

Consensus-ADMM for General Quadratically Constrained Quadratic Programming

Kejun Huang, *Student Member, IEEE*, and Nicholas D. Sidiropoulos, *Fellow, IEEE*

Abstract—Nonconvex quadratically constrained quadratic programming (QCQP) problems have numerous applications in signal processing, machine learning, and wireless communications, albeit the general QCQP is NP-hard, and several interesting special cases are NP-hard as well. This paper proposes a new algorithm for general QCQP. The problem is first reformulated in consensus optimization form, to which the alternating direction method of multipliers can be applied. The reformulation is done in such a way that each of the subproblems is a QCQP with only one constraint (QCQP-1), which is efficiently solvable irrespective of (non)convexity. The core components are carefully designed to make the overall algorithm more scalable, including efficient methods for solving QCQP-1, memory efficient implementation, parallel/distributed implementation, and smart initialization. The proposed algorithm is then tested in two applications: multicast beamforming and phase retrieval. The results indicate superior performance over prior state-of-the-art methods.

Index Terms—Alternating direction method of multipliers (ADMM), feasible point pursuit, multicast beamforming, phase retrieval, non-convex quadratically constrained quadratic programming (QCQP), semi-definite relaxation (SDR).

I. INTRODUCTION

QUADRATICALLY constrained quadratic programming (QCQP) is an optimization problem that minimizes a quadratic function subject to quadratic inequality and equality constraints [1]. We write it in the most general form as follows:

$$\begin{aligned} & \underset{\mathbf{x} \in \mathcal{C}^n}{\text{minimize}} && \mathbf{x}^H \mathbf{A}_0 \mathbf{x} - 2\Re \{ \mathbf{b}_0^H \mathbf{x} \}, \\ & \text{subject to} && \mathbf{x}^H \mathbf{A}_i \mathbf{x} - 2\Re \{ \mathbf{b}_i^H \mathbf{x} \} \leq c_i, \\ & && \forall i = 1, \dots, m. \end{aligned} \quad (1)$$

Notice that for simplicity we only write the constraints as inequalities, but they can be equalities as well (each can be expressed as two inequalities).

A QCQP is in general NP-hard, except for some special cases, for example when all the $\{\mathbf{A}_i\}_{i=1}^m$ in “ \leq ” inequality constraints are positive semi-definite [2, § 4.4], m is “small” [3]–[6], or if the quadratic terms are all homogeneous and $\{\mathbf{A}_i\}_{i=0}^m$ are all Toeplitz [7].

Manuscript received January 4, 2016; revised June 10, 2016; accepted July 9, 2016. Date of publication July 20, 2016; date of current version August 19, 2016. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Sergios Theodoridis. This work was supported in part by National Science Foundation under Project CIF-1525194.

The authors are with the Department of Electrical and Computer Engineering, University of Minnesota Twin Cities, Minneapolis, MN 55455 USA (e-mail: huang663@umn.edu; nikos@umn.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2016.2593681

For general non-convex QCQPs, the prevailing method to tackle the problem is through semi-definite relaxation (SDR), where the following semi-definite programming (SDP) problem is solved instead

$$\begin{aligned} & \underset{\mathbf{X} \in \mathcal{H}^n, \mathbf{x} \in \mathcal{C}^n}{\text{minimize}} && \text{Trace} \{ \mathbf{A}_0 \mathbf{X} \} - 2\Re \{ \mathbf{b}_0^H \mathbf{x} \}, \\ & \text{subject to} && \begin{bmatrix} \mathbf{X} & \mathbf{x} \\ \mathbf{x}^H & 1 \end{bmatrix} \succeq 0, \\ & && \text{Trace} \{ \mathbf{A}_i \mathbf{X} \} - 2\Re \{ \mathbf{b}_i^H \mathbf{x} \} \leq c_i, \\ & && \forall i = 1, \dots, m, \end{aligned} \quad (2)$$

where \mathcal{H}^n denotes the set of $n \times n$ complex Hermitian matrices. Problem (2) is obtained by relaxing the otherwise equivalent constraint $\mathbf{X} = \mathbf{x} \mathbf{x}^H$ to a convex one $\mathbf{X} \succeq \mathbf{x} \mathbf{x}^H$. After solving the convex relaxation problem (2), we not only get a non-trivial lower bound on the optimal cost of (1), but also possibly a solution of (1) if the solution \mathbf{X}_* of (2) turns out to be rank one. If this is not the case, in certain cases there is an efficient way to generate approximate solutions to the original problem in (1) from the higher-rank solution of (2). Let $(\mathbf{x}_*, \mathbf{X}_*)$ be an optimal solution of (2), then by drawing random points $\mathbf{x} \sim \mathcal{CN}(\mathbf{x}_*, \mathbf{X}_*)$, possibly followed by a simple projection or scaling if applicable, one can obtain approximate solutions not far away from the SDR lower bound, for certain kinds of QCQP problems. That is, in certain cases it has been shown that this randomization step (with sufficient draws) is guaranteed to yield some quantified sub-optimality, see for example [8], [9].

If the problem dimension in (1) is large, then squaring the number of variables as in (2) makes the latter very difficult to solve. If a general purpose SDP solver is used to solve (2) using the interior point method, the worst case complexity can be as high as $\mathcal{O}(n^{6.5})$. Another downside of SDR is that, if it is not obvious how to obtain a feasible point for the constraint set of (1), in a lot of cases SDR randomization will not give us a feasible point either.

Another way to tackle problem (1) is through convex restriction, also known as successive convex approximation (SCA) or convex-concave procedure (CCP) for the more general difference of convex programming (DCP) problem [10]. Noticing that any Hermitian matrix can be written as the difference of two positive semi-definite matrices, we can denote

$$\mathbf{A}_i = \mathbf{A}_i^{(+)} + \mathbf{A}_i^{(-)}, \forall i = 0, 1, \dots, m,$$

where $\mathbf{A}_i^{(+)} \succeq 0$ and $\mathbf{A}_i^{(-)} \preceq 0$. Then for each quadratic term, we have that

$$\mathbf{x}^H \mathbf{A}_i \mathbf{x} \leq \mathbf{x}^H \mathbf{A}_i^{(+)} \mathbf{x} + 2\Re \left\{ \mathbf{x}^{(0)H} \mathbf{A}_i^{(-)} \mathbf{x} \right\} - \mathbf{x}^{(0)H} \mathbf{A}_i^{(-)} \mathbf{x}^{(0)},$$

for any point $\mathbf{x}^{(0)} \in \mathcal{C}^n$. Therefore, starting with an initial point $\mathbf{x}^{(0)}$, we can iteratively solve the following convex QCQP until we obtain an approximate solution of (1)

$$\begin{aligned} \mathbf{x}^{(t+1)} \leftarrow & \\ \arg \min_{\mathbf{x}} & \mathbf{x}^H \mathbf{A}_0^{(+)} \mathbf{x} + 2\Re \left\{ \mathbf{x}^{(t)H} \mathbf{A}_0^{(-)} \mathbf{x} \right\} - 2\Re \left\{ \mathbf{b}_0^H \mathbf{x} \right\} \\ \text{subject to} & \mathbf{x}^H \mathbf{A}_i^{(+)} \mathbf{x} + 2\Re \left\{ \mathbf{x}^{(t)H} \mathbf{A}_i^{(-)} \mathbf{x} \right\} - 2\Re \left\{ \mathbf{b}_i^H \mathbf{x} \right\} \\ & \leq c_i + \mathbf{x}^{(t)H} \mathbf{A}_i^{(-)} \mathbf{x}^{(t)}, \forall i = 1, \dots, m. \end{aligned}$$

This is a convex restriction because each quadratic function is replaced with its convex upper bound function. If we start with a feasible $\mathbf{x}^{(0)}$, then it is easy to show that the sequence $\{\mathbf{x}^{(t)}\}$ will remain feasible afterwards. However, if we start with an infeasible $\mathbf{x}^{(0)}$, it is possible (and often the case) that the restricted constraint set becomes empty, thus the iterates cannot proceed. Recently [11] proposed *feasible point pursuit - successive convex approximation* (FPP-SCA) to address this issue, by adding a slack variable to each constraint and an ℓ_1 penalty on the slacks to the cost. FPP-SCA produces good results in both finding a feasible point and approaching closer to the SDR lower bound. The potential disadvantage of FPP-SCA is that in each iteration we still need to solve a non-trivial convex optimization problem, which may take a lot of time even for a moderate number of iterations, if a general-purpose convex optimization solver is used for large-scale problems.

In this paper, we propose a rather different approach to handle QCQPs. The proposed algorithm is based on two building blocks:

- 1) Any QCQP with only one constraint (QCQP-1) can be solved to optimality, and in various cases this can be done efficiently;
- 2) Adopting the alternating direction method of multipliers (ADMM) for consensus optimization as the general algorithmic framework, problem (1) can be decomposed into m QCQP-1's in each iteration, thus leading to efficient optimal updates.

In the rest of this section, we briefly review the basics of the two aforementioned building blocks. Then the general algorithmic framework is introduced in Section II. In Section III, we look into one specific step of the algorithm, and explain how this seemingly non-trivial QCQP-1 sub-problem can be solved very efficiently. Some detailed implementation issues are described in Section IV, including a memory-efficient implementation for certain types of constraints, an empirical way of initialization that works very well in practice, and discussions on parallel and distributed implementations with small communication overhead. Simulation results are presented in Section V, where the proposed algorithm is used for feasible point pursuit, multicast beamforming, and phase retrieval, showing great performance and versatility in various kinds of non-convex QCQP problems. Conclusions are drawn in Section VI.

A. QCQP With Only One Constraint (QCQP-1)

One of the most fundamental results in QCQP is that, any QCQP with only one constraint (QCQP-1) can be solved

optimally, despite the fact that the quadratic terms may be indefinite. The fundamental idea behind this result is the following lemma [2, Appendix B].

Lemma 1: For all $\mathbf{X}, \mathbf{A}, \mathbf{B} \in \mathcal{H}^n$, and $\mathbf{X} \succeq 0$, there exists an $\mathbf{x} \in \mathcal{C}^n$ such that

$$\mathbf{x}^H \mathbf{A} \mathbf{x} = \text{Trace} \{ \mathbf{A} \mathbf{X} \}, \mathbf{x}^H \mathbf{B} \mathbf{x} = \text{Trace} \{ \mathbf{B} \mathbf{X} \}.$$

This means that after we find a solution for the SDR of a QCQP-1, regardless of its rank, we can always find an equivalent rank one solution. Our experience from simulations is that SDR seems to always return a rank one solution for a QCQP-1. Even if not, one can resort to rank reduction as in [5], which handles a more general rank reduction problem. This result is also closely related to the *generalized eigenvalue problem* [12, §8.7] in linear algebra and the *S-procedure* [13, §2.6] in control.

B. Consensus Optimization Using ADMM

Now we briefly introduce the algorithmic tool to be used in this paper, which is based on the ADMM [14]. Consider the following optimization problem

$$\underset{\mathbf{x}}{\text{minimize}} \sum_{i=1}^m f_i(\mathbf{x}) + r(\mathbf{x}),$$

in which the main objective is to minimize a sum of cost functions f_1, \dots, f_m , subject to some additional regularization r on \mathbf{x} . To solve it using ADMM, we first reformulate it into a *consensus* form by introducing m auxiliary variables $\mathbf{z}_1, \dots, \mathbf{z}_m$, as

$$\begin{aligned} \underset{\mathbf{x}, \{\mathbf{z}_i\}_{i=1}^m}{\text{minimize}} & \sum_{i=1}^m f_i(\mathbf{z}_i) + r(\mathbf{x}), \\ \text{subject to} & \mathbf{z}_i = \mathbf{x}, \forall i = 1, \dots, m. \end{aligned}$$

Then we can easily write down the (scaled-form) ADMM iterates for this problem as

$$\begin{aligned} \mathbf{x} & \leftarrow \arg \min_{\mathbf{x}} r(\mathbf{x}) + \rho \sum_{i=1}^m \|\mathbf{z}_i - \mathbf{x} + \mathbf{u}_i\|^2, \\ \mathbf{z}_i & \leftarrow \arg \min_{\mathbf{z}_i} f_i(\mathbf{z}_i) + \rho \|\mathbf{z}_i - \mathbf{x} + \mathbf{u}_i\|^2, \forall i = 1, \dots, m, \\ \mathbf{u}_i & \leftarrow \mathbf{u}_i + \mathbf{z}_i - \mathbf{x}, \forall i = 1, \dots, m, \end{aligned}$$

where \mathbf{x} is treated as the first block, the set of auxiliary variables $\{\mathbf{z}_i\}$ are treated as the second block, and \mathbf{u}_i is the scaled dual variable corresponding to the equality constraint $\mathbf{z}_i = \mathbf{x}$.

