamental \mathcal{NP} -hard global optimization problem [1]. It can be used to formulate numerous applications (e.g., the elastic-plastic torsion problem [2], the molecular conformation analysis [3], the optimal design problems [4]), or can be a sub-problem of more complex problems in various areas (see for instance [5, 6]). Although (1) is the simplest non-convex quadratic optimization program, finding its global optimum remains very challenging even for medium sized instances. Since f is not a convex function, many local optima may not be global; standard approaches for solving (QBP) to global optimality include spatial branch-and-bound algorithms [7–9] combined with convex relaxations to determine lower bounds (see for instance [10, 11]). These convex relaxations are typically either linear, quadratic convex or semi-definite. The solution space is partitioned by branching along the spatial branch-and-bound execution, so as to tighten the convex relaxation as the nodes cover a smaller and smaller feasible area. This family of methods also apply to more general classes of quadratic problems involving quadratic constraints or mixed-integer variables (see for instance [11–14]).

1.1 Main relaxations from the literature

1.1.1 Linear and semidefinite relaxations

Many relaxations of (QBP) express the quadratic function in an extended space of variables, introducing new variables Y_{ij} that are meant to satisfy $Y_{ij} = x_i x_j$, for all $(i, j) \in \mathcal{I}^2$, where \mathcal{I}^2 is the cartesian product of set \mathcal{I} . This standard approach was first used for linearizing f , obtaining the following reformulation of (QBP):

$$(LP) \begin{cases} \min f_L(x, Y) \equiv \langle Q, Y \rangle + c^\top x \\ \ell_i \leq x_i \leq u_i \\ Y = xx^\top \end{cases} \quad \forall i \in \mathcal{I} \quad (2a)$$

$$(2b)$$

Problem (LP) is equivalent to problem (QBP), since when (x^*, Y^*) is a feasible solution to (LP), x^* is also feasible for (QBP), and both problems have the same objective value (i.e. $f_L(x^*, Y^*) = f(x^*)$). Thus, (QBP) can be solved by a spatial branch-and-bound based on the linear relaxation obtained by relaxing the non-convex feasible set

(2a)–(2b) by its convex hull. Since it may be hard to completely describe this convex hull (see [15–17]), it is preferable to construct an outer approximation based on the McCormick’s envelopes [18] captured by the following set \mathcal{M} :

$$\mathcal{M} := (x, Y) \in \mathbb{R}^n \times \mathcal{S}_n : \begin{cases} Y_{ij} \leq u_j x_i + \ell_i x_j - \ell_i u_j & (i, j) \in \mathcal{I}^2 & (3a) \\ Y_{ij} \leq \ell_j x_i + u_i x_j - u_i \ell_j & (i, j) \in \mathcal{I}^2 & (3b) \\ Y_{ij} \geq u_j x_i + u_i x_j - u_i u_j & (i, j) \in \mathcal{I}^2 & (3c) \\ Y_{ij} \geq \ell_j x_i + \ell_i x_j - \ell_i \ell_j & (i, j) \in \mathcal{I}^2 & (3d) \\ \ell_i \leq x_i \leq u_i & \forall i \in \mathcal{I} & (3e) \end{cases}$$

In the resulting linear relaxation called (\overline{LP}) , inequalities (3a)–(3e) involve the lower and upper bounds (ℓ and u) on the original variables x . Then, since the branching rules update the interval $[\ell, u]$ at each node of the branching tree, these inequalities become tighter in the course of the spatial branch-and-bound; this improves the value of the relaxation along the search, as deeper sub-nodes with shorter intervals $[\ell, u]$ are generated. This approach is used by several authors (see for instance [18–20]). In order to tighten (\overline{LP}) , several families of valid linear inequalities were introduced (see for instance [10, 21]), and added to the formulation in a cutting-planes manner. This idea is used by most software implementing the methods described above, see, *e.g.*, Baron ([22]), GloMIQO ([13, 23, 24]), or Gurobi ([25]). Although the evaluation of a linear relaxation is fast, the associated bound is often too weak, and the use of (\overline{LP}) or its extensions in a spatial branch-and-bound may in practice fail to solve medium-sized problems to global optimality.

In order to get tighter relaxations of (QBP) , the use of semi-definite relaxations within branch-and-bound frameworks was also widely studied ([8, 9, 11, 26–28]). A semi-definite relaxation of (QBP) can be obtained by lifting x to a symmetric matrix $X = xx^T$ where these non-convex constraints are relaxed into $X - xx^T \succeq 0$ – notation $M \succeq 0$ means that M is positive semidefinite. By using the Schur complement, $X - xx^T \succeq 0$ is equivalent to (4e) below. After linking variables X and x with the McCormick constraints (4a)–(4d) and linearizing the objective function, we obtain the following model (introduced in [26]) referred to as the “Shor’s plus RLT” relaxation of (QBP) . The associated bound is very tight, but solving (SDP) in practice may be prohibitively slow even for medium-sized programs; this makes its direct and full integration into a branch and bound framework rather impractical.

$$(SDP) \left\{ \begin{array}{ll} \min & f(X, x) \equiv \langle Q, X \rangle + c^T x \\ & X_{ij} \leq u_j x_i + \ell_i x_j - \ell_i u_j \quad (i, j) \in \mathcal{I}^2 & (4a) \\ & X_{ij} \leq \ell_j x_i + u_i x_j - u_i \ell_j \quad (i, j) \in \mathcal{I}^2 & (4b) \\ & X_{ij} \geq u_j x_i + u_i x_j - u_i u_j \quad (i, j) \in \mathcal{I}^2 & (4c) \\ & X_{ij} \geq \ell_j x_i + \ell_i x_j - \ell_i \ell_j \quad (i, j) \in \mathcal{I}^2 & (4d) \\ & \begin{pmatrix} 1 & x^T \\ x & X \end{pmatrix} \succeq 0 & (4e) \\ & x \in \mathbb{R}^n \quad X \in \mathcal{S}_n & (4f) \end{array} \right.$$

The idea of strengthening such an SDP relaxation using cuts comes with a price that has been already discussed in the literature: if the SDP program integrates too many cuts, the computational cost of solving it may become prohibitively high. Thus, it is important to determine the most appropriate and effective cuts, while limiting their number. Different ideas were implemented in this direction, *e.g.*, the polarity cuts and the lift-and-project methodology from [29] or the polynomial separation approach from [30]. We do share the idea that we should find only the most appropriate cuts and to limit their number, but our context is not exactly the same as in above work because we will actually strengthen a convex quadratic program and not an SDP one.

1.1.2 Convex quadratic relaxations

A third family of classical approaches relies on convex quadratic relaxations, including in particular methods MIQCR (Mixed Integer Quadratic Convex Reformulation) and extensions ([31–33]). In these approaches, a quadratic convex relaxation of (QBP) , referred to as $(\overline{P}_{S_0^*})$ and formally defined below, is calculated by exploiting the above (SDP) relaxation. The strength of this approach is that the optimal value of $(\overline{P}_{S_0^*})$ is equal to that of (SDP) . In fact, any SDP matrix $S_0 \succeq \mathbf{0}$ can produce a relaxation (\overline{P}_{S_0}) , but the optimal one, in the sense that it leads to the tighter lower bound, is obtained by constructing matrix S_0^* from the optimal dual solution of (SDP) . **Hence, one needs** to solve (SDP) only once for this very purpose.

Once the optimal S_0^* is determined, the original problem is then solved by a branch-and-bound based on the relaxation $(\overline{P}_{S_0^*})$. The problem is also expressed in the extended (x, Y) space; the reformulated objective can be seen as a quadratic surface (quadric) described by:

$$f_{S_0}(x, Y) = \langle S_0, xx^\top \rangle + c^\top x + \langle Q - S_0, Y \rangle, \quad (5)$$

where $S_0 \succeq \mathbf{0}$ can be any SDP matrix. It is easy to check that for any $S_0 \succeq \mathbf{0}$, the function f_{S_0} is convex; we also have $f_{S_0}(x, Y) = f(x)$ if $Y = xx^\top$. Note moreover that if S_0 is the positive semi-definite null matrix, we have $f_{S_0}(x, Y) = f_L(x, Y)$, which corresponds to simply linearising f . In other words, the reformulated objective function above includes the linear relaxation as special case (when taking $S_0 = \mathbf{0}$).

Putting all above information together, **we obtain** a model (\overline{P}_{S_0}) **indexed by matrix** $S_0 \succeq \mathbf{0}$, **which** is equivalent to (QBP) **but has a convex objective**.

$$(P_{S_0}) \begin{cases} \min f_{S_0}(x, Y) \equiv \langle S_0, xx^\top \rangle + c^\top x + \langle Q - S_0, Y \rangle & (6a) \\ Y = xx^\top & (6b) \\ \ell_i \leq x_i \leq u_i & i \in \mathcal{I} & (6c) \\ Y \in \mathcal{S}_n & (6d) \end{cases}$$

Then, as for the linear relaxation, the non-convex feasible set can be relaxed with the outer approximation based on the set \mathcal{M} of McCormick envelopes, leading to the following quadratic convex relaxation:

$$(\overline{P}_{S_0}) \begin{cases} \min f_{S_0}(x, Y) \equiv \langle S_0, xx^\top \rangle + c^\top x + \langle Q - S_0, Y \rangle & (7a) \\ (x, Y) \in \mathcal{M} & (7b) \end{cases}$$

One strength of this method is that (SDP) is solved only once at the root node of the spatial branch-and-bound leading to a quadratic convex relaxation with the same value as the optimum of (SDP) . The equivalence between both problems holds at the root node, i.e. for the initial lower and upper bounds ℓ and u only. The classical MIQCR uses the same S_0^* along the spatial branch-and-bound: the same non-SDP convex program $(\bar{P}_{S_0^*})$ is solved at each sub-node of the branching tree, even if the interval $[\ell, u]$ is updated along the execution. This is significantly faster than solving a semi-definite optimization problem at each sub-node of the tree, i.e., for each $[\ell, u]$. However, for a given sub-node (i.e., when ℓ and u are updated), the optimal value of $(P_{S_0^*})$ is no longer equal to that of the (SDP) program associated to the updated ℓ and u .

1.2 Our contributions

Our goal is to design an approach that aims at reaching the value of (SDP) for any interval $[\ell, u]$ corresponding to any sub-node of the branching tree. As mentioned previously, problem $(P_{S_0^*})$ captures the tightness of (SDP) at the root node (i.e. for the initial values of ℓ and u), but not for each sub-node (i.e. with updated values of ℓ and u). We aim at improving the value of $(P_{S_0^*})$ at each sub-node by strengthening it with specifically tailored quadratic cuts, to make the convexification at the local sub-node for the current $[\ell, u]$ tend to the associated value of (SDP) .

For this purpose, our main idea is to replace the unique function f_{S_0} with multiple functions $f_{S_k}, S_k \succeq 0$, with $k = 0, 1, 2, \dots, p$ and p a given integer. We then minimize, over all $(x, Y) \in \mathcal{M}$, the following function:

$$f^*(x, Y) = \max_{k=\{1, \dots, p\}} f_{S_k}(x, Y)$$

Let us focus on Figure 1. Since each function f_{S_k} is quadratic and convex, notice we obtain a piecewise-quadratic convex underestimator f^* . In our main algorithm, the idea is to generate the functions f_k one by one in a cutting-planes fashion, since each function f_{k+1} aims at cutting the current optimal solution (x^k, Y^k) at iteration k by making it sub-optimal. In our new general scheme, the relaxation (P_{S_0}) corresponds to the case $p = 1$. Figure 1 also illustrates how this approach with multiple quadratic cuts ($p > 1$) may generate a tighter more refined convexification than MIQCR ($p = 1$) even without necessarily improving the value of the general lower bound.

The above piece-wise quadratic function f^* is convex, but it is particularly computationally demanding to optimize it repeatedly. Each time f^* integrates a new quadric (at each iteration), we have to call a convex QCQP (quadratically constrained quadratic programming) solver to determine its new optimal solution. This is the *main computational bottleneck* of the overall approach and it is very important to accelerate this solver in practice. We thus propose the following *speed-up technique*: identify the variables that do not change too much from one QCQP call to another and strongly limit their variation at each new QCQP call. The effect is similar to reducing the number of variables. If the resulting optimal solution stays strictly inside the box associated to the imposed variation limit, we considered the speed-up technique was

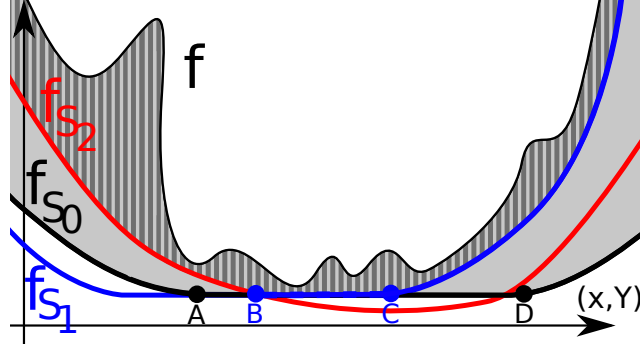


Fig. 1: Convex function f_{S_0} in black may reach its minimum over many solutions of the (x, Y) space (see the long flat segment $[A, D]$), because S_0 may have a large null space, often of dimension close to $\frac{n}{2}$. The resulting convexification in light-gray is weaker than the piecewise-quadratic convexification from the striped area (\max of f_{S_0} , f_{S_1} and f_{S_2}) that has the same minimum; this latter convex function reaches its minimum over a segment $[B, C]$ shorter than $[A, D]$.

successful. Otherwise, the technique failed and we may have remove the box and call the full solver; this is described in Section 4.2.

The road-map of the paper is as follows. Section 2 introduces $(\bar{P}_{\mathcal{K}})$, a parameterized family of piecewise-quadratic and convex relaxations of (QBP) . We then show that for any $[\ell, u]$, the associated value of (SDP) is equal to the optimal value of $(\bar{P}_{\mathcal{K}})$, *i.e.*, it is equal to the best relaxation within this family. Section 3 introduces an iterative **Cutting Quadratics Algorithm (CQA)** that iteratively refines convexification $(\bar{P}_{\mathcal{K}})$ by adding convex quadratic cuts one by one. This algorithm proceeds in cutting-plane fashion, using quadratic hyper-surfaces (or quadratics) instead of hyperplanes to iteratively tighten the relaxation. In Section 4, we integrate Algorithm CQA within a branch-and-bound scheme to solve (QBP) to global optimality. Finally, Section 5 presents experimental results on the *boxqp* instances, suggesting that our new approach is faster than state-of-the-art solvers, and is able to significantly reduce the number of nodes in comparison to the basic MIQCR algorithm.

2 A family of convex piecewise-quadratic relaxations

Given a set $\mathcal{K} = \{S_k \succeq \mathbf{0}, k = 0, 1, \dots, p\}$ of SDP matrices, the multi-cut version of (\bar{P}_{S_0}) from (6a)–(6d) takes the form below, forming a family of equivalent formulations of (QBP) indexed by set \mathcal{K} :

$$(P_{\mathcal{K}}) \begin{cases} \min t & \\ t \geq \langle S_k, xx^\top \rangle + c^\top x + \langle Q - S_k, Y \rangle & S_k \in \mathcal{K} & (8a) \\ Y = xx^\top & & (8b) \\ \ell_i \leq x_i \leq u_i & i \in \mathcal{I} & (8c) \\ Y \in \mathcal{S}_n, t \in \mathbb{R} & & (8d) \end{cases}$$

Like in the mono-cut version (6a)–(6d), the only non-convexity of $(P_{\mathcal{K}})$ comes from constraints (8b) that we can classically relax using the set \mathcal{M} of McCormick envelopes (3a)–(3d). We thus obtain $(\overline{P}_{\mathcal{K}})$ a family of convex relaxations of (QBP) indexed by the same set \mathcal{K} as above:

$$(\overline{P}_{\mathcal{K}}) \begin{cases} \min t & (9a) \\ t \geq \langle S_k, xx^\top \rangle + c^\top x + \langle Q - S_k, Y \rangle, S_k \in \mathcal{K} & (9b) \\ (x, Y) \in \mathcal{M} & (9c) \\ Y \in \mathcal{S}_n, t \in \mathbb{R} \end{cases}$$

Clearly, for any set \mathcal{K} , the problem $(\overline{P}_{\mathcal{K}})$ is a relaxation of (QBP) , since for any solution \bar{x} of (QBP) of value t , the solution $(\bar{x}, \bar{x}\bar{x}^\top, t)$ is feasible for $(\overline{P}_{\mathcal{K}})$ with the same value t . Moreover, since all matrices $S_k \in \mathcal{K}$ are positive semi-definite, each constraint (9a) define a quadratic convex set, and so, $(\overline{P}_{\mathcal{K}})$ is a convex problem. Now, given an integer p , we consider the problem (LB_p) of determining the best set of matrices $\mathcal{K}^* = \{S_0, S_1, \dots, S_p\}$, in the sense of leading to the tightest lower bound of (QBP) :

$$(LB_p) \left\{ \max_{S_0, S_1, \dots, S_p \succeq 0} v(\overline{P}_{\mathcal{K}=\{S_0, S_1, \dots, S_p\}}) \right.$$

where $v(P)$ stands for the optimal value of problem (P) .

This connects our work with previous convex relaxations [31–33], since when we restrict $(\overline{P}_{\mathcal{K}})$ to $p = 1$ (i.e. $\mathcal{K} = \{S_0\}$), we obtain the original MIQCR method from [32]. It is proven in [33], that, the optimal solution of (LB_1) can be derived from the optimal dual solution of (SDP) from (4a)–(4f). It is, moreover, proven in [33] that, if strong duality holds for (SDP) , the optimal value of (LB_1) equals the optimal value of (SDP) which is always the case in a quadratic box-constrained problem. We now state Proposition 1

Proposition 1. $v(LB_\infty) = v(LB_p) = \dots = v(LB_1) = v(SDP)$.

Proof. The last equality (for $p = 1$) follows from [33, Theorem 3.1]. Moreover, by construction, we obviously have $v(LB_\infty) \geq v(LB_p) \geq \dots \geq v(LB_1) = v(SDP)$. We now show $v(LB_\infty) \leq v(SDP)$. Take a feasible solution (\bar{x}, \bar{X}) of (SDP) of objective value $f(\bar{X}, \bar{x}) = \langle Q, \bar{X} \rangle + c^\top \bar{x}$. The solution $(Y = \bar{X}, x = \bar{x})$ is feasible for $(\overline{P}_{\mathcal{K}})$ with a value of

$$\max_{k \in \mathcal{K}} \langle S_k, \bar{x}\bar{x}^\top - \bar{X} \rangle + c^\top \bar{x} + \langle Q, \bar{X} \rangle \leq f(\bar{X}, \bar{x})$$

since $\langle S_k, \bar{x}\bar{x}^\top - \bar{X} \rangle \leq 0$, which is true for any $S_k \succeq 0$ given that $\bar{x}\bar{x}^\top - \bar{X} \preceq 0$ by virtue of (4e). \square

Our idea in the rest of the paper is to use Proposition 1 to reinforce at each sub-node of the spatial branch-and-bound the value of the lower bound, if possible up to reaching the optimum (SDP) . Starting from any initial set \mathcal{K} , we introduce a cutting quadrics algorithm, which adds convex quadratic hyper-surfaces (which amounts to adding positive semi-definite matrices to set \mathcal{K}) at each iteration. We will next present the **Cutting Quadrics Algorithm** (CQA) algorithm as well as its proof of convergence.

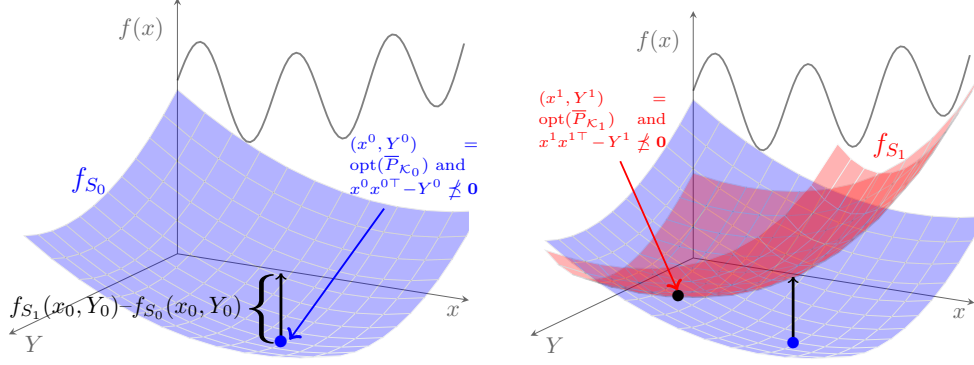


Fig. 2: The Cutting Quadrics Algorithm (CQA) starts from an initial convex function f_{S_0} (blue surface) of problem $(\bar{P}_{\mathcal{K}_0})$ whose optimal solution is (x^0, Y^0) , see the small blue disk. Then it generates matrix S_1 to make point (x^0, Y^0) reach a penalized higher objective value of $f_{S_1}(x^0, Y^0)$, as marked by the black vertical arrow in both figures. Thus, after adding function f_{S_1} (red surface) to $(\bar{P}_{\mathcal{K}_0})$, the optimal solution of $(\bar{P}_{\mathcal{K}_1})$ moves to (x^1, Y^1) .

3 An algorithm for computing a tight quadratic convex relaxation

We now introduce our Cutting Quadrics Algorithm (CQA) that aims at solving (LB_∞) by extending the cutting-planes idea to the case of quadratic convex hyper-surfaces. Our idea is to start from an initial relaxation $(\bar{P}_{\mathcal{K}_0})$ associated to an initial set \mathcal{K}_0 of positive semi-definite matrices. We consider the assumption that the set \mathcal{K}_0 contains at least the null matrix (denoted by $\mathbf{0}_n$), but it can also be composed of any other positive semi-definite matrices like for instance the best matrix S_0^* determined by MIQCR. At each iteration k , CQA calls a convex QCQP (quadratically constrained quadratic programming) solver to determine the optimal solution (x^k, Y^k, t^k) of $(\bar{P}_{\mathcal{K}_k})$, where t^k is its optimal value. It then constructs a matrix $S_{k+1} \succeq 0$ such that solution (x^k, Y^k, t^k) becomes sub-optimal for $(\bar{P}_{\mathcal{K}_{k+1}})$, *i.e.*, for the new program (9a)-(9c) enriched with the following convex quadric:

$$t \geq f_{S_{k+1}} = \langle S_{k+1}, xx^\top \rangle + c^\top x + \langle Q - S_{k+1}, Y \rangle$$

Let us focus on Figure 2 please. Each new iteration $k+1$ aims at cutting solution (x^k, Y^k) when solving $(\bar{P}_{\mathcal{K}_{k+1}})$. We can say (x^k, Y^k) is separated in a cutting-planes fashion, but we use convex quadratic cuts instead of separating hyperplanes.