There are several advantages of this consensus-ADMM algorithm. First and foremost, it is designed for distributed optimization, since each \mathbf{z}_i can be updated in parallel; through careful splitting of the sum of the cost functions, we can also make each update very efficient (possibly in closed-form), which may not be the case for the batch problem. Finally, since it falls into the general algorithmic framework of ADMM, it converges as long as the problem is convex, for all $\rho > 0$.

II. GENERAL ALGORITHMIC FRAMEWORK

We now describe how to apply consensus-ADMM for general QCQPs. Let us first transform (1) into a consensus form

$$\begin{aligned} & \underset{\mathbf{x}, \{\mathbf{z}_i\}_{i=1}^m}{\text{minimize}} && \mathbf{x}^H \mathbf{A}_0 \mathbf{x} - 2\Re \{ \mathbf{b}_0^H \mathbf{x} \}, \\ & \text{subject to} && \mathbf{z}_i^H \mathbf{A}_i \mathbf{z}_i - 2\Re \{ \mathbf{b}_i^H \mathbf{z}_i \} \leq c_i, \\ & && \mathbf{z}_i = \mathbf{x}, \forall i = 1, \dots, m, \end{aligned} \quad (3)$$

then the corresponding consensus-ADMM algorithm takes the form of the following iterations:

$$\begin{aligned} \mathbf{x} & \leftarrow (\mathbf{A}_0 + m\rho\mathbf{I})^{-1} \left(\mathbf{b}_0 + \rho \sum_{i=1}^m (\mathbf{z}_i + \mathbf{u}_i) \right), \\ \mathbf{z}_i & \leftarrow \arg \min_{\mathbf{z}_i} \|\mathbf{z}_i - \mathbf{x} + \mathbf{u}_i\|^2, \\ & \text{subject to } \mathbf{z}_i^H \mathbf{A}_i \mathbf{z}_i - 2\Re \{ \mathbf{b}_i^H \mathbf{z}_i \} \leq c_i \\ \mathbf{u}_i & \leftarrow \mathbf{u}_i + \mathbf{z}_i - \mathbf{x}, \end{aligned} \quad (4)$$

assuming the updates of \mathbf{z}_i exist (see the next section for more discussions on the existence of the \mathbf{z}_i -updates).

The reason we put our algorithm into this form is based on the fact that each update of \mathbf{z}_i is a QCQP-1, thus we know it can be updated optimally, despite the fact that the quadratics may be indefinite. The update for \mathbf{x} is an unconstrained quadratic minimization, and for an indefinite \mathbf{A}_0 we need to choose a large enough ρ to ensure that the minimum is not unbounded; if $\mathbf{A}_0 + \rho m\mathbf{I} \succeq 0$ is satisfied then the solution is simply given by solving a linear equation, and we can cache the Cholesky factorization of $\mathbf{A}_0 + m\rho\mathbf{I}$ to save computations in the subsequent iterations.

A. Convergence

ADMM was first designed for convex problems, for which it is known to converge under mild conditions [15]. Despite the lack of theoretical guarantees, ADMM has also been used for non-convex problems, see for example [14, §9]. In [16], ADMM was applied to non-negative matrix factorization (a non-convex problem) with missing values, and it was shown that, *if* ADMM converges for this non-convex problem, *then* it converges to a KKT point. Some follow-up works on other non-convex problems have made similar claims [17], [18]. A stronger result was recently shown in [19], where it was proven that (there exists a convergent subsequence and) *every limit point is a stationary point* for a class of non-convex consensus and sharing problems. The proof in [19] assumes Lipschitz continuity of the non-convex cost functions (not constraints) to establish that the augmented Lagrangian function is non-increasing, provided the parameter ρ is large enough.

Unfortunately, the convergence result in [19] cannot be applied to our algorithm here, even though both are dealing with non-convex *consensus* problems. The very first step in the proof of [19] shows that the augmented Lagrangian is monotonically non-increasing under certain conditions. These conditions include Lipschitz continuity of the non-convex cost functions (but not the constraints) and that the parameter ρ is large enough. If we want to borrow the arguments made in [19], we would need

to first establish the monotonicity of the augmented Lagrangian. However, our numerical experience is that the augmented Lagrangian is not monotonic, even if we set ρ to be very large. Therefore, we limit ourselves to the following weaker convergence result.

Theorem 1: Denote \mathbf{x}^t and \mathbf{z}_i^t the updates obtained at the t -th iteration of Algorithm (4). Assume that the \mathbf{z}_i^t 's are well-defined for all t and i , and that

$$\lim_{t \rightarrow +\infty} (\mathbf{z}_i^t - \mathbf{x}^t) = 0, \forall i = 1, \dots, m,$$

and

$$\lim_{t \rightarrow +\infty} (\mathbf{x}^{t+1} - \mathbf{x}^t) = 0,$$

then any limit point of $\{\mathbf{x}^t\}$ is a KKT point of (1)

Proof: See Appendix A. ■

III. EFFICIENT \mathbf{z}_i -UPDATES

Now let us focus on the update of \mathbf{z}_i . From our previous discussion on QCQP-1 we know that the update of \mathbf{z}_i can always be solved to optimality, by strong duality; in other words, if we solve the SDR of a QCQP-1, we are guaranteed to obtain a rank one solution. However, with a number of \mathbf{z}_i to be updated iteratively, it is not desirable to rely on general SDP algorithms to update \mathbf{z}_i . Therefore, we now take a detailed look into QCQP-1, and show how to solve it efficiently. For ease of notation, let us drop the subscript, define $\boldsymbol{\zeta} = \mathbf{x} - \mathbf{u}$, and denote the sub-problem as

$$\begin{aligned} & \underset{\mathbf{z}}{\text{minimize}} && \|\mathbf{z} - \boldsymbol{\zeta}\|^2 \\ & \text{subject to} && \mathbf{z}^H \mathbf{A} \mathbf{z} - 2\Re \{ \mathbf{b}^H \mathbf{z} \} = c \end{aligned} \quad (5)$$

We changed the constraint to equality here to simplify subsequent derivations. For an inequality constraint, we first check whether $\boldsymbol{\zeta}$ is feasible: if yes, then $\boldsymbol{\zeta}$ is the solution; if not, then the constraint must be satisfied as equality, according to complementary slackness, thus the following method can be applied.

We start from simpler cases, and gradually build up to the most general case.

A. $\text{rank}(\mathbf{A}) = 1, \mathbf{b} = \mathbf{0}$

For this simple case, the constraint can be equivalently written as

$$|\mathbf{a}^H \mathbf{z}|^2 = c,$$

or simply as a linear constraint with an unknown phase

$$\mathbf{a}^H \mathbf{z} = \sqrt{c} e^{j\theta}.$$

Assuming we know θ , problem (5) becomes a simple projection onto an affine subspace, for which we know the solution is given by

$$\mathbf{z} = \boldsymbol{\zeta} + \left(\frac{\sqrt{c} e^{j\theta} - \mathbf{a}^H \boldsymbol{\zeta}}{\|\mathbf{a}\|^2} \right) \mathbf{a}.$$

Plugging this back to the objective, it is easy to see that the minimum is attained if we choose θ to be the angle of $\mathbf{a}^H \boldsymbol{\zeta}$.

Therefore, the update of \mathbf{z} in this case is given by

$$\mathbf{z} = \zeta + \frac{\sqrt{c - |\mathbf{a}^H \zeta|}}{\|\mathbf{a}\|^2 |\mathbf{a}^H \zeta|} \mathbf{a} \mathbf{a}^H \zeta. \quad (6)$$

In the real case, the unknown phase becomes an unknown sign, and similar steps can be made to result in the same closed-form solution (6).

B. $\text{rank}(\mathbf{A}) > 1, \mathbf{b} = \mathbf{0}$

For \mathbf{A} with a higher rank, there is in general no closed-form solution for (5). However, it is still possible to efficiently update \mathbf{z} . Let the eigen-decomposition of \mathbf{A} be $\mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^H$, where $\mathbf{\Lambda}$ is diagonal real and \mathbf{Q} is unitary, because \mathbf{A} is Hermitian. Define $\tilde{\mathbf{z}} = \mathbf{Q}^H \mathbf{z}$, $\tilde{\zeta} = \mathbf{Q}^H \zeta$, then the problem is equivalent to

$$\begin{aligned} & \underset{\tilde{\mathbf{z}}}{\text{minimize}} \quad \|\tilde{\mathbf{z}} - \tilde{\zeta}\|^2, \\ & \text{subject to} \quad \tilde{\mathbf{z}}^H \mathbf{\Lambda} \tilde{\mathbf{z}} = c. \end{aligned}$$

The corresponding Lagrangian is

$$L = \|\tilde{\mathbf{z}} - \tilde{\zeta}\|^2 + \mu (\tilde{\mathbf{z}}^H \mathbf{\Lambda} \tilde{\mathbf{z}} - c),$$

with a single Lagrange multiplier μ . A necessary condition for optimality is that $\nabla L = 0$, i.e.,

$$\nabla L = 2(\tilde{\mathbf{z}} - \tilde{\zeta}) + 2\mu\mathbf{\Lambda}\tilde{\mathbf{z}} = 0,$$

therefore, $\tilde{\mathbf{z}} = (\mathbf{I} + \mu\mathbf{\Lambda})^{-1} \tilde{\zeta}$. Plugging this solution back into the equality constraint, we have

$$\tilde{\zeta}^H (\mathbf{I} + \mu\mathbf{\Lambda})^{-1} \mathbf{\Lambda} (\mathbf{I} + \mu\mathbf{\Lambda})^{-1} \tilde{\zeta} = c,$$

or equivalently

$$\sum_{k=1}^n \frac{\lambda_k}{(1 + \mu\lambda_k)^2} |\tilde{\zeta}_k|^2 = c,$$

which means the correct Lagrange multiplier μ can be numerically found by solving this nonlinear equation, via for example bisection or Newton's method. In fact, we can also show that the desired solution is unique, leaving no ambiguity to the value of μ . From the dual of a QCQP-1, we have that $\mathbf{I} + \mu\mathbf{\Lambda} \succeq \mathbf{0}$ [2, Appendix B], i.e.,

$$1 + \mu\lambda_k \geq 0, \forall k = 1, \dots, n.$$

This can give us a first possible region where the correct μ can be: $\mu \leq -1/\lambda_{\min}$ if $\lambda_{\min} < 0$, and $\mu \geq -1/\lambda_{\max}$ if $\lambda_{\max} > 0$.

Moreover, if we define

$$\phi(\mu) = \sum_{k=1}^n \frac{\lambda_k}{(1 + \mu\lambda_k)^2} |\tilde{\zeta}_k|^2 - c,$$

then

$$\phi'(\mu) = -2 \sum_{k=1}^n \frac{\lambda_k^2}{(1 + \mu\lambda_k)^3} |\tilde{\zeta}_k|^2,$$

and for all μ such that $\mathbf{I} + \mu\mathbf{\Lambda} \succeq \mathbf{0}$, $\phi'(\mu) < 0$, which means $\phi(\mu)$ is monotonically decreasing (strictly) in that region, therefore the solution for $\phi(\mu) = 0$ is unique. In fact, we can show that there exists a root within that interval, as long as the

constraint set is not empty: if $-1/\lambda_{\max} \leq \mu \leq -1/\lambda_{\min}$, then $\phi(-1/\lambda_{\max}) = +\infty$, and $\phi(-1/\lambda_{\min}) = -\infty$, which together with the monotonicity imply that a root always exists in between. If $\mathbf{\Lambda} \succeq \mathbf{0}$, the interval becomes $-1/\lambda_{\max} \leq \mu \leq +\infty$, consequently $-c \leq \phi(\mu) \leq +\infty$, so a root exists if and only if $c \geq 0$, but if $c < 0$ then $\tilde{\mathbf{z}}^T \mathbf{\Lambda} \tilde{\mathbf{z}} = c$ is infeasible. A similar argument applies to the case when $\mathbf{\Lambda} \preceq \mathbf{0}$. Once the value of μ is found, we can plug it back to obtain $\tilde{\mathbf{z}}$, and the desired update of \mathbf{z} is simply given by $\mathbf{z} = \mathbf{Q}\tilde{\mathbf{z}}$.

To save computation, we can cache the eigen-decomposition of \mathbf{A} . Then in the subsequent ADMM iterations the computation is dominated by the matrix-vector multiplication $\mathbf{Q}\tilde{\mathbf{z}}$, since evaluating either $\phi(\mu)$ or $\phi'(\mu)$ (if Newton's method is used) only takes $O(n)$ complexity.

C. $\text{rank}(\mathbf{A}) > 1, \mathbf{b} \neq \mathbf{0}$

Now we have reached the most general case when \mathbf{A} can have higher rank and \mathbf{b} can be non-zero. The idea is very similar to the previous case, although the expressions are a little more complicated. Again let $\mathbf{A} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^H$ be the eigen-decomposition, problem (5) is equivalent to

$$\begin{aligned} & \underset{\tilde{\mathbf{z}}}{\text{minimize}} \quad \|\tilde{\mathbf{z}} - \tilde{\zeta}\|^2, \\ & \text{subject to} \quad \tilde{\mathbf{z}}^H \mathbf{\Lambda} \tilde{\mathbf{z}} - 2\Re\{\tilde{\mathbf{b}}^H \tilde{\mathbf{z}}\} = c, \end{aligned}$$

Where $\tilde{\mathbf{z}} = \mathbf{Q}^H \mathbf{z}$, $\tilde{\zeta} = \mathbf{Q}^H \zeta$, and $\tilde{\mathbf{b}} = \mathbf{Q}^H \mathbf{b}$. Setting the gradient of the Lagrangian equal to zero, we have

$$\tilde{\mathbf{z}} = (\mathbf{I} + \mu\mathbf{\Lambda})^{-1} (\tilde{\zeta} + \mu\tilde{\mathbf{b}}).$$

Plugging it back to the equality constraint, it becomes a nonlinear equation with respect to μ ,

$$\phi(\mu) = \sum_{k=1}^n \lambda_k \left| \frac{\tilde{\zeta}_k + \mu\tilde{b}_k}{1 + \mu\lambda_k} \right|^2 - 2\Re \left\{ \sum_{k=1}^n \tilde{b}_k^* \frac{\tilde{\zeta}_k + \mu\tilde{b}_k}{1 + \mu\lambda_k} \right\} - c,$$

and its derivative

$$\phi'(\mu) = -2 \sum_{k=1}^n \frac{|\tilde{b}_k - \lambda_k \tilde{\zeta}_k|^2}{(1 + \mu\lambda_k)^3} < 0,$$

for all μ such that $\mathbf{I} + \mu\mathbf{\Lambda} \succeq \mathbf{0}$, which is necessary for optimality of (5). Therefore, $\phi(\mu)$ is monotonic in the possible region of solution, and any local solution (for example found by bisection or Newton's method) is guaranteed to be the unique (thus correct) solution, which always exists for a non-empty constraint set, similar to the previous case. Notice that if $\mathbf{b} = \mathbf{0}$, $\phi(\mu)$ and $\phi'(\mu)$ reduce to the simpler expression that we derived in the previous subsection. Detailed implementation of bisection and Newton's method to solve $\phi(\mu) = 0$ is given in Algorithms 1 and 2. In practice, bisection converges linearly (~ 20 iterations) while Newton's method converges quadratically (~ 5 iterations), but bisection is numerically more stable, so the best choice is application-specific.

An interesting observation from this most general case is that, solving a QCQP-1 always boils down to solving a scalar nonlinear equation $\phi(\mu) = 0$. It is easy to see that if \mathbf{A} has p distinct eigenvalues, solving $\phi(\mu) = 0$ is equivalent to solving

Algorithm 1: Solving $\phi(\mu) = 0$ using bisection.

```

1  $\lceil \mu \rceil \leftarrow$  max number available in float representation;
2  $\lfloor \mu \rfloor \leftarrow$  min number available in float representation;
3 if  $\lambda_{\min} < 0$  then  $\lceil \mu \rceil \leftarrow -1/\lambda_{\min}$ ;
4 if  $\lambda_{\max} > 0$  then  $\lfloor \mu \rfloor \leftarrow -1/\lambda_{\max}$ ;
5 repeat
6    $\mu = (\lceil \mu \rceil + \lfloor \mu \rfloor)/2$ ;
7   if  $\phi(\mu) > 0$  then  $\lfloor \mu \rfloor \leftarrow \mu$ ;
8   else  $\lceil \mu \rceil \leftarrow \mu$ ;
9 until  $\lceil \mu \rceil - \lfloor \mu \rfloor < \varepsilon$ ;
10 return  $\mu = (\lceil \mu \rceil + \lfloor \mu \rfloor)/2$ 

```

Algorithm 2: Solving $\phi(\mu) = 0$ using Newton's method.

```

1  $\mu \leftarrow -(\lambda_{\min} + \lambda_{\max})/2\lambda_{\min}\lambda_{\max}$ ;
2 repeat
3    $\mu = \mu - \phi(\mu)/\phi'(\mu)$ ;
4 until  $-\phi(\mu)^2/\phi'(\mu) < \varepsilon$ ;
5 return  $\mu$ 

```

a polynomial of degree $2p + 1$ ($2p$ if $\mathbf{b} = \mathbf{0}$). Polynomials of order ≥ 5 do not admit closed-form expressions for their roots, necessitating the use of numerical methods like bisection or Newton's method.

Remark: So far we have assumed that $\mathbf{I} + \mu\mathbf{A}$ is invertible, which may not always be the case. However, recall that the duality of QCQP-1 implies $\mathbf{I} + \mu\mathbf{A} \geq 0$, therefore there are at most two possible values of μ that can make the matrix $\mathbf{I} + \mu\mathbf{A}$ singular: $\mu = -1/\lambda_{\min}$ if $\lambda_{\min} < 0$, and $\mu = -1/\lambda_{\max}$ if $\lambda_{\max} > 0$, so for completeness one may first check these two values of μ , although this situation never occurred in our experiments.

D. Bound Constraint

The basic idea of making the \mathbf{z} -updates equivalent to solving a QCQP-1 is that the latter is always efficiently solvable. In some cases this efficiency can be maintained even if we incorporate some more constraints. One such case is that of a quadratic term that is bounded from both sides, i.e.,

$$\begin{aligned} & \underset{\mathbf{z}}{\text{minimize}} \quad \|\mathbf{z} - \boldsymbol{\zeta}\|^2, \\ & \text{subject to} \quad c - \epsilon \leq \mathbf{z}^H \mathbf{A} \mathbf{z} - 2\Re\{\mathbf{b}^H \mathbf{z}\} \leq c + \epsilon. \end{aligned}$$

Using the same idea as before, we can write down the Lagrangian and the solution again takes the form

$$\mathbf{z} = \mathbf{Q}(\mathbf{I} + \mu\mathbf{A})^{-1} \mathbf{Q}^H (\boldsymbol{\zeta} + \mu\mathbf{b}),$$

where $\mathbf{A} = \mathbf{Q}\boldsymbol{\Lambda}\mathbf{Q}^H$ is the eigen-decomposition of \mathbf{A} , and μ is such that

$$\begin{cases} \mu = 0, & \text{if } -\epsilon \leq \phi(0) \leq \epsilon, \\ \mu < 0, & \text{then } \phi(\mu) = -\epsilon, \\ \mu > 0, & \text{then } \phi(\mu) = \epsilon. \end{cases}$$

In fact, since we know $\phi(\mu)$ is monotonically decreasing within the feasible region, if $\phi(0) > \epsilon$, both the solution of $\phi(\mu) = \pm\epsilon$

are positive, therefore we must take the solution of $\phi(\mu) = \epsilon$, and vice versa. In other words, if

$$c - \epsilon \leq \boldsymbol{\zeta}^H \mathbf{A} \boldsymbol{\zeta} - 2\Re\{\mathbf{b}^H \boldsymbol{\zeta}\} \leq c + \epsilon,$$

then $\mathbf{z} = \boldsymbol{\zeta}$; if it is greater than $c + \epsilon$, the upper-bound constraint must be active, and like-wise if it is less than $c - \epsilon$. This is very intuitive, since we are just ‘‘rounding’’ the constraint to the closest bound.

IV. IMPLEMENTATION ISSUES

So far we have derived an ADMM algorithm for general QCQP problems, which features straightforward iterations and efficient per-iteration updates. In this section we revisit the entire algorithm and discuss detailed implementations to make it more actionable.

A. Memory-Efficient Implementation

An apparent disadvantage of our algorithm is that we need to introduce an auxiliary variable \mathbf{z}_i and the corresponding dual variable \mathbf{u}_i for every single quadratic constraint. For $\mathbf{x} \in \mathcal{C}^n$ and m such constraints, we need $\mathcal{O}(mn)$ memory just to store the intermediate variables. Depending on the application, this memory requirement may be too demanding. For example, if $\mathbf{A}_i = \mathbf{a}_i \mathbf{a}_i^H, \forall i = 1, \dots, m$, it only takes $\mathcal{O}(mn)$ memory to describe the problem, or even as small as $\mathcal{O}(m)$ if the \mathbf{a}_i 's are highly structured, e.g., obtained from the rows of the discrete Fourier transform (DFT) matrix. In such cases $\mathcal{O}(mn)$ intermediate memory seems very unappealing for large m and n . This is less of an issue when the \mathbf{A}_i 's are all full rank, since then we need $\mathcal{O}(mn^2)$ memory to just specify the problem, and if that is affordable, then $\mathcal{O}(mn)$ memory for intermediate variables seems relatively reasonable.

Consider the following special QCQP, which occurs frequently in practice:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} \quad \mathbf{x}^H \mathbf{A}_0 \mathbf{x} - 2\Re\{\mathbf{b}_0^H \mathbf{x}\} \\ & \text{subject to} \quad |\mathbf{a}_i^H \mathbf{x}|^2 = c_i, \forall i = 1, \dots, m. \end{aligned} \quad (7)$$

Again, the algorithm that follows can be easily modified to tackle inequality constraints or bound constraints, but we start with equality constraints here for clarity. According to our previous discussion, we can write down explicitly the consensus-ADMM iterations as

$$\begin{aligned} \mathbf{x} & \leftarrow (\mathbf{A}_0 + m\rho\mathbf{I})^{-1} \left(\mathbf{b}_0 + \rho \sum_{i=1}^m (\mathbf{z}_i + \mathbf{u}_i) \right), \\ \mathbf{z}_i & \leftarrow \mathbf{x} - \mathbf{u}_i + \frac{\sqrt{c_i} - |\mathbf{a}_i^H (\mathbf{x} - \mathbf{u}_i)|}{\|\mathbf{a}_i\|^2 |\mathbf{a}_i^H (\mathbf{x} - \mathbf{u}_i)|} \mathbf{a}_i \mathbf{a}_i^H (\mathbf{x} - \mathbf{u}_i), \\ \mathbf{u}_i & \leftarrow \mathbf{u}_i + \mathbf{z}_i - \mathbf{x}. \end{aligned} \quad (8)$$

Define $\mathbf{z}_s = \sum_i \mathbf{z}_i$ and $\mathbf{u}_s = \sum_i \mathbf{u}_i$, then we can simplify the algorithm as

$$\begin{aligned}\mathbf{x} &\leftarrow (\mathbf{A}_0 + m\rho\mathbf{I})^{-1}(\mathbf{b}_0 + \mathbf{z}_s + \mathbf{u}_s), \\ \mathbf{z}_s &\leftarrow m\mathbf{x} - \mathbf{u}_s + \mathbf{A}_s\boldsymbol{\nu}, \\ \mathbf{u}_s &\leftarrow \mathbf{u}_s + \mathbf{z}_s - m\mathbf{x},\end{aligned}$$

where $\mathbf{A}_s = [\mathbf{a}_1 \mathbf{a}_2 \dots \mathbf{a}_m]$ is a $n \times m$ matrix formed by parallel stacking all the \mathbf{a}_i vectors as its columns, and $\boldsymbol{\nu}$ is a vector of length m with its elements defined as

$$\nu_i = \frac{\sqrt{c_i} - |\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|}{\|\mathbf{a}_i\|^2 |\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|} \mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i).$$

If we are given the vector $\boldsymbol{\nu}$ at every iteration, then we can simply work with the summation of the local variables \mathbf{z}_s and \mathbf{u}_s without the possible memory explosion. To compute the vector $\boldsymbol{\nu}$, we notice that it is sufficient to know the value of $\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)$ for each ν_i , and since we keep track of \mathbf{x} explicitly, the only difficulty is to keep track of $\mathbf{a}_i^H \mathbf{u}_i$ without the actual value of \mathbf{u}_i . At the end of each iteration, by combining the updates of \mathbf{z}_i and \mathbf{u}_i , we have that

$$\mathbf{u}_i \leftarrow \frac{\sqrt{c_i} - |\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|}{\|\mathbf{a}_i\|^2 |\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|} \mathbf{a}_i \mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i),$$

therefore

$$\mathbf{a}_i^H \mathbf{u}_i \leftarrow \frac{\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)}{|\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|} (\sqrt{c_i} - |\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|).$$

Now if we define $\alpha_i = \mathbf{a}_i^H \mathbf{u}_i$, it is apparent that we can update α_i iteratively as

$$\alpha_i \leftarrow \frac{\mathbf{a}_i^H \mathbf{x} - \alpha_i}{|\mathbf{a}_i^H \mathbf{x} - \alpha_i|} (\sqrt{c_i} - |\mathbf{a}_i^H \mathbf{x} - \alpha_i|).$$

To sum up, a memory-efficient way to implement consensus-ADMM for problem (7) takes the form

$$\begin{aligned}\mathbf{x} &\leftarrow (\mathbf{A}_0 + m\rho\mathbf{I})^{-1}(\mathbf{b}_0 + \rho(\mathbf{z}_s + \mathbf{u}_s)), \\ \boldsymbol{\xi} &\leftarrow \mathbf{A}_s^H \mathbf{x}, \\ \nu_i &\leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} \frac{\sqrt{c_i} - |\xi_i - \alpha_i|}{\|\mathbf{a}_i\|^2}, \\ \mathbf{z}_s &\leftarrow m\mathbf{x} - \mathbf{u}_s + \mathbf{A}_s \boldsymbol{\nu}, \\ \mathbf{u}_s &\leftarrow \mathbf{u}_s + \mathbf{z}_s - m\mathbf{x}, \\ \alpha_i &\leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} (\sqrt{c_i} - |\xi_i - \alpha_i|).\end{aligned}\quad (9)$$

The explicit variables are $\mathbf{x}, \mathbf{z}_s, \mathbf{u}_s \in \mathcal{C}^n$ and $\boldsymbol{\xi}, \boldsymbol{\nu}, \boldsymbol{\alpha} \in \mathcal{C}^m$, so the total memory consumption now is $\mathcal{O}(m+n)$, compared to $\mathcal{O}(mn)$ in the original form.