The key idea is to determine a new matrix S_{k+1} such that the additional quadric penalizes (very heavily) the optimal solution from the previous iteration (x^k, Y^k) . For a given $r^k > 0$, we take $S_{k+1} = r^k \cdot v_{max} v_{max}^\top$, where v_{max} is the eigenvector of matrix $(x^k x^{k\top} - Y^k)$ of maximum eigenvalue (which is positive when $x^k x^{k\top} - Y^k \not\leq \mathbf{0}$). With an appropriate setting of the parameter r^k , this choice ensures that solution (x^k, Y^k) will become sub-optimal in the new program $(\bar{P}_{\mathcal{K}_{k+1}})$, *i.e.*, defining S_{k+1} as above

with a very large r^k will induce a prohibitively large penalty $\langle S_{k+1}, x^k x^{k\top} - Y^k \rangle$ on point (x^k, Y^k) . The choice of matrix S_{k+1} is not unique, and in practice, we can also define $S_{k+1} = r^k \sum v_i v_i^\top$, where the sum is carried out over all eigenvectors v_i having a positive eigenvalue. This will enable the new quadric to penalize larger areas of the (x, Y) space, *i.e.*, more solutions (x, Y) such that $\langle S_{k+1}, x x^\top - Y \rangle > 0$.

An important slowdown of the overall method comes from the fact that we solve each new program $(\bar{P}_{\mathcal{K}_{k+1}})$ from scratch, without being able to re-use any information from the previous program $(\bar{P}_{\mathcal{K}_k})$ constructed before integrating matrix S_{k+1} . In other words, we did not find any solver implementing a reoptimization step to incrementally compute $(\bar{P}_{\mathcal{K}_{k+1}})$ by exploiting the solution of $(\bar{P}_{\mathcal{K}_k})$ – similarly to performing a dual pivot step in the Simplex algorithm after adding a cut to an LP. Any future progress of commercial solvers in this sense may speed-up the proposed ideas. Still, we propose in Section 4.2 a reoptimization feature that proved useful to speed-up the (re-)optimization $(\bar{P}_{\mathcal{K}_{k+1}})$ by limiting the variation of certain variables that did not change in previous iterations.

The overall solution method is summed up in Algorithm 1, and we prove its convergence in Theorem 1. Note that, as previously mentioned, we start with the assumption that set \mathcal{K}_0 will at least integrate the null matrix $\mathbf{0}_n$ (see the very first instruction).

Algorithm 1: The Cutting-Quadrics Algorithm (CQA)

Input : variable bounds ℓ and u , initial penalty parameter r^0 , precision parameter δ , (optional) initial matrices \mathcal{K}_0

Output: The best lower bound on (LB_∞)

$\mathcal{K}_0 \leftarrow \mathcal{K}_0 \cup \{\mathbf{0}_n\}$ // \mathcal{K}_0 may contain the optimal matrix used by MIQCR

$(x^0, Y^0, t^0) \leftarrow \text{Solve}(\bar{P}_{\mathcal{K}_0})$ // variable t^0 is the optimum obj. value

$k \leftarrow 0$

while $(x^k x^{k\top} - Y^k \not\leq 0)$ // In practice we use $\lambda_{\max}(x^k x^{k\top} - Y^k) \geq \delta$

do

$r^k = r^0 + k$ // A penalty parameter;

$v_k \leftarrow$ the eigenvector of $x^k x^{k\top} - Y^k$ of maximum eigenvalue

$S_{k+1} \leftarrow r^k \cdot v_k v_k^\top$ // Or $r^k \cdot \sum v_i v_i^\top$, where the sum is carried

// out over all eigenvectors v_i having a positive eigenvalue

$\mathcal{K}_{k+1} = \mathcal{K}_k \cup S_{k+1}$

$(x^{k+1}, Y^{k+1}, t^{k+1}) \leftarrow \text{Solve}(\bar{P}_{\mathcal{K}_{k+1}})$

$k \leftarrow k + 1$

return t^k as the optimal solution of (QBP) if $Y^k = x^k x^{k\top}$, or as a lower bound otherwise

We now state Theorem 1 ensuring that our algorithm stops when t^k reaches the optimal value of (SDP) .

Theorem 1. *When $k \rightarrow \infty$, the value of the solutions t^k generated by Algorithm 1 converge to the optimal value of (SDP) .*

Proof. In order to prove that **QCA** converges to the optimal value of (SDP) , i.e. that $t^{k^*} = v(SDP)$, we decompose the proof in three steps:

- i) we first show that each intermediate optimal solution (x^k, Y^k, t^k) such that $\lambda_{\max}(x^k x^{k\top} - Y^k) > 0$ will be separated by **CQA**;
- ii) we prove that $\lambda_{\max}(x^k x^{k\top} - Y^k)$ converges to 0 as $k \rightarrow \infty$;
- iii) we finally prove that t^k reaches the optimal value of (SDP) when $k \rightarrow \infty$.

(i) Recall (x^k, Y^k, t^k) is the optimal solution to $(\bar{P}_{\mathcal{K}_k})$ of iteration k of objective value t^k . We first prove that the new quadric $f_{S_{k+1}}$ forces **CQA** to change the current optimal solution (x^k, Y^k, t^k) , or equivalently that the value of $(\bar{P}_{\mathcal{K}_{k+1}})$ at point (x^k, Y^k) is always greater or equal than t^k . Let v_k be the eigenvector of matrix $(x^k x^{k\top} - Y^k)$ of maximum eigenvalue λ_{\max} . By definition, we have $S_{k+1} = r^k \cdot v_k v_k^\top$ and we can develop:

$$\begin{aligned}
f_{S_{k+1}}(x^k, Y^k) - f_{S_k}(x^k, Y^k) &= \langle S_{k+1}, x^k x^{k\top} \rangle + c^\top x^k + \langle Q - S_{k+1}, Y^k \rangle \\
&\quad - \langle S_k, x^k x^{k\top} \rangle - c^\top x^k - \langle Q - S_k, Y^k \rangle \\
&= \langle S_{k+1}, x^k x^{k\top} - Y^k \rangle - \langle S_k, x^k x^{k\top} - Y^k \rangle \\
&= \langle r^k \cdot v_k v_k^\top, x^k x^{k\top} - Y^k \rangle - \langle r^{k-1} \cdot v_{k-1} v_{k-1}^\top, x^k x^{k\top} - Y^k \rangle \\
&= r^k \lambda_{\max} - \langle r^{k-1} \cdot v_{k-1} v_{k-1}^\top, x^k x^{k\top} - Y^k \rangle \\
&> r^k \lambda_{\max} - r^{k-1} \lambda_{\max},
\end{aligned}$$

where the last inequality holds for any S_k of the form $r^{k-1} v_{k-1} v_{k-1}^\top$, based on the following well-known property [34, §3.2]:

$$\lambda_{\max} = \max_{u \in \mathbb{R}^n, \|u\|=1} \langle uu^\top, x^k x^{k\top} - Y^k \rangle$$

This proves $f_{S_{k+1}}(x^k, Y^k) > f_{S_k}(x^k, Y^k) = t^k$, because $r^k > r^{k-1}$. Recall the first line of the **while** loop constructs $r^k = r^{k-1} + 1$ (and after another i iterations r^{k+i} is even higher) so that the penalty imposed on (x^k, Y^k) by the use of S_{k+1} forces **CQA** to move from (x^k, Y^k, t^k) to another solution.

We still need to show that $f_{S_{k+1}}(x^k, Y^k) > f_{S_k}(x^k, Y^k)$ also holds for the very first step, when $k = 0$ and S_k no longer has a form $r^{k-1} v_{k-1} v_{k-1}^\top$. The initial quadrics belong to a set \mathcal{K}_0 provided by the user. The inequality remains true using a high enough r_0 since

$$\begin{aligned}
f_{\mathcal{K}_0}(x^0, Y^0) &= \max_{S_i \in \mathcal{K}_0} \{ \langle S_i, x^0 x^{0\top} \rangle + c^\top x^0 + \langle Q - S_i, Y^0 \rangle \} \\
f_{\mathcal{K}_0}(x^0, Y^0) &< \langle r^0 \cdot v_0 v_0^\top, x^0 x^{0\top} \rangle + c^\top x^0 + \langle Q - r^0 \cdot v_0 v_0^\top, Y^0 \rangle,
\end{aligned} \tag{10}$$

where in the last inequality we used $\langle v_0 v_0^\top, x^0 x^{0\top} - Y^0 \rangle > 0$ which is true because vector v_0 corresponds to the maximum eigenvalue of $x^0 x^{0\top} - Y^0$ which is considered positive here. Still, depending on the initial set \mathcal{K}_0 , we may need a very large initial r^0 (in theory).

(ii) We will now prove that for any $\delta > 0$ no matter how small, there exists some \bar{k} such that $\lambda_{max}(x^k x^{k\top} - Y^k) < \delta$ for any $k \geq \bar{k}$.

Suppose for the sake of contradiction that this is not the case. This means that the algorithm generates an infinite number of matrices $(x^k x^{k\top} - Y^k) \in \mathcal{S}_n$ that all satisfy $\lambda_{max}(x^k x^{k\top} - Y^k) \geq \delta$. The Bolzano-Weierstrass theorem states that any infinite sequence in a bounded set (recall that (x, Y) belong to the unit hypercube) contains a convergent subsequence. This means that there exists a subsequence (k_i) such that $(x^{k_i} x^{k_i\top} - Y^{k_i})$ converges to some fixed point $\bar{x} \bar{x}^\top - \bar{Y}$ when $i \rightarrow \infty$. By assumption, this convergence point satisfies $\lambda_{max}(\bar{x} \bar{x}^\top - \bar{Y}) \geq \delta$.

Let us study how the above subsequence (k_i) stays in a box very close to (\bar{x}, \bar{Y}) . For any infinitesimal ϵ , there exists some k_ϵ such that for any $k_i > k_\epsilon$ in the above subsequence, any element in matrix $(x^{k_i} x^{k_i\top} - Y^{k_i})$ is at a distance smaller than ϵ from the corresponding element of $(\bar{x} \bar{x}^\top - \bar{Y})$, i.e., it stays in a neighborhood N_ϵ of $(\bar{x} \bar{x}^\top - \bar{Y})$ of ∞ -norm below ϵ . Moreover, since λ_{max} is a continuous function, we also have for a sufficiently small ϵ , $\lambda_{max}(x x^\top - Y) > \delta' > 0 \forall (x, Y) \in N_\epsilon$ for some δ' arbitrarily close to δ . All points $k_i > k_\epsilon$ in above subsequence will satisfy $\lambda_{max}(x^{k_i} x^{k_i\top} - Y^{k_i}) > \delta'$.

But here comes the contradiction : for a sufficiently large \hat{k}_i , we have $r^{\hat{k}_i} = r^0 + \hat{k}_i$ so that $f_{S_{\hat{k}_i}}(x, Y)$ will include a penalty of at least $r^{\hat{k}_i} \delta'$ that can become large enough to make any $(x, Y) \in N_\epsilon$ sub-optimal. In other words, at some iteration \hat{k}_i , the weight $r^{\hat{k}_i}$ will become large-enough to cut the whole box, since the associated penalty $r^{\hat{k}_i} \lambda_{max}(x^{\hat{k}_i} x^{\hat{k}_i\top} - Y^{\hat{k}_i}) > r^{\hat{k}_i} \delta'$ is very large. Thus, no iteration after \hat{k}_i of above subsequence can stay in N_ϵ and this is a contradiction.

(iii) We now prove that $t^k \rightarrow v(SDP)$.