Finally, we show that the modified iterates in (9) can handle some variations in the constraints. Suppose the i -th constraint is an inequality $|\mathbf{a}_i^H \mathbf{x}|^2 \leq c_i$, then the update of \mathbf{z}_i in (8) should be

$$\mathbf{z}_i \leftarrow \begin{cases} \mathbf{x} - \mathbf{u}_i, & \text{if } |\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|^2 \leq c_i \\ \mathbf{x} - \mathbf{u}_i + \frac{\sqrt{c_i} - |\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|}{\|\mathbf{a}_i\|^2 |\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|} \mathbf{a}_i \mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i), & \text{otherwise,} \end{cases}$$

or simply

$$\mathbf{z}_i \leftarrow \mathbf{x} - \mathbf{u}_i + \frac{[\sqrt{c_i} - |\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|]_-}{\|\mathbf{a}_i\|^2 |\mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i)|} \mathbf{a}_i \mathbf{a}_i^H(\mathbf{x} - \mathbf{u}_i).$$

This means the corresponding ν_i and α_i updates can be similarly modified as

$$\begin{aligned}\nu_i &\leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} \frac{[\sqrt{c_i} - |\xi_i - \alpha_i|]_-}{\|\mathbf{a}_i\|^2}, \\ \alpha_i &\leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} [\sqrt{c_i} - |\xi_i - \alpha_i|]_-, \end{aligned}$$

and the rest of the updates in (9) stays the same. Conversely, if the constraint is a \geq inequality, we only keep the nonnegative part of $\sqrt{c_i} - |\xi_i - \alpha_i|$ in the updates of ν_i and α_i . If it is a bound constraint

$$c_i - \epsilon \leq |\mathbf{a}_i^H \mathbf{x}|^2 \leq c_i + \epsilon,$$

according to our previous discussion on ‘‘rounding’’ for this kind of constraint, we can define τ_i as

$$\tau_i \leftarrow \begin{cases} \sqrt{c_i - \epsilon} - |\xi_i - \alpha_i|, & \text{if } |\xi_i - \alpha_i|^2 < c_i - \epsilon, \\ \sqrt{c_i + \epsilon} - |\xi_i - \alpha_i|, & \text{if } |\xi_i - \alpha_i|^2 > c_i + \epsilon, \\ 0, & \text{otherwise.} \end{cases}$$

Then the corresponding updates of ν_i and α_i are

$$\begin{aligned}\nu_i &\leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} \frac{\tau_i}{\|\mathbf{a}_i\|^2}, \\ \alpha_i &\leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} \tau_i.\end{aligned}$$

As we will see later, this type of memory efficient implementation can even be extended to cases when the constraints are not exactly homogeneous rank one quadratics. Furthermore, recall our previous discussion that a homogeneous rank one quadratic constraint is simply a linear constraint with an unknown phase (or sign in the real case), implying that if we have actual linear constraints in the QCQP problem, a similar idea can also be applied to avoid explicitly introducing a huge number of auxiliary variables, while still maintaining the simplicity of the updates.

B. Initialization and Parameter Setting

At this point we need to remind the reader that, although the consensus ADMM algorithm we derived for non-convex QCQPs has an appealing form and cheap per-iteration complexity, it is after all a heuristic for what is in general an NP-hard problem. We may then anticipate that appropriate initialization and judicious parameter tuning will be more important than in standard applications of ADMM to convex problems. Nevertheless, we have devised practical rules that seem to work well in most cases, as discussed below.

The only parameter that needs to be tuned is ρ , which is only involved in the update of \mathbf{x} if we have an explicit objective. Clearly a smaller ρ steers the \mathbf{x} -update towards putting more emphasis on decreasing the cost function, whereas a bigger ρ puts more weight on agreeing with the auxiliary variables $\{\mathbf{z}_i\}$,

each guaranteed to satisfy one constraint. We found empirically that if we start with a feasible \mathbf{x} , then we can afford to have a relatively small value of ρ for faster decrease of the cost, while preventing \mathbf{x} from diverging towards infeasibility.

How can we find a feasible point for initialization? In some cases it is easy, for example when all the \mathbf{A}_i 's are positive semi-definite, and all the inequality constraints are homogeneous and one-sided, then a simple scaling suffices to make an arbitrary point feasible. In general, finding a feasible point is also NP-hard. In our context, we can attempt to find a feasible point by using the same consensus ADMM algorithm for the following feasibility problem, rewritten in the consensus form

$$\begin{aligned} & \text{find } \mathbf{x}, \{\mathbf{z}_i\}_{i=1}^m, \\ & \text{such that } \mathbf{z}_i^H \mathbf{A}_i \mathbf{z}_i - 2\Re\{\mathbf{b}_i^H \mathbf{z}_i\} \leq c_i, \\ & \mathbf{z}_i = \mathbf{x}, \forall i = 1, \dots, m, \end{aligned}$$

Applying consensus ADMM, we obtain the following updates

$$\begin{aligned} \mathbf{x} & \leftarrow \frac{1}{m} \sum_{i=1}^m (\mathbf{z}_i + \mathbf{u}_i), \\ \mathbf{z}_i & \leftarrow \arg \min_{\mathbf{z}_i} \|\mathbf{z}_i - \mathbf{x} + \mathbf{u}_i\|^2, \\ & \text{subject to } \mathbf{z}_i^H \mathbf{A}_i \mathbf{z}_i - 2\Re\{\mathbf{b}_i^H \mathbf{z}_i\} \leq c_i \\ \mathbf{u}_i & \leftarrow \mathbf{u}_i + \mathbf{z}_i - \mathbf{x}, \end{aligned}$$

which are completely independent of ρ ¹. This type of iterates with random initialization usually converges much faster in finding a feasible point, if one exists. The result can then serve as initialization for subsequent ADMM updates with the cost function brought back into consideration. If the first phase fails to find a feasible point even after multiple trials, then we have, to some extent, numerical evidence that the problem may be infeasible, and for practical purposes there is often no point in proceeding further anyway.

C. Parallel and Distributed Implementation

Consensus ADMM is by its very nature highly parallelizable, since the update of each auxiliary variable \mathbf{z}_i is independent of all others. This nice property is thankfully maintained even in the squeezed form (9), since all the operations involved are element-wise, except for two matrix vector multiplications $\mathbf{A}_s^H \mathbf{x}$ and $\mathbf{A}_s \mathbf{v}$, which can also be parallelized easily. This means that the proposed algorithm can easily achieve p -fold acceleration by using p processors on a shared-memory system.

A more interesting case is when a large amount of data is stored in distributed storage, and different agents need to coordinate with a master node with small communication overheads. Suppose the data for the constraints $\{\mathbf{A}_i, \mathbf{b}_i\}$ are stored across p agents, all connected to the master node which is in charge of the cost function. Since we assign each constraint an individual variable \mathbf{z}_i and dual \mathbf{u}_i , suppose the k -th agent is in charge of m_k constraints, a naive implementation would require the j -th agent to send m_j of the \mathbf{z}_i 's and \mathbf{u}_i 's to the central node in

each iteration. This is not necessary, as a matter of fact, since for the update of \mathbf{x} only the sum of all the \mathbf{z}_i 's and the \mathbf{u}_i 's is required. Therefore, to minimize communication overheads, the j -th agent can simply define $\mathbf{x}_j = \sum_{i \in \Omega_j} (\mathbf{z}_i + \mathbf{u}_i)$, where Ω_j is the index set of the constraints handled by the j -th agent. At the master node, another summation over all the \mathbf{x}_j 's is carried out for the exact update of \mathbf{x} .

V. APPLICATIONS AND NUMERICAL EXPERIMENTS

So far we have introduced the general idea of applying consensus ADMM to QCQPs with efficient per-iteration updates, and explored memory-efficient and parallel/distributed implementation issues. In this section, we will look into some important QCQP applications, write down explicitly the algorithm, and compare its numerical performance with some state-of-the-art algorithms. All simulations were performed in MATLAB on a Linux desktop with 8 Intel i7 cores and 32 GB of RAM.

A. Feasible Point Pursuit

One of the main drawbacks of the SDR approach for non-convex QCQPs is that when it is not obvious how to find a feasible point that satisfies the constraint set, there is a high chance that SDR, followed by taking the principal component and/or Gaussian randomization, will not satisfy all the constraints either. Recently, a new algorithm called FPP-SCA [11] was proposed to address this issue by iteratively linearizing the non-convex part of the problem, while adding nonnegative slacks to each constraint and penalizing the sum of slacks in the cost function as well. Simulations in [11] suggest that FPP-SCA works well with high probability, even when SDR fails.

Consensus ADMM can also be used to find feasible points, and it is possible to aim it towards finding a feasible point having smallest ℓ_2 norm. Instead of giving each constraint a slack and trying to minimize the sum of the slacks, consensus ADMM gives each constraint a local variable and tries to drive these local variables to consensus. Explicitly, let us consider the following problem:

$$\begin{aligned} & \underset{\mathbf{x} \in \mathcal{C}^n}{\text{minimize}} \quad \|\mathbf{x}\|^2 \\ & \text{subject to} \quad \mathbf{x}^H \mathbf{A}_i \mathbf{x} \leq c_i, \forall i = 1, \dots, m, \end{aligned} \quad (10)$$

where $\mathbf{A}_1, \dots, \mathbf{A}_m$ are in general Hermitian indefinite and full rank. Following our discussion in Section III-B, the detailed consensus ADMM algorithm for (10) is given in Algorithm 3, where we have applied the two stage approach described in Section IV-B: we attempt to find a feasible point in the first phase, followed by stably decreasing its norm in the second phase. We found empirically that simply setting $\rho = 1$ works very well for the second phase in this context.

Now let us compare consensus-ADMM with FPP-SCA on some synthetically generated problems. After fixing the problem dimension n and m , we first generate $\mathbf{x}_{\text{feas}} \sim \mathcal{CN}(0, \mathbf{I})$. A Hermitian indefinite matrix \mathbf{A}_i is generated by first randomly drawing a $n \times n$ matrix from $\mathcal{CN}(0, 1)$, and then taking the average of its Hermitian and itself. The corresponding c_i is set to be $\mathbf{x}_{\text{feas}}^H \mathbf{A}_i \mathbf{x}_{\text{feas}} - |v_i|$ where v_i is randomly generated

¹Or one can interpret this as ρ being $+\infty$.

Algorithm 3: Consensus-ADMM for (10).

```

1 initialize  $\mathbf{x}$ ,  $\mathbf{z}_i$  and  $\mathbf{u}_i$ ;
2 for  $i = 1, \dots, m$  do
3   Take the eigen-decomposition of  $\mathbf{A}_i = \mathbf{Q}_i \mathbf{\Lambda}_i \mathbf{Q}_i^H$ ;
4 end
5 repeat
6    $\mathbf{x} \leftarrow \frac{1}{m} \sum_{i=1}^m (\mathbf{z}_i + \mathbf{u}_i)$ ;
7   for  $i = 1, \dots, m$  do
8      $\mu_i \leftarrow \arg_{\mu} \phi_i(\mu) = 0$  using Alg. 1 or 2;
9      $\mu_i \leftarrow \max\{0, \mu_i\}$ ; // for  $\leq$  constraints
10     $\mathbf{z}_i \leftarrow \mathbf{Q}_i (\mathbf{I} + \mu_i \mathbf{\Lambda}_i)^{-1} \mathbf{Q}_i^H (\mathbf{x} - \mathbf{u}_i)$ ;
11     $\mathbf{u}_i \leftarrow \mathbf{u}_i + \mathbf{z}_i - \mathbf{x}$ ;
12  end
13 until  $\mathbf{x}$  feasible;
14  $\rho = 1$ ;
15 repeat
16    $\mathbf{x} \leftarrow \frac{1}{m+\rho-1} \sum_{i=1}^m (\mathbf{z}_i + \mathbf{u}_i)$ ;
17   for  $i = 1, \dots, m$  do
18      $\mu_i \leftarrow \arg_{\mu} \phi_i(\mu) = 0$  using Alg. 1 or 2;
19      $\mu_i \leftarrow \max\{0, \mu_i\}$ ; // for  $\leq$  constraints
20      $\mathbf{z}_i \leftarrow \mathbf{Q}_i (\mathbf{I} + \mu_i \mathbf{\Lambda}_i)^{-1} \mathbf{Q}_i^H (\mathbf{x} - \mathbf{u}_i)$ ;
21      $\mathbf{u}_i \leftarrow \mathbf{u}_i + \mathbf{z}_i - \mathbf{x}$ ;
22  end
23 until The successive difference of  $\mathbf{x}$  is smaller than  $\varepsilon$ ;
```

TABLE I
AVERAGED PERFORMANCE, OVER 100 MONTE-CARLO TRIALS, IN FEASIBLE POINT PURSUIT

	feasible point		avg. loss (dB)		avg. time (sec.)	
$m = 32$	100%	100%	0.376	0.375	4.5	31.9
$m = 40$	100%	100%	0.503	0.526	5.4	37.7
$m = 48$	100%	100%	0.600	0.597	7.9	44.5

In each column, consensus-ADMM is on the left, and FPP-SCA is on the right. The average loss is defined as $10 \log_{10} (\|\mathbf{x}\|^2 / \text{Trace}\{\mathbf{X}\})$, where \mathbf{X} is the solution of the SDR.

from $\mathcal{N}(0, 1)$. The constructed constraint set is therefore guaranteed to be non-empty, because we know \mathbf{x}_{feas} is feasible, similar to the problem setting considered in [11]. For $n = 20$ and $m \in \{32, 40, 48\}$, the averaged results over 100 Monte-Carlo trials are presented in Table I, and in each trial both ADMM and FPP-SCA are initialized with the same point, which is randomly generated from $\mathcal{CN}(0, \mathbf{I})$. As we can see, our proposed algorithm is able to produce similar performance with a much shorter execution time. It is possible to develop specialized solvers for FPP-SCA to accelerate it, but it is a non-trivial task which may require a lot of thinking, whereas our algorithm, readily available in Algorithm 3, only requires elementary operations and simple iterations, thus it is also easy to code in a lower-level language.

To illustrate the scalability of our algorithm, we tested it on a larger problem with $n = 100$ and $m = 200$, and our algorithm took about 8 minutes to find a feasible point with smallest norm, which took about 10^4 iterations. As shown in Fig. 1, the final result is not very far away from the generally unattainable lower bound provided by the SDR, with loss only about 0.45 dB. If all we need is a feasible point, then it only requires about 200 iterations, showing great efficiency in finishing the most

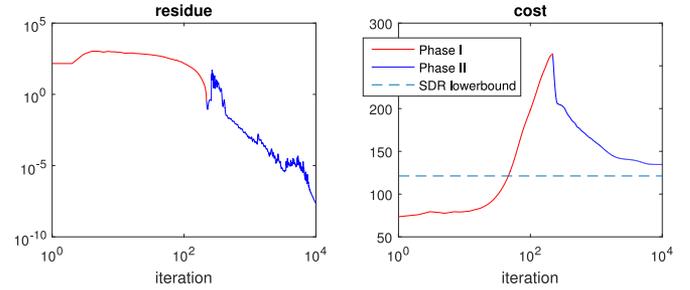


Fig. 1. Residual of the equality constraints $\sum_{i=1}^m \|\mathbf{z}_i - \mathbf{x}\|^2$ (left) and cost function $\|\mathbf{x}\|^2$ (right) vs. iteration number of Algorithm 3 for one random problem instance.

important task. In comparison, FPP-SCA requires more than 25 minutes to achieve a similar result.

B. Multicast Beamforming

Transmit beamforming is a wireless communication technique for transmitting signals to one or more users in a spatially selective way. A transmit beamforming system comprises a base station equipped with n antennas, transmitting signals to a set of m users within a certain service area, each having a single antenna. Assuming the transmitter is able to acquire channel state information (CSI) for all the users, multicast beamforming corresponds to the case when the base station *selectively broadcasts* a common information-bearing signal to many users, ensuring a minimum received signal-to-noise ratio (SNR) at each user terminal, with the goal that the total transmission power is much smaller than the traditional method of radiating power isotropically around its service area. Clearly, this also controls interference to other nearby systems, e.g., in neighboring cells.

There are various formulations of multicast beamforming, ranging from single-group to multiple-groups, perfect CSI at the base station to channel second order statistics only, to name just a few; cf. [20] and the references therein. Almost all formulations are within the range of non-convex QCQP (and NP-hard [21]), therefore it makes sense to test the performance of our proposed algorithm in this application. For brevity, we only consider the case when perfect CSI is available at the base station transmitter, corresponding to a fixed wireless scenario.

1) *Single-Group Multicast Beamforming*: The most basic multicast beamforming formulation takes the following form [21]:

$$\begin{aligned}
 & \underset{\mathbf{w} \in \mathcal{C}^n}{\text{minimize}} && \|\mathbf{w}\|^2, \\
 & \text{subject to} && |\mathbf{h}_i^H \mathbf{w}|^2 \geq 1, \forall i = 1, \dots, m,
 \end{aligned} \tag{11}$$

where each \mathbf{h}_i corresponds to the channel coefficients scaled according to the additive noise power and receive SNR requirement. Given $\{\mathbf{h}_i\}$, we wish to guarantee a certain SNR to all the receivers, while minimizing the transmit power $\|\mathbf{w}\|^2$.

Problem (11) is exactly in the form of (7), except that the constraints are with inequalities, thus we can directly use the memory-efficient updates (9). For initialization, we found em-

Algorithm 4: Consensus-ADMM for (11).

- 1 initialize $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I})$;
 - 2 $\mathbf{w} \leftarrow \mathbf{w} / \min(\|\mathbf{H}_s^H \mathbf{w}\|)$;
 - 3 $\mathbf{z}_s \leftarrow m\mathbf{w}$;
 - 4 $\mathbf{u}_s \leftarrow 0$;
 - 5 $\rho = 2\sqrt{m}$;
 - 6 **repeat**
 - 7 $\mathbf{w} \leftarrow \frac{1}{m+\rho^{-1}}(\mathbf{z}_s + \mathbf{u}_s)$;
 - 8 $\boldsymbol{\xi} \leftarrow \mathbf{H}_s^H \mathbf{w}$;
 - 9 $\nu_i \leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} \frac{[1 - |\xi_i - \alpha_i|]_+}{\|\mathbf{h}_i\|^2}$;
 - 10 $\mathbf{z}_s \leftarrow m\mathbf{w} - \mathbf{u}_s + \mathbf{H}_s \boldsymbol{\nu}$;
 - 11 $\mathbf{u}_s \leftarrow \mathbf{u}_s + \mathbf{z}_s - m\mathbf{w}$;
 - 12 $\alpha_i \leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} [1 - |\xi_i - \alpha_i|]_+$;
 - 13 **until** The successive difference of \mathbf{w} is smaller than ε ;
-

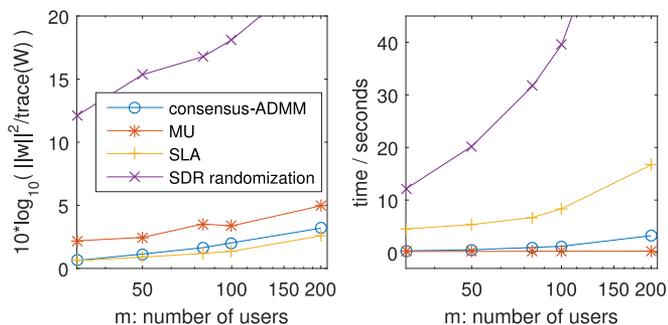


Fig. 2. Averaged performance of various methods for (11) over 100 Monte-Carlo trials, with $n = 100$ and m from 30 to 200. Performance gap relative to the generally unattainable SDR lower bound on the left, and computation time on the right.

pirically that it is better to initialize with a feasible point, so that we can afford to use a smaller ρ to provide faster convergence, reducing the risk of having unstable sequences. Fortunately it is easy to find a feasible point for (11) – for any random initialization point, one only needs to scale it up until all the constraints are satisfied. The complete algorithm for (11), including the initialization strategy and our choice of ρ is given in Algorithm 4.

We test the numerical performance of Algorithm 4 on problem (11), and we compare it with the standard SDR followed by Gaussian randomization [21], successive linear approximation (SLA) [22], and the recently proposed multiplicative update (MU) [23], which uses an approximate formulation and leads to highly efficient approximate solutions to (11). CVX [24] is used to solve the SDR and SLA in our experiment. We fix $n = 100$, and vary $m \in \{30, 50, 80, 100, 200\}$, with each channel vector \mathbf{h}_i randomly generated from $\mathcal{CN}(0, \mathbf{I})$. The averaged gap to the SDR lower bound, and the averaged computation time for all algorithms, with averages taken over 100 Monte-Carlo trials, are shown in Fig. 2, with each algorithm initialized at the same randomly generated points. As we can see, ADMM is able to give slightly worse performance than SLA in terms of transmission power, while requiring far smaller execution time.

TABLE II
AVERAGED PERFORMANCE OF VARIOUS METHODS FOR (11) OVER 100 MONTE-CARLO TRIALS, WITH $n = 500$ AND $m = 100$

	$\ \mathbf{w}\ ^2$	computation time
consensus-ADMM	0.1131	0.5235 sec.
1 step of SLA	0.1213	0.6204 sec.
≤ 10 steps of SLA	0.1125	6.2115 sec.

To test the scalability of our algorithm, we also applied it to a massive MIMO multicast scenario with $n = 500$ antennas and $m = 100$ users. In this case, CVX is not able to solve the SDR within reasonable amount of time, so we can only compare the transmission power $\|\mathbf{w}\|^2$ without knowing how far it is from the SDR lower bound. It is reported in [23] that MU followed by one step of SLA gives the best result in both minimizing the transmission power and keeping computation time low, so we compare our algorithm with this two-step strategy here, with one or up to ten SLA iterations (unless $\|\mathbf{w}\|^2$ converges, i.e., per-iteration improvement is less than 10^{-5}). For fair comparison, ADMM is initialized with the same point generated by MU. ADMM is able to compute a beamformer that is better than one step of SLA and do so in less time; in fact the transmission power obtained via ADMM is only slightly worse than ten steps of SLA. Notice that the update rule of ADMM is as simple as that of MU; both are simple enough to implement in real communication hardware, whereas SLA requires a full-blown convex optimization solver – which seems unrealistic for base station deployment as of this writing.

2) *Secondary User Multicast Beamforming:* We now consider adding primary user interference constraints to the basic multicast beamforming formulation in (11). This secondary multicast underlay scenario has been considered in [25], and the problem of interest can be formulated as

$$\begin{aligned}
 & \underset{\mathbf{w} \in \mathcal{C}^n}{\text{minimize}} && \|\mathbf{w}\|^2, \\
 & \text{subject to} && |\mathbf{h}_i^H \mathbf{w}|^2 \geq \tau, \forall i = 1, \dots, m, \\
 & && |\mathbf{g}_k^H \mathbf{w}|^2 \leq \eta, \forall k = 1, \dots, l,
 \end{aligned} \tag{12}$$

where in this case we have l additional primary users who should be protected for excess interference cause by the secondary multicast transmission, and \mathbf{g}_k denotes the channel vector from the multicast transmitter to the k -th (single-antenna) primary user receiver.