Let us first take any convergence point $(\bar{x}, \bar{Y}, \bar{t})$ of Algorithm 1 that has to satisfy $\lambda_{max}(\bar{x} \bar{x}^\top - \bar{Y}) \leq 0$ according to point (ii) above. This means $\bar{x} \bar{x}^\top - \bar{Y} \preceq 0$. We will use this \bar{x} and \bar{Y} to build a feasible solution of (SDP) such that $\bar{t} \geq v(SDP)$. Since $(\bar{x} \bar{x}^\top - \bar{Y}) \preceq 0$, the solution $(x = \bar{x}, Y = \bar{Y})$ is feasible for (SDP) with a value of $c^\top \bar{x} + \langle Q, \bar{Y} \rangle \geq v(SDP)$. The optimal value \bar{t} can be written as below (regardless of how exactly \mathcal{K} was constructed):

$$\bar{t} = \max_{S \in \mathcal{K}} \langle S, \bar{x} \bar{x}^\top - \bar{Y} \rangle + c^\top \bar{x} + \langle Q, \bar{Y} \rangle = c^\top \bar{x} + \langle Q, \bar{Y} \rangle$$

The last equality holds since $\max_{S \in \mathcal{K}} \langle S, \bar{x} \bar{x}^\top - \bar{Y} \rangle = \langle \mathbf{0}_n, \bar{x} \bar{x}^\top - \bar{Y} \rangle = 0$ which is true because $\mathbf{0}_0 \in \mathcal{K}_0$ and because $\langle S, \bar{x} \bar{x}^\top - \bar{Y} \rangle \leq 0$ for any positive semi-definite matrix $S \in \mathcal{K}$.

Since (\bar{x}, \bar{Y}) is feasible for (SDP) , this proves that $\bar{t} \geq v(SDP)$.

We now show $\bar{t} \leq v(SDP)$. Take the optimal solution (\tilde{x}, \tilde{Y}) of (SDP) and notice that it is feasible for $(\bar{P}_{\mathcal{K}_k})$, i.e., for program (9a)-(9c) using any quadrics $\mathcal{K} = \mathcal{K}_k$

constructed by the algorithm. The value of \bar{t} has to be lower or equal to the objective value of (\tilde{x}, \tilde{Y}) in $(\bar{P}_{\mathcal{K}_k})$ which is

$$\max_{S \in \mathcal{K}_k} \langle S, \tilde{x}\tilde{x}^\top - \tilde{Y} \rangle + c^\top \tilde{x} + \langle Q, \tilde{Y} \rangle = \max_{S \in \mathcal{K}_k} \langle S, \tilde{x}\tilde{x}^\top - \tilde{Y} \rangle + v(SDP) \leq v(SDP),$$

where we simply used $\langle S, \tilde{x}\tilde{x}^\top - \tilde{Y} \rangle \leq 0$. This follows from $S \succeq \mathbf{0}$ and $\tilde{x}\tilde{x}^\top - \tilde{Y} \preceq \mathbf{0}$, which holds by virtue of (4e). \square

The proof of Theorem 1 implies the following Corollary.

Corollary 1. *Algorithm 1 solves problem (SDP). In particular, if $(\bar{x}, \bar{Y}, \bar{t})$ is a convergence point of Algorithm 1, (\bar{Y}, \bar{x}) is an optimal solution to (SDP) with a value of $\bar{t} = v(SDP)$.*

Theorem 1 thus states that Algorithm 1 computes a quadratic convex relaxation that reaches the value of (SDP). Within a branch and bound framework, we obtain the same lower bound as MIQCR at the root node but we may obtain improved lower bounds at the subnodes discovered by branching. Moreover, as mentioned previously, the additional quadrics may tighten the convexification by reducing the set of optimal equivalent solutions at each sub-node (like in Figure 1). This increases the potential of the new convexification when integrated within a spatial branch-and-bound algorithm. The numerical experiments will show that the total number of nodes may be halved even if Algorithm 1 only produces very few quadrics (few iterations of the **while** loop).

Recall that Algorithm 1 can start from any set \mathcal{K}_0 . To improve its convergence, we describe hereafter several ways to populate \mathcal{K}_0 , whose size is not limited. The first one consists of adding to \mathcal{K}_0 the optimal matrix calculated by MIQCR, which is $S_0^* = Q + \Phi^1 + \Phi^2 - \Phi^3 - \Phi^4$, where $\Phi^1, \Phi^2, \Phi^3, \Phi^4$ are the symmetric matrices built from the optimal dual variables associated with the McCormick constraints (4a)–(4d) of (SDP). While solving a large SDP problem at the root node is computationally costly, it increases the quality of all resulting bounds and reduces the number of branch-and-bound nodes. Another choice relies on extracting the convex part of Q , by constructing the matrix

$$S_0^+ = \sum \lambda_i v_i v_i^T \quad (11)$$

where the sum is carried over all non-negative eigenvalues λ_i of Q associated to eigenvectors v_i . Preliminary tests show that integrating these SDP matrices in \mathcal{K}_0 is a good choice for all CQA variants.

However, the practical utility of integrating a particular SDP matrix in \mathcal{K}_0 can only be evaluated on an instance by instance basis, by trying and testing it. We here mention two other potentially useful ideas. A very classical way of transforming Q into an easy-to-compute SDP matrix consists of taking $S^{\lambda_{\min}} = Q - \lambda_{\min}(Q) \cdot I_n$, where $\lambda_{\min}(Q)$ is the minimum eigenvalue of Q . A more general approach consists of restoring definiteness [35]: starting from any $\hat{X} \succeq \mathbf{0}$ take the SDP matrix obtained by solving $\max \left\{ t : \hat{X} + t \cdot (Q - \hat{X}) \succeq \mathbf{0} \right\}$; this reduces to finding the intersection between the

segment $[\hat{X}, Q]$ and the boundary of the SDP cone. Observe that, matrix $S^{\lambda_{\min}}$ may have been obtained using this restoring definiteness methodology with $\hat{X} = I_n$.

3.1 A running example of CQA

We consider a toy instance:

$$(P^{EX}) \left\{ \min_{0 \leq x \leq 1} \langle Q, xx^\top \rangle + c^\top x \right.$$

where $c = -(8, 10, 3)^\top$ and $Q = \frac{1}{2} \begin{pmatrix} 5 & 6 & -1 \\ 6 & 4 & -1 \\ -1 & -1 & 7 \end{pmatrix}$.

Running the classical MIQCR

We start by solving (P^{EX}) with algorithm **MIQCR**. For this, we solve the associated SDP relaxation from which we compute matrix S_0^* as described in [33]. Providing only two decimals to lighten the text, this matrix is:

$$S_0^* = \begin{pmatrix} 2.50 & 1.25 & -0.50 \\ 1.26 & 1.26 & -0.84 \\ 0.50 & -0.84 & 3.50 \end{pmatrix}$$

This leads to the following quadratic convex relaxation of (P^{EX}) :

$$(P_{S_0^*}^{EX}) \left\{ \begin{array}{l} \min t \\ t \geq \langle S_0^*, xx^\top \rangle + c^\top x + \langle Q - S_0^*, Y \rangle \\ (x, Y) \in \mathcal{M} \\ Y \in \mathcal{S}_n, t \in \mathbb{R} \end{array} \right.$$

The optimal solution (\bar{Y}, \bar{x}) of $(P_{S_0^*}^{EX})$ is $\bar{x} = (0.53, 1, 0.65)^\top$, $\bar{Y} = \begin{pmatrix} 0.53 & 0.53 & 0.35 \\ 0.53 & 1.00 & 0.65 \\ 0.35 & 0.65 & 0.65 \end{pmatrix}$,

and the objective function value is $t = -9.82357704$.

In such a solution, we have $\max_{i,j} \{\bar{x}_i \bar{x}_j - \bar{Y}_{ij}\} = 0.2491$. Hence, since $\bar{Y} \neq \bar{x} \bar{x}^\top$, we cannot prove the global optimality of (\bar{Y}, \bar{x}) , and we thus have to start a branch-and-bound process to solve (P^{EX}) to global optimality.

Running the new QCA

We consider algorithm **QCA** that starts with three quadratic cuts associated to initial matrices $\mathcal{K}_0 = \{S_0^*, \mathbf{0}_3, S_{\min}^\lambda\}$, where S_0^* is computed above by **MIQCR**, $\mathbf{0}_3$ is the zero matrix, and S_{\min}^λ is constructed as described in the last paragraph above the heading

of Section 3.1. We obtain the following quadratic convex relaxation:

$$(P_{\mathcal{K}_0}^{EX}) \left\{ \begin{array}{l} \min t \\ t \geq \langle S_0^*, xx^\top \rangle + c^\top x + \langle Q - S_0^*, Y \rangle \\ t \geq c^\top x + \langle Q, Y \rangle \\ t \geq \langle S_{\min}^\lambda, xx^\top \rangle + c^\top x + \langle Q - S_{\min}^\lambda, Y \rangle \\ (x, Y) \in \mathcal{M} \\ Y \in \mathcal{S}_n, t \in \mathbb{R} \end{array} \right.$$

Solving $(P_{\mathcal{K}_0}^{EX})$ gives an initial solution (Y^0, x^0) not very different from that of **MIQCR**, and even identical up to a tolerance of 2 decimals, *i.e.* $(Y^0 \cong \bar{Y}, x^0 \cong \bar{x})$, with a slightly different objective value, $t^0 = -9.82357704$.

Contrary to **MIQCR**, **CQA** now starts generating more and more quadratic cuts:
Iteration 1 $\lambda_{\max}(x^0 x^{0\top} - Y^0) = 0.20$, with eigenvector $v = (0.38, 0, -0.92)^\top$. Using $r = 10$, we generate a new quadratic cut associated to matrix:

$$S_1 = \begin{pmatrix} 14.65 & 0 & -35.36 \\ 0 & 0 & 0 \\ -35.36 & 0 & 85.35 \end{pmatrix}$$

This leads to solution (Y^1, x^1) , with $x^1 = x^0$, $Y^1 = \begin{pmatrix} 0.38 & 0.53 & 0.25 \\ 0.53 & 1.00 & 0.65 \\ 0.24 & 0.65 & 0.32 \end{pmatrix}$,

and $t^1 = -9.82352931$.

Iteration 2 $\lambda_{\max}(x^1 x^{1\top} - Y^1) = 0.14$, with eigenvector $v = (0.38, 0, 0.92)^\top$. Using $r = 10$, we generate a new quadratic cut associated to matrix:

$$S_2 = \begin{pmatrix} 14.64 & 0 & 35.35 \\ 0 & 0 & 0 \\ 35.35 & 0 & 85.36 \end{pmatrix}$$

This leads to solution (Y^2, x^2) , with $x^2 = x^1$, $Y^2 = \begin{pmatrix} 0.28 & 0.53 & 0.34 \\ 0.53 & 1.00 & 0.65 \\ 0.34 & 0.65 & 0.42 \end{pmatrix}$,

and $t^2 = -9.82352914$.

Iteration 3 We have $Y^2 - x^2 x^{2\top} = 0$. We thus stop not only because $x^2 x^{2\top} - Y^2 \preceq 0$, but also because we reach the equality $Y^2 = x^2 x^{2\top}$, certifying at the root node of the branching tree that (Y^2, x^2) is an optimal solution of the initial program (P^{EX}) .

This toy example confirms the logic from Figure 1: at the root node of the tree, we do not improve the lower bound (the infinitesimal increase $t^0 < t^1 < t^2$ is a phenomenon arising due to the tolerances of the quadratic solver), but the overall convexification

is tighter. We could numerically verify that the final convex relaxation has a unique optimal solution – identical to that of the original problem (over the x variables).