Again, the efficient updates in (8) for rank one quadratic constraints can be used. As for initialization, there is no obvious way to find a feasible point in this case, so the two-stage procedure we discussed before is used, which ignores the cost function first to find a feasible point, and then uses this feasible point to initialize the complete updates with a relatively small ρ to accelerate convergence. The complete algorithm is given in Algorithm 5.

Similar to the previous simulation settings, we fix $n = 100$ and $l = 10$, and vary $m \in \{30, 50, 80, 100\}$, with channel coefficients randomly generated from $\mathcal{CN}(0, \mathbf{I})$. For $\tau = 10$ and $\eta = 1$, the averaged performance over 100 Monte-Carlo runs comparing to the SDR lower bound and FPP-SCA as described

Algorithm 5: Consensus-ADMM for (12).

```

1 initialize  $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I})$ ;
2  $\mathbf{z}_s \leftarrow (m+l)\mathbf{w}$ ;
3  $\mathbf{u}_s \leftarrow 0$ ;
4 repeat
5    $\mathbf{w} \leftarrow \frac{1}{m+l}(\mathbf{z}_s + \mathbf{u}_s)$ ;
6    $\xi \leftarrow \mathbf{H}_s^H \mathbf{w}$ ;
    $\bar{\xi} \leftarrow \mathbf{G}_s^H \mathbf{w}$ ;
7    $\nu_i \leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} \frac{[\sqrt{\tau} - |\xi_i - \alpha_i|]_+}{\|\mathbf{h}_i\|^2}$ ;
    $\bar{\nu}_i \leftarrow \frac{\bar{\xi}_i - \bar{\alpha}_i}{|\bar{\xi}_i - \bar{\alpha}_i|} \frac{[\sqrt{\eta} - |\bar{\xi}_i - \bar{\alpha}_i|]_-}{\|\mathbf{g}_i\|^2}$ ;
8    $\mathbf{z}_s \leftarrow (m+l)\mathbf{w} - \mathbf{u}_s + \mathbf{H}_s \boldsymbol{\nu} + \mathbf{G}_s \bar{\boldsymbol{\nu}}$ ;
9    $\mathbf{u}_s \leftarrow \mathbf{u}_s + \mathbf{z}_s - (m+l)\mathbf{w}$ ;
    $\alpha_i \leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} [\sqrt{\tau} - |\xi_i - \alpha_i|]_+$ ;
10   $\bar{\alpha}_i \leftarrow \frac{\bar{\xi}_i - \bar{\alpha}_i}{|\bar{\xi}_i - \bar{\alpha}_i|} [\sqrt{\eta} - |\bar{\xi}_i - \bar{\alpha}_i|]_-$ ;
11 until  $\mathbf{w}$  feasible;
12  $\rho = 2\sqrt{m+l}$ ;
13 repeat
14    $\mathbf{w} \leftarrow \frac{1}{m+l+\rho-1}(\mathbf{z}_s + \mathbf{u}_s)$ ;
15    $\xi \leftarrow \mathbf{H}_s^H \mathbf{w}$ ;
    $\bar{\xi} \leftarrow \mathbf{G}_s^H \mathbf{w}$ ;
16    $\nu_i \leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} \frac{[\sqrt{\tau} - |\xi_i - \alpha_i|]_+}{\|\mathbf{h}_i\|^2}$ ;
    $\bar{\nu}_i \leftarrow \frac{\bar{\xi}_i - \bar{\alpha}_i}{|\bar{\xi}_i - \bar{\alpha}_i|} \frac{[\sqrt{\eta} - |\bar{\xi}_i - \bar{\alpha}_i|]_-}{\|\mathbf{g}_i\|^2}$ ;
17    $\mathbf{z}_s \leftarrow (m+l)\mathbf{w} - \mathbf{u}_s + \mathbf{H}_s \boldsymbol{\nu} + \mathbf{G}_s \bar{\boldsymbol{\nu}}$ ;
18    $\mathbf{u}_s \leftarrow \mathbf{u}_s + \mathbf{z}_s - (m+l)\mathbf{w}$ ;
    $\alpha_i \leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} [\sqrt{\tau} - |\xi_i - \alpha_i|]_+$ ;
19    $\bar{\alpha}_i \leftarrow \frac{\bar{\xi}_i - \bar{\alpha}_i}{|\bar{\xi}_i - \bar{\alpha}_i|} [\sqrt{\eta} - |\bar{\xi}_i - \bar{\alpha}_i|]_-$ ;
20 until The successive difference of  $\mathbf{w}$  is smaller than  $\varepsilon$ ;
```

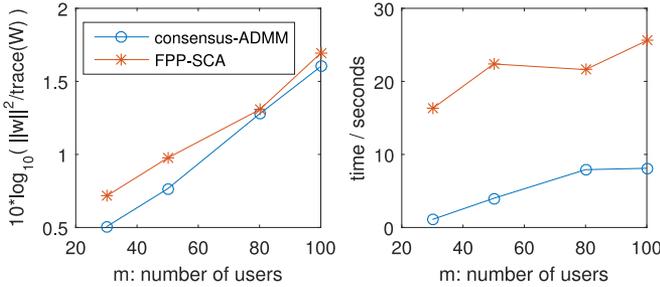


Fig. 3. Averaged performance of ADMM and FPP-SCA for (12) over 100 Monte-Carlo trials, with performance degradation relative to the generally unattainable SDR lower bound on the left, and computation time on the right.

in [11] is shown in Fig. 3. SDR randomization almost always fails to find a feasible point in this case, thus not considered in this experiment. Both methods are initialized with a random point from $\mathcal{CN}(0, \mathbf{I})$, and manage to obtain a good feasible point in all problem instances that we tried. We allow ADMM to take multiple initializations if the first stage of Algorithm 5 does not find a feasible point after 10^3 iterations, thus the computation time of ADMM is more likely to vary (as seen on the right panel of Fig. 3 for $m = 80$ and 100), although much smaller than that of FPP-SCA, which requires a general-purpose convex optimization solver, in our case CVX [24]. Note that ADMM also

yields better performance than FPP-SCA in terms of transmit power $\|\mathbf{w}\|^2$.

C. Phase Retrieval

Phase retrieval is the problem of estimating a signal from the magnitude of complex linear measurements, without access to the corresponding phases. This problem arises in various applications like crystallography, microscopy, and optical imaging [26]. Specifically, let \mathbf{s} be the desired signal, the measurements $\{y_i\}_{i=1}^m$ are collected via $y_i = |\mathbf{a}_i^H \mathbf{s}|^2$, possibly perturbed by noise. In the sequel we will see that for a number of noise models this problem can be written as non-convex QCQP, therefore we can test our algorithm together with other state-of-the-art phase retrieval methods. Notice that for some specific measurement systems the problem actually has hidden convexity, e.g., this is the case for 1-D over-sampled Fourier measurements [27], but here we focus on a general measurement setup.

1) *Noiseless Case:* Assuming all the measurements are exact, we can write the phase retrieval problem as the following feasibility problem

$$\begin{aligned} & \text{find} && \mathbf{x} \in \mathcal{C}^n, \\ & \text{such that} && |\mathbf{a}_i^H \mathbf{x}|^2 = y_i, \forall i = 1, \dots, m. \end{aligned} \quad (13)$$

This is exactly in the form of (7) except that there is no explicit cost function, so we can apply the memory efficient implementation of consensus-ADMM to obtain the following updates

$$\begin{aligned} \mathbf{x} & \leftarrow \frac{1}{m}(\mathbf{z}_s + \mathbf{u}_s), \\ \boldsymbol{\xi} & \leftarrow \mathbf{A}_s^H \mathbf{x}, \\ \nu_i & \leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} \frac{[\sqrt{y_i} - |\xi_i - \alpha_i|]}{\|\mathbf{a}_i\|^2}, \\ \mathbf{z}_s & \leftarrow m\mathbf{x} - \mathbf{u}_s + \mathbf{A}_s \boldsymbol{\nu}, \\ \mathbf{u}_s & \leftarrow \mathbf{u}_s + \mathbf{z}_s - m\mathbf{x}, \\ \alpha_i & \leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} (\sqrt{y_i} - |\xi_i - \alpha_i|), \end{aligned} \quad (14)$$

where $\mathbf{A}_s = [\mathbf{a}_1 \mathbf{a}_2 \dots \mathbf{a}_m]$ is obtained by stacking all the \mathbf{a}_i vectors as its columns. Notice that since we do not have an explicit cost function, it does not matter what value of ρ we choose – they all work the same for this problem.

2) *Bounded Noise:* In practice the measurements are seldom perfect, so we need to incorporate uncertainties in the measurements. A simple assumption is that measurements are quantized at relatively high resolution, in which case we can model the measurements as being corrupted by noise that is uniformly distributed between $[-\epsilon, \epsilon]$. We can modify the noiseless formulation (13) as follows (similar to the B-FPP formulation proposed in [28])

$$\begin{aligned} & \text{find} && \mathbf{x} \in \mathcal{C}^n, \\ & \text{such that} && y_i - \epsilon \leq |\mathbf{a}_i^H \mathbf{x}|^2 \leq y_i + \epsilon, \forall i = 1, \dots, m. \end{aligned} \quad (15)$$

Again we can apply the memory efficient implementation with the following updates

$$\begin{aligned}
 \mathbf{x} &\leftarrow \frac{1}{m}(\mathbf{z}_s + \mathbf{u}_s), \\
 \boldsymbol{\xi} &\leftarrow \mathbf{A}_s^H \mathbf{x}, \\
 \tau_i &\leftarrow \begin{cases} \sqrt{y_i - \epsilon} - |\xi_i - \alpha_i|, & \text{if } |\xi_i - \alpha_i|^2 < y_i - \epsilon, \\ \sqrt{y_i + \epsilon} - |\xi_i - \alpha_i|, & \text{if } |\xi_i - \alpha_i|^2 > y_i + \epsilon, \\ 0, & \text{otherwise,} \end{cases} \\
 \nu_i &\leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} \frac{\tau_i}{\|\mathbf{a}_i\|^2}, \\
 \mathbf{z}_s &\leftarrow m\mathbf{x} - \mathbf{u}_s + \mathbf{A}_s \boldsymbol{\nu}, \\
 \mathbf{u}_s &\leftarrow \mathbf{u}_s + \mathbf{z}_s - m\mathbf{x}, \\
 \alpha_i &\leftarrow \frac{\xi_i - \alpha_i}{|\xi_i - \alpha_i|} \tau_i.
 \end{aligned} \tag{16}$$

3) *Gaussian Noise*: Another interesting scenario is that where measurements are corrupted by additive white Gaussian noise, in which case maximum likelihood estimation can be cast as the following non-convex QCQP (similar to the LS-FPP formulation proposed in [28])

$$\begin{aligned}
 &\underset{\mathbf{x} \in \mathcal{C}^n, \mathbf{w} \in \mathcal{R}^m}{\text{minimize}} && \frac{1}{2} \|\mathbf{w}\|^2 \\
 &\text{subject to} && |\mathbf{a}_i^H \mathbf{x}|^2 = y_i + w_i, \forall i = 1, \dots, m.
 \end{aligned} \tag{17}$$

This kind of constraint is not covered in our previous discussions, so we study this case in a bit more detail here. Let us first rewrite (17) into a consensus optimization form by introducing m auxiliary variables z_1, \dots, z_m replicating \mathbf{x} for each constraint

$$\begin{aligned}
 &\underset{\mathbf{x}, \{z_i\}, \mathbf{w}}{\text{minimize}} && \frac{1}{2} \|\mathbf{w}\|^2 \\
 &\text{subject to} && |\mathbf{a}_i^H z_i|^2 = y_i + w_i, \\
 &&& z_i = \mathbf{x}, \forall i = 1, \dots, m.
 \end{aligned}$$

The plain vanilla version of ADMM, treating \mathbf{x} as the first block and $\{z_i\}$ and \mathbf{w} as the second block, takes the following form:

$$\begin{aligned}
 \mathbf{x} &\leftarrow \frac{1}{m} \sum_{i=1}^m (z_i + \mathbf{u}_i), \\
 (z_i, w_i) &\leftarrow \arg \min_{|\mathbf{a}_i^H \mathbf{x}|^2 = y_i + w_i} \frac{1}{2} |w_i|^2 + \rho \|z_i - \mathbf{x} + \mathbf{u}_i\|^2, \\
 \mathbf{u}_i &\leftarrow \mathbf{u}_i + z_i - \mathbf{x}.
 \end{aligned}$$

The main difficulty boils down to an efficient method for the second update, which can be written explicitly as

$$\begin{aligned}
 &\underset{z_i, w_i}{\text{minimize}} && \frac{1}{2} |w_i|^2 + \rho \|z_i - \mathbf{x} + \mathbf{u}_i\|^2 \\
 &\text{subject to} && |\mathbf{a}_i^H z_i|^2 = y_i + w_i.
 \end{aligned} \tag{18}$$

The same idea of using the Lagrangian can be applied to solve (18). Constructing the Lagrangian with a single multiplier μ_i and setting its derivative with respect to w_i and z_i equal to

0, we have

$$w_i = \mu_i,$$

which, interestingly, means that the optimal multiplier is actually equal to the estimated noise term, and

$$\begin{aligned}
 z_i &= (\rho \mathbf{I} + \mu_i \mathbf{a}_i \mathbf{a}_i^H)^{-1} \rho (\mathbf{x} - \mathbf{u}_i) \\
 &= \mathbf{x} - \mathbf{u}_i - \frac{\mu_i \mathbf{a}_i^H (\mathbf{x} - \mathbf{u}_i)}{\rho + \mu_i \|\mathbf{a}_i\|^2} \mathbf{a}_i.
 \end{aligned} \tag{19}$$

Plugging them back into the equality constraint, we end up with an equation with respect to μ_i

$$\frac{\rho^2 |\mathbf{a}_i^H (\mathbf{x} - \mathbf{u}_i)|^2}{(\rho + \|\mathbf{a}_i\|^2 \mu_i)^2} = y_i + \mu_i. \tag{20}$$

Equation (20) is equivalent to a cubic equation, for which we know the formula for the three roots. Moreover, since we know the three roots of a real cubic equation are either all real or one real and two complex conjugates, and that the correct μ_i we are looking for is real, we can deliberately select the value of ρ so that the latter case happens, resolving the ambiguity in solving (20). Detailed derivation of the formula for solving (20) is given in Appendix B, where it is also shown that by setting $\rho = 1.1 \max_i y_i \|\mathbf{a}_i\|^2$, each equation (20) is guaranteed to have a unique real root, thus being the correct multiplier we are looking for.

Memory efficient implementation is again applicable here, which eventually leads to the following updates:

$$\begin{aligned}
 \mathbf{x} &\leftarrow \frac{1}{m}(\mathbf{z}_s + \mathbf{u}_s), \\
 \boldsymbol{\xi} &\leftarrow \mathbf{A}_s^H \mathbf{x}, \\
 \mu_i &\leftarrow \text{steps described in Appendix B,} \\
 \nu_i &\leftarrow -\frac{\mu_i}{\rho + \mu_i \|\mathbf{a}_i\|^2} (\xi_i - \alpha_i), \\
 \mathbf{z}_s &\leftarrow m\mathbf{x} - \mathbf{u}_s + \mathbf{A}_s \boldsymbol{\nu}, \\
 \mathbf{u}_s &\leftarrow \mathbf{u}_s + \mathbf{z}_s - m\mathbf{x}, \\
 \alpha_i &\leftarrow -\frac{\mu_i \|\mathbf{a}_i\|^2}{\rho + \mu_i \|\mathbf{a}_i\|^2} (\xi_i - \alpha_i).
 \end{aligned} \tag{21}$$

The only unclear part in (21) is the update of μ_i . However, since μ_i is a solution of (20), which only depends on $\mathbf{a}_i^H (\mathbf{x} + \mathbf{u}_i) = \xi_i + \alpha_i$, it is indeed possible to implement (21) without explicitly calculating the individual z_i 's and \mathbf{u}_i 's.

4) *Adding Priors*: In a lot of cases there is useful prior information available about the signal that can help enhance the estimation performance. For example, one may know a priori that the signal to be estimated is real, non-negative, and/or sparse. All of these type of prior information can easily be incorporated into the x -update, which usually boils down to very simple projections, like zeroing out the imaginary part and/or zeroing out the negative values.

For sparsity, a popular method is to add an ℓ_1 penalty to the cost, which in our case leads to a *soft-thresholding* to the update of \mathbf{x} . However, recall that we are dealing with NP-hard problems and there is no guarantee that

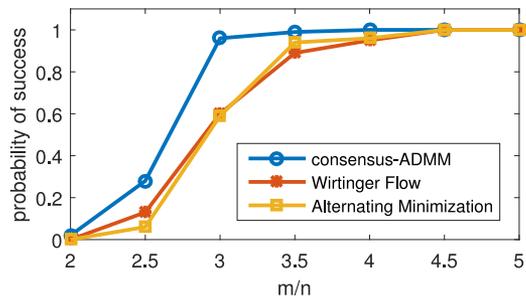


Fig. 4. Empirical probability of resolution based on 100 Monte-Carlo trials for various number of measurements.

our algorithm is always going to reach a global solution. Therefore, for practical purposes it is sometimes better to just use the straight-forward cardinality constraint, which is easy to tune for the desired sparsity level, and has an equally simple *hard-thresholding* update.

5) *Simulations*: Now we show some numerical results on the performance of the proposed algorithms. There exist many algorithms designed for phase retrieval under specific scenarios, for example, the classical Gerchberg-Saxton algorithm [29] and Fienup's algorithm [30] were designed for Fourier-based measurements, and have been successfully applied in the phase retrieval community for decades. More recently, inspired by the success of compressive sensing, random Gaussian measurements have gained more attention, and the state-of-the-art algorithms include alternating minimization [31] and Wirtinger flow [32], both of which contain a special initialization step to help convergence. SDR-based methods have also been developed, including PhaseLift [33] and PhaseCut [34], however, they again suffer from effectively squaring the problem dimension, and how to recover a good approximate solution to the original problem when they return a higher rank matrix is an open question.

We consider the columns of the sensing matrix $\{\mathbf{a}_i\}$ to be generated from an i.i.d. complex Gaussian distribution $\mathcal{CN}(0, \mathbf{I})$. In a noiseless scenario, we apply the updates (14) to the problem of phase retrieval with random Gaussian measurements, and compare it with two state-of-the-art algorithms designed for this setting: Wirtinger flow [32] and alternating minimization [31]. Notice that algorithm (14) only requires two matrix-vector multiplications, thus the per-iteration complexity is the same as that of Wirtinger flow and alternating minimization.

We randomly generate a desired signal $\mathbf{s} \in \mathcal{C}^n$ where $n = 128$, then take m phase-less measurements $|\mathbf{a}_i^H \mathbf{s}|^2$ with m ranging from $2n$ to $5n$. Using the same initialization proposed in [31] and [32], we let consensus-ADMM (Algorithm (14)), Wirtinger flow, and alternating minimization run for at most 10^3 iterations, and the probability of resolution over 100 Monte-Carlo trials is given in Fig. 4, where we declare that the signal has been resolved if (after adjusting for the global phase ambiguity of the estimate \mathbf{x}),

$$\min_{\theta} \|e^{j\theta} \mathbf{x} - \mathbf{s}\|^2 < 10^{-5}.$$

TABLE III
PERFORMANCE OF QUANTIZED PHASE RETRIEVAL

	# of violations	MSE
consensus-ADMM	0	-34 dB
Wirtinger flow	13.3	-37 dB
Alternating Min.	143.4	-31 dB

TABLE IV
PERFORMANCE OF GAUSSIAN NOISE PHASE RETRIEVAL

	$\ \mathbf{y} - \mathbf{A}_s \mathbf{x} ^2\ ^2$	MSE
consensus-ADMM	6.9e3	-2.4 dB
Wirtinger flow	6.9e3	-2.4 dB
Alternating Min.	1.1e4	0.4 dB

It is very satisfying to see that consensus-ADMM has higher empirical probability of resolution than both Wirtinger flow and alternating minimization. Due to the similar per-iteration complexities of the three algorithms, running 10^5 iterations take approximately the same time for all of them.

Finally, we briefly show the performance of Algorithm (16) and (21), under their corresponding noise models, and we only show the results for $n = 128$ and $m = 5n$, with \mathbf{s} and \mathbf{A}_s generated as before. Consider the quantized measurements $\mathbf{y} = \lfloor |\mathbf{A}_s^H \mathbf{s}|^2 \rfloor$, where $\lfloor \cdot \rfloor$ rounds the argument to the nearest integer, we can use formulation (15) with $\epsilon = 0.5$, and apply Algorithm (16). The number of constraint violations and the mean squared error (MSE) are given in Table III, each averaged over 100 Monte Carlo trials, where MSE is defined as

$$\text{MSE} = 10 \log_{10} (\min_{\theta} \|e^{j\theta} \mathbf{x} - \mathbf{s}\|^2).$$

As we can see, Algorithm (16) is able to give a solution that is consistent with all the measurements in all cases, whereas the other two algorithms cannot, even though their MSE performance is still pretty good. For additive white Gaussian noise $\mathbf{y} = |\mathbf{A}_s^H \mathbf{s}|^2 + \mathbf{w}$, we fix the SNR to be 20 dB, and the averaged performance over 100 Monte Carlo trials is shown in Table IV. Algorithm (21) performs almost as well as Wirtinger flow, and both perform better than alternating minimization in this case. This is as expected, since alternating minimization aims to solve a different formulation, which is not the maximum likelihood one for this model. To sum up, consensus-ADMM is able to achieve similar (if not better) performance to the state-of-the-art methods for phase retrieval with random Gaussian measurements.

VI. CONCLUSION

In this paper, we have proposed a new algorithm for general non-convex QCQPs, which is very different from any existing methods, general or specialized, for such problems. The main ideas behind this proposed algorithm are:

- Any QCQP-1 can be optimally solved, irrespective of (non-)convexity;
- Consensus ADMM can be used to solve general QCQPs, in such a way that each update requires to solve a number of QCQP-1's.

For the type of QCQP-1's encountered here, we showed that they can be solved very efficiently, and made important steps towards enhancing scalability of the overall algorithm, including

- Memory-efficient implementations for certain types of “simple” constraints, e.g., rank one quadratics;
- Parallel/distributed implementations with small communication overhead.

The proposed algorithm and various custom implementations were fully fleshed out and applied to various important non-convex QCQP problems, from pure feasibility pursuit to two real-world engineering tasks: multicast beamforming and phase retrieval. The proposed algorithms consistently exhibited favorable performance compared to the prior state of art, including classical and more modern methods for phase retrieval, which has drawn renewed interest in recent years. We believe that the general applicability of the proposed algorithm has the potential to bring significant performance improvements to many other applications of non-convex QCQP as well.

APPENDIX A PROOF OF THEOREM 1

A KKT point \mathbf{x}_* of (1), together with the corresponding dual variable $\boldsymbol{\mu}_*$, satisfies that

$$\mathbf{A}_0 \mathbf{x}_* - \mathbf{b}_0 + \sum_{i=1}^m \mu_{i*} (\mathbf{A}_i \mathbf{x}_* - \mathbf{b}_i) = 0, \quad (22a)$$

$$\boldsymbol{\mu}_* \geq 0, \quad (22b)$$

$$\mathbf{x}_*^H \mathbf{A}_i \mathbf{x}_* - 2\Re\{\mathbf{b}_i^H \mathbf{x}_*\} \leq c_i, \quad (22c)$$

$$\begin{aligned} \mu_{i*} (\mathbf{x}_*^H \mathbf{A}_i \mathbf{x}_* - 2\Re\{\mathbf{b}_i^H \mathbf{x}_*\}) &= 0, \\ \forall i &= 1, \dots, m. \end{aligned} \quad (22d)$$

Let us use a superscript t to denote the point obtained at iteration t by the update rule (4), then at iteration $t+1$, since we assume that each $\{\mathbf{z}_i^t\}$ is well defined, we have that

$$\mathbf{A}_0 \mathbf{x}^{t+1} + m\rho \mathbf{x}^{t+1} = \mathbf{b}_0 + \rho \sum_{i=1}^m (\mathbf{z}_i^t + \mathbf{u}_i^t), \quad (23a)$$

$$\begin{aligned} \exists \mu_i^{t+1} \geq 0 \text{ s.t. } \mathbf{u}_i^{t+1} + \mu_i^{t+1} (\mathbf{A}_i \mathbf{z}_i^{t+1} - \mathbf{b}_i) &= 0, \\ \forall i &= 1, \dots, m \end{aligned}$$

where in the second equation we used the fact that

$$\mathbf{u}^{t+1} = \mathbf{u}^t + \mathbf{z}_i^{t+1} - \mathbf{x}^{t+1}.$$

Now by assuming $\mathbf{z}_i^{t+1} - \mathbf{x}^{t+1} \rightarrow 0$, we trivially have that

$$\mathbf{u}_i^{t+1} - \mathbf{u}_i^t \rightarrow 0.$$

Therefore, for t sufficiently large, we have

$$\mathbf{A}_0 \mathbf{x}^{t+1} + m\rho \mathbf{x}^{t+1} = \mathbf{b}_0 + \rho \sum_{i=1}^m (\mathbf{x}^t - \mu_i^{t+1} (\mathbf{A}_i \mathbf{x}^{t+1} - \mathbf{b}_i)).$$

By further assuming

$$\mathbf{x}^{t+1} - \mathbf{x}^t \rightarrow 0,$$

this becomes exactly (22a) by setting $\boldsymbol{\mu}_* = \rho \boldsymbol{\mu}^{t+1}$. The rest of the KKT conditions are guaranteed by the feasibility of \mathbf{z}_i^{t+1} for the i -th constraint, and our assumption that $\mathbf{z}^{t+1} - \mathbf{x}^{t+1} \rightarrow 0$.