4 The spatial branch-and-bound to optimally solve (QBP)

We are now interested in computing the optimal solution of (QBP) . A classical way is to use a branch-and-bound algorithm (see [36] for a complete description), where the lower bound is computed at each node using the new **CQA** algorithm. We proved in Section 3 that **CQA** converges to an optimal solution of (SDP) , but each (non-final) iteration of **CQA** does provide a valid lower bound of (QBP) . Moreover, since each **CQA** iteration solves a computationally-expensive convex optimization problem $(\bar{P}_{\mathcal{K}_k})$, we need to find the best trade-off between the tightness of the convexification and its computation time. Thus, in our numerical experiments, we will only consider **CQA** variants that perform relatively few iterations at each branch-and-bound node.

We first provide in Section 4.1 below the general framework of the branch-and-bound. We then present in Section 4.2 two techniques to speed up the solution of $(\bar{P}_{\mathcal{K}})$, because the calls to the convex QCQP solver that optimizes $(\bar{P}_{\mathcal{K}})$ represent the main computational bottleneck of the overall method. Finally, we describe our upper bounding procedure in Section 4.3.

4.1 The dynamics of the branch-and-bound tree construction

We consider a branching tree that contains only fully evaluated nodes, meaning that the lower and upper bounds of each node are known at all times. Each lower bound is determined by calling **CQA** and each upper bound is determined using the heuristic described in Section 4.3. A branching decision is automatically taken at the end of the node evaluation, unless the node is pruned. However, there are two types of nodes:

- A node is closed if the branching decision was implemented and the node produced two evaluated child nodes.
- A node is open or childless if, although evaluated, its branching decision was not yet implemented to produce child nodes.

Each branch-and-bound step performs the following: (i) seek an open node; (ii) implement the branching decisions determined when the parent node was evaluated (at a previous moment) and (iii) evaluate the two resulting child nodes. The selected parent node is thus no longer childless and is marked closed. Its two child nodes are considered open. After evaluating the child nodes, we take a branching decision **right away, to** be implemented later.

4.1.1 How many iterations per node

The evaluation of the lower bound **at each node of the branching tree** is performed from scratch, starting from the same \mathcal{K}_0 . We did try to enable certain nodes to inherit quadratic cuts from some ancestor nodes, but the observed speed-up is not large

enough to warrant complicating the overall branch-and-bound. The most important decision for making **CQA** reach its full potential in practice **concerns** the number of iterations that should be executed **at** each node. More iterations means fewer nodes in the general branching tree, but this comes at the cost of needing more computational work per node.

We thus consider two different stopping criteria for **CQA**. Since the algorithm is supposed to stop when matrix $(x^k x^{k\top} - Y^k) \preceq 0$, we replace this condition with $\lambda_{\max}(x^k x^{k\top} - Y^k) > \delta$. Thus, by varying the value δ , we can choose to make an additional iteration only if matrix $(x^k x^{k\top} - Y^k)$ is far to be **semidefinite negative**, *i.e.*, if its largest eigenvalue can be considered large enough. Our second stopping criteria is simply a bound on the number of iterations of the **while** loop in Algorithm 1.

4.1.2 The branching decision

The variable selection strategy is as follows. Let $\epsilon_b > 0$ be a very small precision parameter for considering an imperfect equality as satisfied. Denoting the solution of $(\bar{P}_{\mathcal{K}_k})$ at the current node by (x^k, Y^k) , two cases are possible:

1. If $-\epsilon_b \leq (Y - xx^\top)_{ij} \leq \epsilon_b$ for all $(i, j) \in \mathcal{I}^2$, then (x, Y) is the optimal solution of the considered branch.
2. Else, we determine i^* by only looking for the diagonal values so that $x_i^2 \neq Y_{ii}$. We first restrict to variables i with a decent range $u_i - \ell_i$ above a given threshold γ fixed to 0.1 in the beginning. We thus solve

$$i^* = \operatorname{argmax}_{u_i - \ell_i > \gamma} |x_i^2 - Y_{ii}|$$

If this scanning finds no satisfactory solution, we divide the above parameter γ by 3. And the process is repeated until we find a satisfactory variable i to be split.

The feasible interval $[\ell_{i^*}, u_{i^*}]$ of the selected variable x_{i^*} is always split in half: the two child nodes are $x_{i^*} \in [\ell_{i^*}, \frac{\ell_{i^*} + u_{i^*}}{2}]$ and $x_{i^*} \in [\frac{\ell_{i^*} + u_{i^*}}{2}, u_{i^*}]$.

Regarding the selection of the next sub-problem to solve, we use the “best-first” strategy, in the sense that we select the node with the highest evaluated lower bound. If this lower bound is higher than the current best-known upper bound, the node is pruned. Otherwise, we implement the branching decision and evaluate the child nodes as described in Section 4.1.

4.2 Using boxes to speed-up the QCQP solver that iteratively optimizes the programs $(\bar{P}_{\mathcal{K}_k})$

We now present a technique that may speed-up the convex solver that is called iteratively at each sub-node to optimize $(\bar{P}_{\mathcal{K}_k})$, because this operation is the main computational bottleneck of the overall method. The solvers we are aware of can not take advantage of the fact that program $(\bar{P}_{\mathcal{K}_{k+1}})$ is simply program $(\bar{P}_{\mathcal{K}_k})$ enriched with a

new quadratic constraint. The Simplex algorithm may simply perform a unique pivoting step when a new linear constraint is added, but – far from such a re-optimization potential – our QCQP solver restarts from scratch when a new constraint is added. Any future research or improvement in QCQP re-optimization may speed-up our algorithm.

Generally, our speed-up idea relies on restricting the allowed variation of certain variables, by putting a box around them. If the optimal solution of the resulting overly-restricted convex program is strictly inside the box, then this is also the optimal solution of the initial relaxation. Otherwise, we remove the artificial box and the convex solver may have to be called a second time.

4.2.1 Boxes in relation to father and ancestor nodes

The optimal solution in a child node is sometimes not fundamentally different from that of the father node, in the sense that many variables stay fixed in both the current and the father nodes. To speed up the first call to the convex solver at each node, we restrict the decision variables as follows. We first identify the variables that never changed in any ancestor node on the genealogical lineage upwards for β levels, where β is a parameter. We first solve the node by strongly limiting the variation of these variables (we almost fix them) compared to their value in the father node. If the optimal solution is within the box we succeed in speeding up this optimization stage. Otherwise, we failed because we have to call the convex solver a second time after removing the box. In many cases, there are multiple optimal solutions. If at least some of them stay inside the box, this technique will succeed.

4.2.2 Boxes acting on the iterations of CQA

While the above technique is applied before starting the **while** loop of Algorithm 1, we now present a technique to be applied inside this **while** loop. We build the following box. First, we compute the maximum variation x^{Max} (or Y^{Max} respectively) over all optimal x (or Y , resp.) values observed during the convex solver calls from previous iterations (including at the QCQP call before the **while** loop). We then limit the variation of variables x and Y within a box of radius x^{Max} and Y^{Max} , respectively. If the obtained optimal solution stays inside the box, we succeeded in speeding-up the QCQP convex solver. Otherwise, we failed and we do not consider the reported solution and node lower bound reliable. In the latter case, we simply do not update the evaluated lower bound of the considered node; yet, we do not call the convex solver a second time.

4.3 A heuristic for computing upper bounds

We use a rather basic coordinate descent heuristic to determine upper bounds. We start from the solution x^k of $(\bar{P}_{\mathcal{K}_k})$ at the current node (ignoring all Y^k components), that is also feasible for the original (QBP). Our coordinate descent begins by iterating over all $i \in [1..n]$; for each i , we fix all variables to their current values except for x_i^k . We then solve an optimization sub-problem with only one decision variable, namely x_i^k . This requires minimizing a simple quadratic function in dimension one that can be determined using elementary mathematics, regardless of its convexity

status. We decided to keep x_i in its range $[l_i, u_i]$, but this could be relaxed. After scanning all variables once, we can repeat the process as long as there is at least one variable i for which we detect a possible improvement.

5 Numerical results

To evaluate our new solution method hereafter called **Cutting-Quadric Branch-and-Bound (CQBB)**, we consider the set of box-constrained quadratic instances labelled *boxqp* and introduced in [8, 9]. We first evaluate the efficiency improvements of CQBB over the original method MIQCR [33], and then compare it with the standard solvers Cplex 22.1 [37], Baron 24.3.19 [22], and Gurobi 11.0.3 [25], as well as with two literature algorithms: QuadProgBB [11] and GloMIQO 2 [24].

Our algorithm versions CQBB-1 and CQBB-2 and their parameter configuration

We use two variants of CQA for computing the lower bound at each node of the branch-and-bound tree, leading to methods CQBB-1 and CQBB-2.

- CQBB-1: the stopping **criterion** of CQA is $\lambda_{\max}(x^k x^{k\top} - Y^k) > \delta$, with $\delta = 0.8$. Additionally, we limit for each node the maximum number of iterations of the **while** loop to 3.
- CQBB-2: We perform exactly one iteration of the **while** loop of CQA, taking $\delta = 0$.

Both CQBB-1 and CQBB-2 start from the initial set $\mathcal{K}_0 = \{S_0^+, S_0^*, \mathbf{0}\}$, where S_0^+ is given by (11) and S_0^* is the optimal matrix of method MIQCR. It is obtained by solving (SDP) heuristically by calling the solver Mosek [38] together with the Conic Bundle library [39] within a Lagrangian duality framework as described in [40]. We only use $r^0 = 1000$ to compute all S_k with $k \geq 1$, a value that proved more than sufficient to make inequality (10) hold. The QCQP solver for optimizing $(\bar{P}_{\mathcal{K}_k})$ at each CQA iteration is the Mosek solver using the C interface. We define $S_{k+1} = r^k \sum v_i v_i^T$, summing over all eigenvectors v_i associated to an eigenvalue no smaller than 0.01. We computed all eigenvalues and eigenvectors using the C++ Eigen library.

Regarding the parameters of the boxes in relation to ancestor nodes (described in Sections 4.2.1), we use $\beta = 3$ for instances with up to 90 variables (*i.e.*, we restrict the variables that did not change for the father, grandfather or great-grandfather node), and $\beta = 12$ for the instances with $n \geq 100$.

Characteristics of the instances

We consider the *boxqp* instances that **require** minimizing a continuous quadratic function within box constraints. The sizes of the instances vary from $n = 20$ to 125 and the densities (% of non-zero elements) of matrix Q from 20% to 100%.

Technical configuration and time limit

We run our **experiments** on a cluster under Open Suze Linux with 2 CPU Intel Xeon of 2.3 GHz. We set the time limit to **one hour for all** instances (for all compared methods, using one thread); the relative optimality gap of the branch-and-bound to

$$\epsilon_b = 10^{-4}.$$

5.1 Comparing CQBB-1 and CQBB-2 with the starting point MIQCR

We start by a detailed comparison of the two versions CQBB-1 and CQBB-2 with the original method MIQCR [33]. Recall that the latter approach consists at each node of the branch-and-bound, in evaluating once the relaxation $(\bar{P}_{\mathcal{K}})$ where the set \mathcal{K} is a static set only composed of S_0^* . To make a fair comparison, we will use a new implementation of MIQCR, with the same branching strategy and upper bound heuristic as that of CQBB. Moreover, for MIQCR, we solve (SDP) as described above for algorithm CQBB.