APPENDIX B SOLVING (20)

We derive the solution for equation (20), which can be equivalently written as the following cubic equation after we drop the subscripts,

$$\begin{aligned} \|\mathbf{a}\|^4 \mu^3 + (2\rho \|\mathbf{a}\|^2 + y \|\mathbf{a}\|^4) \mu^2 + (2y\rho \|\mathbf{a}\|^2 + \rho^2) \mu \\ + y\rho^2 - \rho^2 \|\mathbf{a}^H(\mathbf{x} - \mathbf{u})\|^2 = 0. \end{aligned} \quad (24)$$

For a general cubic equation

$$\gamma_3 \mu^3 + \gamma_2 \mu^2 + \gamma_1 \mu + \gamma_0 = 0,$$

the three roots can be found with the following formulas (assuming all of these quantities are non-zero, which can be ensured by our specific choice of ρ presented in the sequel)

$$\Delta_0 = \gamma_2^2 - 3\gamma_3\gamma_1,$$

$$\Delta_1 = 2\gamma_2^3 - 9\gamma_3\gamma_2\gamma_1 + 27\gamma_3^2\gamma_0,$$

$$C = \sqrt[3]{\frac{\Delta_1 + \sqrt{\Delta_1^2 - 4\Delta_0^3}}{2}},$$

$$\hat{\mu}_k = -\frac{1}{3\gamma_3} \left(\gamma_2 + \iota_k C + \frac{\Delta_0}{\iota_k C} \right),$$

where $\iota_1 = 1, \iota_2 = \frac{-1+j\sqrt{3}}{2}, \iota_3 = \frac{-1-j\sqrt{3}}{2}$ are the three cubic roots of 1. Furthermore, if all the coefficients are real, then there is at least one real root, and the other two are either complex conjugates or both real, depending on whether $\Delta_1^2 - 4\Delta_0^3$ is negative or positive.

Now, let us plug in the coefficients of (24) into the formula, to get

$$\Delta_0 = (\rho \|\mathbf{a}\|^2 - y \|\mathbf{a}\|^4)^2,$$

$$\Delta_1 = -2(\rho \|\mathbf{a}\|^2 - y \|\mathbf{a}\|^4)^3 - 27\rho^2 \|\mathbf{a}\|^8 |\mathbf{a}^H(\mathbf{x} + \mathbf{u})|^2,$$

$$\begin{aligned} \Delta_1^2 - 4\Delta_0^3 &= \left(27\rho^2 \|\mathbf{a}\|^8 |\mathbf{a}^H(\mathbf{x} + \mathbf{u})|^2 \right)^2 \\ &+ \left(27\rho^2 \|\mathbf{a}\|^8 |\mathbf{a}^H(\mathbf{x} + \mathbf{u})|^2 \right) \\ &\times \left(4(\rho \|\mathbf{a}\|^2 - y \|\mathbf{a}\|^4)^3 \right). \end{aligned}$$

Before we proceed, recall that the coefficients of (20) are all real, and the root that we are looking for, which is the optimal Lagrange multiplier, is also real. Therefore, to make our life easier, we can choose the value of ρ to ensure that $\Delta_1^2 - 4\Delta_0^3 > 0$, for example by setting $\rho > y \|\mathbf{a}\|^2$. Then we can proceed to the rest of the steps and choose the solution to be $\hat{\mu}_1$, the only real root of (20). Putting the subscripts back and considering there are m of them, this means we should set $\rho > \max_i y_i \|\mathbf{a}_i\|^2$. In practice smaller ρ usually leads to faster convergence, we therefore recommend setting

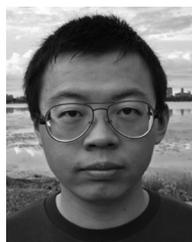
$$\rho = 1.1 \max_i y_i \|\mathbf{a}_i\|^2.$$

ACKNOWLEDGMENT

The authors would like to thank Prof. V. Elser from Cornell University for correcting a mistake made in an earlier version of this paper.

REFERENCES

- [1] Z.-Q. Luo, W.-K. Ma, M.-C. So, Y. Ye, and S. Zhang, "Semidefinite relaxation of quadratic optimization problems," *IEEE Signal Process. Mag.*, vol. 27, no. 3, pp. 20–34, 2010.
- [2] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [3] Y. Ye and S. Zhang, "New results on quadratic minimization," *SIAM J. Optim.*, vol. 14, no. 1, pp. 245–267, 2003.
- [4] A. Beck and Y. C. Eldar, "Strong duality in nonconvex quadratic optimization with two quadratic constraints," *SIAM J. Optim.*, vol. 17, no. 3, pp. 844–860, 2006.
- [5] Y. Huang and D. P. Palomar, "Rank-constrained separable semidefinite programming with applications to optimal beamforming," *IEEE Trans. Signal Process.*, vol. 58, no. 2, pp. 664–678, Feb. 2010.
- [6] W. Ai, Y. Huang, and S. Zhang, "New results on Hermitian matrix rank-one decomposition," *Math. Program.*, vol. 128, nos. 1/2, pp. 253–283, 2011.
- [7] A. Konar and N. D. Sidiropoulos, "Hidden convexity in QCQP with Toeplitz-Hermitian quadratics," *IEEE Signal Process. Lett.*, vol. 22, no. 10, pp. 1623–1627, Oct. 2015.
- [8] M. X. Goemans and D. P. Williamson, "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *J. ACM*, vol. 42, no. 6, pp. 1115–1145, 1995.
- [9] Z.-Q. Luo, N. D. Sidiropoulos, P. Tseng, and S. Zhang, "Approximation bounds for quadratic optimization with homogeneous quadratic constraints," *SIAM J. Optim.*, vol. 18, no. 1, pp. 1–28, 2007.
- [10] A. L. Yuille and A. Rangarajan, "The concave-convex procedure," *Neural Comput.*, vol. 15, no. 4, pp. 915–936, 2003.
- [11] O. Mehanna, K. Huang, B. Gopalakrishnan, A. Konar, and N. D. Sidiropoulos, "Feasible point pursuit and successive approximation of non-convex QCQPs," *IEEE Signal Process. Lett.*, vol. 22, no. 7, pp. 804–808, Jul. 2015.
- [12] G. H. Golub and C. F. Van Loan, *Matrix Computations*, vol. 3. Baltimore, MD, USA: The Johns Hopkins Univ. Press, 1996.
- [13] S. P. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*, vol. 15. Philadelphia, PA, USA: SIAM, 1994.
- [14] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2011.
- [15] J. Eckstein and D. P. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Math. Program.*, vol. 55, nos. 1–3, pp. 293–318, 1992.
- [16] Y. Xu, W. Yin, Z. Wen, and Y. Zhang, "An alternating direction algorithm for matrix completion with nonnegative factors," *Frontiers Math. China*, vol. 7, no. 2, pp. 365–384, 2012.
- [17] B. Jiang, S. Ma, and S. Zhang, "Alternating direction method of multipliers for real and complex polynomial optimization models," *Optimization*, vol. 63, no. 6, pp. 883–898, 2014.
- [18] A. P. Liavas and N. D. Sidiropoulos, "Parallel algorithms for constrained tensor factorization via the alternating direction method of multipliers," *IEEE Trans. Signal Process.*, vol. 63, no. 20, pp. 5450–5463, Oct. 2015.
- [19] M. Hong, Z.-Q. Luo, and M. Razaviyayn, "Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems," *SIAM J. Optim.*, vol. 26, no. 1, pp. 337–364, 2016.
- [20] A. B. Gershman, N. D. Sidiropoulos, S. Shahbazpanahi, M. Bengtsson, and B. Ottersten, "Convex optimization-based beamforming: From receive to transmit and network designs," *IEEE Signal Process. Mag.*, vol. 27, no. 3, pp. 62–75, May 2010.
- [21] N. D. Sidiropoulos, T. N. Davidson, and Z.-Q. Luo, "Transmit beamforming for physical-layer multicasting," *IEEE Trans. Signal Process.*, vol. 54, no. 6, pp. 2239–2251, Jun. 2006.
- [22] L.-N. Tran, M. F. Hanif, and M. Juntti, "A conic quadratic programming approach to physical layer multicasting for large-scale antenna arrays," *IEEE Signal Process. Lett.*, vol. 21, no. 1, pp. 114–117, Jan. 2014.
- [23] B. Gopalakrishnan and N. D. Sidiropoulos, "High performance adaptive algorithms for single-group multicast beamforming," *IEEE Trans. Signal Process.*, vol. 63, no. 16, pp. 4373–4384, Aug. 2015.
- [24] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," Mar. 2014. [Online]. Available: <http://cvxr.com/cvx>
- [25] K. T. Phan, S. Vorobyov, N. D. Sidiropoulos, and C. Tellambura, "Spectrum sharing in wireless networks via QoS-aware secondary multicast beamforming," *IEEE Trans. Signal Process.*, vol. 57, no. 6, pp. 2323–2335, Jun. 2009.
- [26] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, "Phase retrieval with application to optical imaging: A contemporary overview," *IEEE Signal Process. Mag.*, vol. 32, no. 3, pp. 87–109, May 2015.
- [27] K. Huang, Y. C. Eldar, and N. D. Sidiropoulos, "On convexity and identifiability in 1-D Fourier phase retrieval," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2016, pp. 3941–3945.
- [28] C. Qian, N. D. Sidiropoulos, K. Huang, L. Huang, and H.-C. So, "Phase retrieval using feasible point pursuit: Algorithms and Cramér-Rao bound," *IEEE Trans. Signal Process.*, vol. 64, no. 20, pp. 5282–5296, 2016.
- [29] R. Gerchberg and W. Saxton, "A practical algorithm for the determination of phase from image and diffraction plane pictures," *Optik*, vol. 35, p. 237, 1972.
- [30] J. R. Fienup, "Reconstruction of an object from the modulus of its Fourier transform," *Opt. Lett.*, vol. 3, no. 1, pp. 27–29, 1978.
- [31] P. Netrapalli, P. Jain, and S. Sanghavi, "Phase retrieval using alternating minimization," *IEEE Trans. Signal Process.*, vol. 63, no. 18, pp. 4814–4826, Sep. 2015.
- [32] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval via Wirtinger flow: Theory and algorithms," *IEEE Trans. Inf. Theory*, vol. 61, no. 4, pp. 1985–2007, 2015.
- [33] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM J. Imag. Sci.*, vol. 6, no. 1, pp. 199–225, 2013.
- [34] I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, MaxCut and complex semidefinite programming," *Math. Program.*, vol. 149, nos. 1/2, pp. 47–81, 2015.



Kejun Huang (S'13) received the Ph.D. degree in electrical engineering from the University of Minnesota, Minneapolis, MN, USA, in 2016, the Bachelor's degree from Nanjing University of Information Science and Technology, Nanjing, China, in 2010. His research interests include signal processing, machine learning, and optimization.



Nicholas D. Sidiropoulos (F'09) received the Diploma degree in electrical engineering from Aristotelian University of Thessaloniki, Thessaloniki, Greece, and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland-College Park, College Park, MD, USA, in 1988, 1990, and 1992, respectively. He served as an Assistant Professor at the University of Virginia, an Associate Professor at the University of Minnesota, and a Professor at Technical University of Crete, Chania, Greece. Since 2011, he has been at the University of Minnesota,

where he currently holds an ADC Chair in digital technology. His research interests include signal processing theory and algorithms, optimization, communications, and factor analysis—with a long-term interest in tensor decomposition and its applications. His current focus is primarily on signal and tensor analytics for learning from big data. He received the NSF/CAREER award in 1998, and the IEEE Signal Processing (SP) Society Best Paper Award in 2001, 2007, and 2011. He served as an IEEE SP Society Distinguished Lecturer (2008–2009), and as the Chair of the IEEE Signal Processing for Communications and Networking Technical Committee (2007–2008). He received the 2010 IEEE SP Society Meritorious Service Award, and the 2013 Distinguished Alumni Award from the Dept. of ECE, University of Maryland. He is a Fellow of EURASIP (2014).