Tables 1-3 report the results of methods MIQCR, CQBB-1 and CQBB-2 on small instances ($20 \leq n \leq 40$), medium-sized instances ($50 \leq n \leq 90$) and respectively large instances ($n \geq 100$). For each of these three tables, the column “*Instance*” represents the instance under the name n - d - k , where n is the number of variables, d is the density in percents, and k is a numeration label. The column “*Nodes*” report the total number of nodes needed by the branch-and-bound; Column “*CPU*” provides the CPU time (in seconds) required for solving the instance to global optimality. When the instance is not solved within one hour, we report the final gap ($g\%$), where $g = \left\lfloor \frac{ub-lb}{ub} \right\rfloor \times 100$, with ub and lb the upper and lower bounds of the algorithms after one hour of CPU time, respectively. Finally, the columns “*#calls*” indicate the total number of calls to the QCQP solver; a “-” symbol means that the instance was not solved within one hour.

The general rankings of the 3 algorithms are the following:

number of nodes: 1) CQBB-2; 2) CQBB-1; 3) MIQCR.

cpu time: 1) CQBB-1; 2) CQBB-2; 3) MIQCR.

More precisely, we observe that MIQCR develops 34% more nodes than CQBB-1 on average over all instances (solved within the time limit), and the number of nodes is further significantly reduced by method CQBB-2 since MIQCR requires 78% more nodes than CQBB-2 on average. As such, CQBB-2 was able to reduce the number of nodes by half (or more) for almost a third of the small and medium instances. The most spectacular improvement is visible on instance 080-25-1: CQBB-2 reduced the number of nodes from 391 to 23.

Regarding the total CPU time, we observe that MIQCR is 60% slower than CQBB-1 on average over all the instances (solved within the time limit), and 20% slower than CQBB-2. In fact, despite a reduced number of nodes, the time spent at each node penalizes the overall branch-and-bound of algorithm CQBB-2. Clearly, the stopping criteria of CQBB-1 that rely on the accuracy of the positivity of the maximum eigenvalue seems to be a good trade-off between bound quality and CPU time.

Finally, it may be interesting to comment on the number of calls to the QCQP solver (Column *#calls*). As expected, method MIQCR performs exactly one call per node. The number of calls for our new approaches is not fixed and depends on the

Table 1: The main results of the two considered CQA versions against the original MIQCR on the smallest instances

instance	MIQCR			CQBB-1			CQBB-2		
	Nodes	CPU	#calls	Nodes	CPU	#calls	Nodes	CPU	#calls
020-100-1	9	0.2	9	3	0.0	3	3	0.0	4
020-100-2	13	0.2	13	9	0.1	9	9	0.1	13
020-100-3	5	0.1	5	3	0.0	3	3	0.0	4
030-060-1	33	1.0	33	31	0.5	31	27	1.2	48
030-060-2	1	0.0	1	1	0.0	1	1	0.0	1
030-060-3	29	0.6	29	19	0.3	19	19	0.6	32
030-070-1	71	1.9	71	63	1.0	66	65	2.2	105
030-070-2	5	0.1	5	3	0.1	3	3	0.1	4
030-070-3	41	0.9	41	37	0.6	41	17	0.5	27
030-080-1	69	2.0	69	67	1.4	100	73	3.0	155
030-080-2	3	0.1	3	5	0.1	5	5	0.2	7
030-080-3	13	0.3	13	7	0.1	7	9	0.3	14
030-090-1	7	0.2	7	3	0.1	3	3	0.1	4
030-090-2	13	0.3	13	11	0.3	17	3	0.1	5
030-090-3	3	0.1	3	3	0.1	3	3	0.1	4
030-100-1	11	0.3	11	9	0.2	9	7	0.2	11
030-100-2	9	0.3	9	9	0.2	9	7	0.3	11
030-100-3	27	0.6	27	21	0.4	23	13	0.4	20
040-030-1	7	0.3	7	3	0.1	3	3	0.1	4
040-030-2	13	0.5	13	5	0.2	5	5	0.3	7
040-030-3	5	0.2	5	3	0.1	3	3	0.2	4
040-040-1	445	15	445	263	7	275	249	12	416
040-040-2	7	0.3	7	5	0.2	5	5	0.3	7
040-040-3	27	1.0	27	23	0.7	25	23	1.3	38
040-050-1	29	1.1	29	25	0.7	25	25	1.3	38
040-050-2	43	1.6	43	35	1.2	43	25	1.5	42
040-050-3	25	1.0	25	17	0.6	21	11	0.6	17
040-060-1	279	11	279	219	7	251	209	11	373
040-060-2	21	0.9	21	13	0.6	19	9	0.5	15
040-060-3	7	0.4	7	5	0.2	5	5	0.4	8
040-070-1	7	0.3	7	5	0.2	5	5	0.3	7
040-070-2	7	0.3	7	7	0.2	7	5	0.4	8
040-070-3	9	0.4	9	5	0.2	5	5	0.4	8
040-080-1	5	0.2	5	5	0.2	5	5	0.3	7
040-080-2	5	0.3	5	7	0.3	9	5	0.4	8
040-080-3	13	0.5	13	15	0.5	15	11	0.7	18
040-090-1	9	0.4	9	5	0.2	5	5	0.3	7
040-090-2	9	0.4	9	9	0.3	9	9	0.5	13
040-090-3	7	0.3	7	5	0.2	5	5	0.3	7
040-100-1	13	0.6	13	7	0.2	7	7	0.4	13
040-100-2	21	1.8	21	17	0.5	17	17	1.0	27
040-100-3	221	18	221	207	6.3	245	177	10	323

value of the solution (x^0, Y^0) computed before the **while** loop of CQA (Algorithm 1). Simply computing this (x^0, Y^0) requires a first call to the QCQP solver. CQBB-1 can call the QCQP solver up to 3 additional times because it performs at most 3 iterations of the **while** loop. CQBB-2 always performs exactly one iteration, leading to 2 calls to the QCQP solver (one before the **while** and one after). However, sometimes the **while** loop may be skipped completely, either because $\lambda_{\max}(x^0 x^{0\top} - Y^0) \leq \delta$, or because the lower bound associated to (x^0, Y^0) is good enough to prune the node. As such, we do not systematically have 2 QCQP solver calls per iteration in the case of CQBB-2.

Table 2: The main results of the two considered CQA variants against the original MIQCR on medium-size instances

instance	MIQCR			CQBB-1			CQBB-2		
	<i>Nodes</i>	<i>CPU</i>	<i>#calls</i>	<i>Nodes</i>	<i>CPU</i>	<i>#calls</i>	<i>Nodes</i>	<i>CPU</i>	<i>#calls</i>
050-030-1	13	0.8	13	9	0.5	11	7	0.8	12
050-030-2	19	1.1	19	17	0.9	19	15	1.4	24
050-030-3	31	1.6	31	27	1.4	35	13	1.2	20
050-040-1	9	0.7	9	9	0.6	12	5	0.6	8
050-040-2	43	2.5	43	31	1.3	31	25	2.1	40
050-040-3	7	0.5	7	7	0.4	7	5	0.6	8
050-050-1	5207	381	5207	4549	218	5088	4447	455	8286
050-050-2	65	3.7	65	67	3.2	77	53	4.5	88
050-050-3	155	8	155	121	7	180	63	5.8	114
060-020-1	13	1.2	13	5	0.5	5	5	0.8	7
060-020-2	11	1.0	11	5	0.5	5	5	0.8	7
060-020-3	61	4.7	61	33	2.1	33	33	4.2	54
070-025-1	77	8	77	41	3.7	41	33	6.5	57
070-025-2	199	22	199	129	11	133	109	21	209
070-025-3	225	25	225	145	15	168	105	17	174
070-050-1	183	21	183	159	17	184	145	26	255
070-050-2	29	3.6	29	25	3.0	31	23	4.7	40
070-050-3	9	1.4	9	11	1.6	15	7	1.7	11
070-075-1	65	7	65	59	6.1	63	55	11	108
070-075-2	1735	203	1735	1435	161	1750	1401	263	2592
070-075-3	819	96	819	729	86	990	585	113	1148
080-025-1	391	61	391	227	41	412	23	8	53
080-025-2	559	88	559	433	62	487	405	114	796
080-025-3	161	25	161	95	12	99	77	19	131
080-050-1	21543	3532	21543	17953	2631	21255	12597	(0.38%)	–
080-050-2	65	10	65	77	15	137	21	6.8	43
080-050-3	327	56	327	293	48	361	279	87	578
080-075-1	113	18	113	101	16	127	103	30	214
080-075-2	333	52	333	309	51	390	301	82	579
080-075-3	1149	183	1149	1049	168	1301	1075	294	2115
090-025-1	3243	668	3243	2141	418	2552	2149	728	3985
090-025-2	2683	561	2683	1703	350	2038	1567	516	2780
090-025-3	623	126	623	405	77	472	391	124	684
090-050-1	1707	387	1707	1413	321	1783	1419	550	2885
090-050-2	21	5.3	21	19	4.2	22	15	7.0	27
090-050-3	493	113	493	389	88	495	393	154	806
090-075-1	6037	1209	6037	5919	1347	7746	6387	2510	13509
090-075-2	5947	1234	5947	5451	1206	6831	5371	2005	10611
090-075-3	1353	282	1353	1251	298	1677	1231	487	2605

Table 3: The main results of the two considered CQA variants against the original MIQCR on the largest instances

instance	MIQCR			CQBB-1			CQBB-2		
	<i>Nodes</i>	<i>CPU</i>	<i>#calls</i>	<i>Nodes</i>	<i>CPU</i>	<i>#calls</i>	<i>Nodes</i>	<i>CPU</i>	<i>#calls</i>
100-025-1	665	182	665	505	121	518	429	209	776
100-025-2	315	85	315	261	65	292	191	85	323
100-025-3	269	75	269	191	45	195	199	90	343
100-050-1	12720	(1.67%)	—	13380	(1.66%)	—	7221	(1.98%)	—
100-050-2	12546	(0.76%)	—	13034	(0.82%)	—	6934	(1.07%)	—
100-050-3	811	235	811	755	190	813	663	309	1225
100-075-1	12001	3490	12001	11929	3107	13170	7199	(0.41%)	—
100-075-2	9479	2723	9479	9517	2563	10694	7099	(0.34%)	—
100-075-3	8867	2511	8867	8829	2268	9575	7437	(0.22%)	—
125-025-1	7319	(1.79%)	—	7832	(1.72%)	—	3824	(2.15%)	—
125-025-2	6994	(0.01%)	—	5629	2577	6144	3970	(0.06%)	—
125-025-3	1733	895	1733	1365	625	1419	1275	1024	2088
125-050-1	7555	(1.59%)	—	7662	(1.55%)	—	4158	(1.81%)	—
125-050-2	7329	(0.90%)	—	7344	(0.90%)	—	4042	(1.05%)	—
125-050-3	7102	(0.84%)	—	7324	(0.82%)	—	4011	(0.07%)	—
125-075-1	4841	2483	4841	4827	2175	5267	4557	(0.03%)	—
125-075-2	6943	(2.48%)	—	7608	(2.36%)	—	4198	(2.94%)	—
125-075-3	6856	(1.35%)	—	7699	(1.26%)	—	4232	(1.44%)	—

5.2 General comparison with existing solvers

We now compare methods CQBB-1, CQBB-2 and MIQCR [33] to the standard solvers cplex 22.1.1 [37], Baron 24.3.19 [22], and Gurobi 11.0.3 [25], as well as to the algorithms QuadProgBB [11] which is a spatial branch-and-bound based on the satisfaction of KKT optimality conditions, and GloMIQO 2 [24] that mixes several algorithmic components for solving MIQCPs.

Solvers and algorithms serving as a basis for comparison

These are the settings of the solvers mentioned above:

- For the solvers cplex 22.1.1 [37] and Gurobi 11.0.3 [25], we use the AMPL [41] interface and the default parameters.
- For the solver Baron 24.3.19 [22], we use the Gams [42] interface and the default parameters.
- For algorithms QuadProgBB [11] and GloMIQO 2 [24], since we do not have the licenses of GloMIQO 2 and of matlab necessary for QuadProgBB, we report the results of the paper [29]. Note that these experiments were carried out on a server with a very similar configuration to our server, belonging to a similar generation.

To compare the considered methods, we use a performance profile of the CPU times (see [43] for a complete description). The basic idea is the following: for each instance i and each solver s , we denote by t_{is} the time for solving instance i by solver s , and we define the *performance ratio* as $r_{is} = \frac{t_{is}}{\min_s t_{is}}$. Let N be the total number of

instances considered; an overall assessment of the performance of solver s for a given τ is given by: $P(r_{is} \leq \tau) = \frac{1}{N} * \text{number of instances } i \text{ such that } r_{is} \leq \tau$.

Thus, in a performance profile of CPU times, each curve corresponds to a solver, where each point of a curve gives, for a given factor τ , the percentage of instances whose CPU time was at most τ times greater than the minimal CPU time reported by any of all the solvers. In particular, for $\tau = 1$, we have the proportion of instances for which the considered solver was the fastest method.

We present the performance profiles of the CPU times for all compared algorithms on the *boxqp* instances with up to 90 variables in Figure 3, and on the largest instances (with $n \geq 100$) in Figure 4. We observe that our new approach CQBB-1 compares well with MIQCR and all standard solvers in terms of the number of instances solved (maximum τ). For the small and medium-sized instances (Figure 3), GloMIQO solves 69 instances, Baron solves 73 instances, Cplex solves 78 instances, Gurobi and CQBB-2 solve 80 instances, and QuadProgBB, MIQCR and CQBB-1 solve 81 instances out of 81 (within the time limit of 1 hour). However, note that the performance profile starts for $\tau = 1$ with a smaller value $P(r_{is} \leq \tau)$ for CQBB-1 than for the commercial solvers. This is because the smallest instances are solved more rapidly by Gurobi, Cplex, GloMIQO, or Baron.

However, the focus of our algorithm is on the largest and denser instances, and, with regards to them, Figure 4 shows that the (dark orange) curve of CQBB-1 is above all other curves except over the segment $\tau < 30$ where it is only dominated by Gurobi.

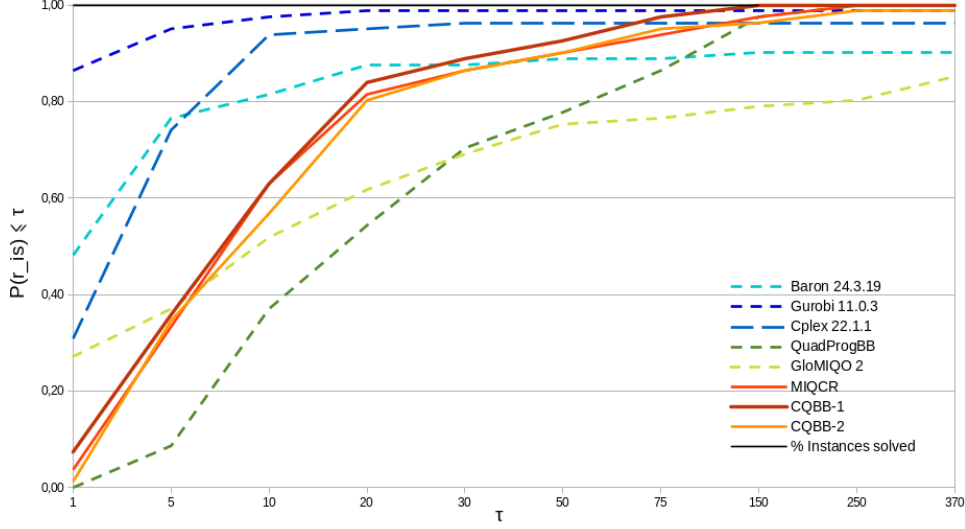


Fig. 3: Performance profile of the total CPU time for the *boxqp* instances with $n = 20$ to 90 within a time limit of 1 hour.

On the long run ($\tau \geq 30$), the winner is CQBB-1 followed by Gurobi and MIQCR. Considering the maximum time limit of 1 hour, CQBB-1 solves 10 instances, MIQCR and Gurobi 9 instances, QuadProgBB and Cplex 6 instances, CQBB-2 5 instances, Baron 3 instances, and GloMIQO 2 instances.

We also noticed that the average final gap of CQBB-1 over the larger instances unsolved within the time limit is roughly 1.4%, while it is around 10% for Gurobi and 32% for Cplex. This can be explained by the tightness of the SDP convexification computed at the root node by CQBB-1.

6 Conclusions

We have presented a generic approach to solve box-constrained quadratic programs to global optimality. The main idea is to combine the strength of quadratic convex relaxations with the cutting-planes logic. Indeed, instead of considering a unique convex quadric to underestimate the objective function, we propose a family of convexifications indexed by an arbitrary number of quadrics.

We have proposed the idea to add these cutting-quadrics one by one, in a cutting-planes fashion. The quadric generated at each iteration is actually a quadratic cut that separates the current optimal solution, acting similarly to a hyper-plane in a cutting-planes method. We proved that this iterative algorithm converges to the optimal value of a tight semidefinite relaxation of (QBP).

This iterative algorithm was integrated into a spatial branch-and-bound method; the lower bound of each branch-and-bound node is determined using the new multi-quadric convexification. By generating quadrics specifically-tailored for each node,

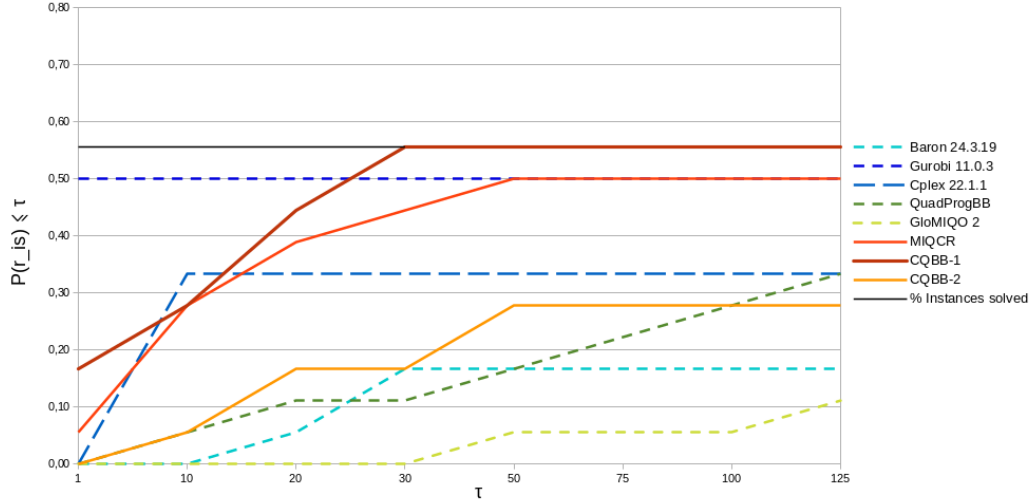


Fig. 4: Performance profile of the total CPU time for the *boxqp* instances with $n \geq 100$ within a time limit of 1 hour

we almost systematically reduced the total number of nodes compared to existing methods that convexify using a unique quadric. The main computational bottleneck comes from the convex QCQP solver that has to be called after adding each new quadric; we managed to accelerate this step by using artificial tentative boxes over the decision variables (that are lifted only if necessary). Numerical results suggest it is more practical to use relatively few quadrics per node ($p \leq 3$); the resulting methods compete well not only with existing solvers based on other convexifications, but also with other solvers that use completely different techniques.

References

- [1] Horst, R., Pardalos, P.M., Thoai, N.V.: Introduction to global optimization. Kluwer, Dordrecht (2000)
- [2] Glowinski, R.: Numerical methods for nonlinear variational problems. Springer-Verlag, Berlin, New York (1984)
- [3] W. Glunt, M.R. T.L. Hayden: Molecular conformations from distance matrices. J. Comput. Chem. **14**, 114–120 (1993)
- [4] E.G. Birgin, J.M.M. I. Chambouleyron: Estimation of the optical constants and the thickness of thin films using unconstrained optimization. J. Comput. Phys. **151**, 862–880 (1999)
- [5] Jain, P., Kar, P.: Non-convex optimization for machine learning. Foundations and Trends® in Machine Learning **10**(3-4), 142–363 (2017) <https://doi.org/10.1561/>

- [6] Wen, F., Chu, L., Liu, P., Qiu, R.C.: A survey on nonconvex regularization-based sparse and low-rank recovery in signal processing, statistics, and machine learning. *IEEE Access* **6**, 69883–69906 (2018) <https://doi.org/10.1109/ACCESS.2018.2880454>
- [7] An, L.T.H., Tao, P.D.: A branch and bound method via d.c. optimization algorithms and ellipsoidal technique for box constrained nonconvex quadratic problems. *Journal of Global Optimization* **13**(2), 171–206 (1998) <https://doi.org/10.1023/A:1008240227198>
- [8] Burer, S., Vandenbussche, D.: Globally solving box-constrained nonconvex quadratic programs with semidefinite-based finite branch-and-bound. *Comput Optim Appl* **43**, 181–195 (2009)
- [9] Vandenbussche, D., Nemhauser, G.: A branch-and-cut algorithm for nonconvex quadratic programs with box constraints. *Mathematical Programming* **102**(3), 259–275 (2005)
- [10] Bonami, P., Günlük, O., Linderoth, J.: Globally solving nonconvex quadratic programming problems with box constraints via integer programming methods. *Mathematical Programming Computation* **10**(3), 333–382 (2018)
- [11] Chen, J., Burer, S.: Globally solving nonconvex quadratic programming problems via completely positive programming. *Mathematical Programming Computation* **4**(1), 33–52 (2012)
- [12] Blik, C., Bonami, P., Lodi, A.: Solving Mixed-Integer Quadratic Programming problems with IBM-CPLEX: a progress report. In: *Proceedings of the Twenty-Sixth RAMP Symposium*, Hosei University, Tokyo, October 16-17 (2014)
- [13] Misener, R., Smadbeck, J.B., Floudas, C.A.: Dynamically generated cutting planes for mixed-integer quadratically constrained quadratic programs and their incorporation into GloMIQO 2. *Optimization Methods and Software* **30**(1), 215–249 (2015)
- [14] Tawarmalani, M., Sahinidis, N.V.: Global optimization of mixed-integer nonlinear programs: A theoretical and computational study. *Mathematical programming* **99**(3), 563–591 (2004)
- [15] Anstreicher, K.M., Burer, S.: Computable representations for convex hulls of low-dimensional quadratic forms. *Mathematical Programming* **124**(1), 33–43 (2010) <https://doi.org/10.1007/s10107-010-0355-9>
- [16] Burer, S., Letchford, A.N.: On nonconvex quadratic programming with box constraints. *SIAM Journal on Optimization* **20**(2), 1073–1089 (2009) <https://doi.org/10.1137/0807073>

- [17] Dong, H., Linderoth, J.: On valid inequalities for quadratic programming with continuous variables and binary indicators. In: Goemans, M., Correa, J. (eds.) *Integer Programming and Combinatorial Optimization*, pp. 169–180. Springer, Berlin, Heidelberg (2013)
- [18] McCormick, G.P.: Computability of global solutions to factorable non-convex programs: Part i - convex underestimating problems. *Mathematical Programming* **10**(1), 147–175 (1976)
- [19] Sherali, H.D., Adams, W.P.: *A Reformulation-linearization Technique for Solving Discrete and Continuous Nonconvex Problems* vol. 31. Springer, ??? (2013)
- [20] Yajima, Y., Fujie, T.: A polyhedral approach for nonconvex quadratic programming problems with box constraints. *Journal of Global Optimization* **13**(2), 151–170 (1998)
- [21] Lambert, A.: Using general triangle inequalities within quadratic convex reformulation method. *Optimization Methods and Software* **38**(3), 626–653 (2023) <https://doi.org/10.1080/10556788.2022.2157002>
- [22] Sahinidis, N.V., Tawarmalani, M.: *Baron 9.0.4: Global optimization of mixed-integer nonlinear programs. User’s Manual* (2010)
- [23] Misener, R., Floudas, C.A.: Global optimization of mixed-integer quadratically-constrained quadratic programs (MIQCQP) through piecewise-linear and edge-concave relaxations. *Math. Program. B* **136**(1), 155–182 (2012). http://www.optimization-online.org/DB_HTML/2011/11/3240.html
- [24] Misener, R., Floudas, C.A.: GloMIQO: Global mixed-integer quadratic optimizer. *Journal of Global Optimization* **57**(1), 3–50 (2013) <https://doi.org/10.1007/s10898-012-9874-7>
- [25] Gurobi Optimization: version 11.0.3 (2024). <https://www.gurobi.com/>
- [26] Anstreicher, K.M.: Semidefinite programming versus the reformulation-linearization technique for nonconvex quadratically constrained quadratic programming. *Journal of Global Optimization* **43**(2), 471–484 (2009) <https://doi.org/10.1007/s10898-008-9372-0>
- [27] Burer, S., Vandenbussche, D.: A finite branch-and-bound algorithm for nonconvex quadratic programming via semidefinite relaxations. *Mathematical Programming* **113**(2), 259–282 (2008) <https://doi.org/10.1007/s10107-006-0080-6>
- [28] Vandenbussche, D., Nemhauser, G.L.: A polyhedral study of nonconvex quadratic programs with box constraints. *Mathematical Programming* **102**(3), 531–557

(2005)

- [29] Saxena, A., Bonami, P., Lee, J.: Convex relaxations of non-convex mixed integer quadratically constrained programs: projected formulations. *Mathematical Programming* **130**, 359–413 (2011)
- [30] Gorge, A., Lisser, A., Zorgati, R.: Generating cutting planes for the semidefinite relaxation of quadratic programs. *Computers & Operations Research* **55**, 65–75 (2015) <https://doi.org/10.1016/j.cor.2014.09.008>
- [31] Billionnet, A., Elloumi, S., Lambert, A.: Extending the QCR method to the case of general mixed integer program. *Mathematical Programming* **131**(1), 381–401 (2012)
- [32] Billionnet, A., Elloumi, S., Lambert, A.: Exact quadratic convex reformulations of mixed-integer quadratically constrained problems. *Mathematical Programming* **158**(1), 235–266 (2016) <https://doi.org/10.1007/s10107-015-0921-2>
- [33] Elloumi, S., Lambert, A.: Global solution of non-convex quadratically constrained quadratic programs. *Optimization Methods and Software* **34**(1), 98–114 (2019) <https://doi.org/10.1080/10556788.2017.1350675>
- [34] Porumbel, D.: Demystifying the characterization of SDP matrices in mathematical programming. arXiv preprint nr. 2210.13072 (2022)
- [35] Higham, N.J., Strabic, N., Sego, V.: Restoring definiteness via shrinking, with an application to correlation matrices with a fixed block. *siam REVIEW* **58**(2), 245–263 (2016)
- [36] Belotti, P., Kirches, C., Leyffer, S., Linderoth, J., Luedtke, J., Mahajan, A.: Mixed-integer nonlinear optimization. *Acta Numerica* **22**, 1–131 (2013)
- [37] IBM: Cplex version 22.1.1 (2022). <https://www.ibm.com/>
- [38] ApS, M.: The MOSEK Optimization Toolbox for MATLAB Manual. Version 9.2. (2019). <http://docs.mosek.com/9.0/toolbox/index.html>
- [39] Helmberg, C.: Conic Bundle V0.3.10. (2011). <http://www-user.tu-chemnitz.de/helmberg/ConicBundle/>
- [40] Billionnet, A., Elloumi, S., Lambert, A., Wiegele, A.: Using a Conic Bundle method to accelerate both phases of a Quadratic Convex Reformulation. *INFORMS Journal on Computing* **29**(2), 318–331 (2017) <https://doi.org/10.1287/ijoc.2016.0731>
- [41] Fourer, R., Gay, D.M., Kernighan, B.W.: AMPL: A Modeling Language for Mathematical Programming, p. 351. The Scientific Press (now an imprint of Boyd & Fraser Publishing Co.), Danvers, MA, USA (1993)

- [42] GAMS Development Corporation. Fairfax, U. VA: General Algebraic Modeling System (GAMS) Release 36.1.0. (2021). <https://www.gams.com/download/>
- [43] Dolan, D., Moré, J.: Benchmarking optimization software with performance profiles. *Mathematical Programming* **91**, 201–213 (1986)

Data Availability Statement

There is no original new data introduced in this work. The instances we use are well known for many years to the experts in the field. We do cite [8, 9] in the beginning of the numerical results (Section 5) for the interested reader to find the source.

Compliance with Ethical Standards

- There is no (financial or non-financial) conflict of interest with any third party
- The research involved no other human or animal participants; the question of their involvement does not apply
- No funding received
- All authors consent to submitting the paper

Response to the reviews of paper JOGO-D-24-00042

Dear editors and reviewers,

Thank you all for taking the time to work on our paper. We have revised it and did our best to reply to all remarks. All new text insertions are marked either in green (a few words added) or in blue (at least a complete new phrase).

Amélie Lambert and Daniel Porumbel

Reviewer 1

>Reviewer #1: The paper presents an approach to solving
>box-constrained quadratic programs to global optimality by
>strengthening the convexification of the program through
>the iterative addition of convex quadratic cuts. The
>authors integrate this approach into a spatial
>branch-and-bound algorithm and present numerical results
>showing that, in some cases, the algorithm outperforms
>well-known solvers. The paper is well-written and contains
>the necessary theoretical details.

Thank you for taking the time to read our paper and provide this summary.

> 1. The paper would benefit from adding an example and
> illustration of the proposed algorithm, similar to
> reference [25].

We added an example as requested (Section 3.1, page 13). This shows that the convexification becomes more tighter as we add more quadratic cuts. Thanks to this, we even find the optimal solution at the root node of the branch-and-bound in this example.

> 2. The discussion on page 12 could be expanded to explore
> strategies for selecting K_0 more effectively.

We added a paragraph (in blue, page 12) just before the heading of Section 3.1 to discuss other ideas of initial SDP matrices to populate K_0 . As indicated there, the effectiveness of any SDP matrix choice can be evaluated only on an instance by instance basis, by trying and testing it.

> 3. Discussing potential applications or providing a case
> study where this method could be particularly effective
> would add practical value to the paper.

We added a sentence (in blue, page 2, the second paragraph of the introduction) listing many applications of box-constrained quadratic programming, also insisting that BQP is often used as a sub-problem of more complex problems. The method we propose is a generic one, which has precisely the strength of solving problems of this class independently of their underlying structures. The proposed techniques are quite costly, and are necessarily more effective on the most difficult instances to solve.

In particular, instances where the number of quadratic terms is limited are usually solved very efficiently using a standard linearization; in such case, our method will be slower. On the other hand, our approach based on quadratic cuts is most efficient for the densest instances with an indefinite Q .

>There are typos that should be fixed:

>P.16: Open Suze Linux -> Suse

>p.16 :experiences -> experiments

>P. 19: "Baron 74 solves instances"

>Section 4.1.2 typo in transpose symbol

Done. Several notations of the form $x^k x^{k\top}$ have been corrected, with the correctly placed transpose symbol.

Reviewer 2

> Reviewer #2: The manuscript proposes a novel approach to
> Quadratic Box-constrained Programming (QBP) by introducing
> cutting planes logic to tighten the bounds. The authors
> present an iterative algorithm named the "Cutting-Quadratics
> Algorithm (CQA)," wherein cuts are introduced to a
> quadratic convex relaxation of the QBP. The paper also
> provides a proof of convergence for the proposed
> algorithm. It's not an unsolved or unknown problem - the
> problem has been studied extensively.

Thank you for taking the time for writing this review of our work.

> This paper uses Gurobi version 9.11. There is significant
> speed-up since then. Additionally, Cplex 20.1, (22 is
> available) was released in December 2020. This is a point
> to be questioned.

As suggested, we have solved the considered instances with the latest versions of the Gurobi (11.0.3) and Cplex (22.1) solvers in order to compare their computation times with those of our new algorithm (Section 5.2).

> I think the parameters and solution strategies in the
> computational results section are a bit confusingly
> explained.

We agree that in the initial submission, we were mixing remarks characterising our algorithms with comments on external solvers or on the instances. The first page of Section 5 was reorganized by separating these remarks under separated mini-headings. Moreover, to simplify the reading, we fixed the time limit to 1 hours for all instances.

> The BB + Cutting proposed by this study doesn't seem to be
> suggested for the first time in the literature. Now,
> here's a statement made by Gorge (2015): ''Semidefinite

> Programming is well-known for providing relaxations of
 > quadratic programs. In practice, only few real-world
 > applications of this approach have been reported. This can
 > be explained by the fact that the standard semidefinite
 > relaxation must generally be tightened with cuts, which
 > increases the substantial computational cost of the
 > semidefinite program. Then, the challenge is to come up
 > with the most effective cuts.’’ Can you elaborate on that?

We introduced references [29] and [30], and a paragraph (in blue, page 4) just before the heading of Section 1.1.2 to better describe the context of our work, with regards to the interesting point of finding the most effective cuts (while limiting their number not to explode the program size).

> Although it is mentioned later in the manuscript, it would
 > benefit the reader if it were clarified earlier that in
 > Algorithm 1, $\overline{P}_{\mathcal{K}_{k+1}}$ is
 > resolved from scratch. It would also be helpful to note
 > that techniques to enhance solver performance are
 > introduced in subsequent sections.

As suggested, we introduced (in blue, page 9) a fourth paragraph in Section 3 to discuss in advance the idea of solving from scratch each new program $(\overline{P}_{\mathcal{K}_{k+1}})$.

> I would have expected a comparison with other existing
 > algorithms rather than just the original algorithm +
 > extension. Not showing it in the computation while
 > discussing it in the literature seems a bit off to me.

We added in Section 5 a comparison in terms of computation time with two existing algorithms from the literature: QuadProgBB [11], which is a spatial branch-and-bound based on the satisfaction of KKT optimality conditions, and GloMIQO 2 [24], which is a spatial branch-and-bound based on quadratic relaxations, also called α -BB.

There are some minor grammatical errors.

"rue saint-Martin" should be "rue Saint-Martin"

Page 1, key words, "piecewise-uadratic" should be "piecewise-quadratic."

Page 2, paragraph 1, "carthesian" should be "cartesian".

Page 5, paragraph 1, after "As mentioned previously" there should be a comma.

Corrected, thanks.