Phase Retrieval from 1D Fourier Measurements: Convexity, Uniqueness, and Algorithms*

Kejun Huang, Yonina C. Eldar, and Nicholas D. Sidiropoulos

July 8, 2016

Abstract

This paper considers phase retrieval from the magnitude of 1D over-sampled Fourier measurements, a classical problem that has challenged researchers in various fields of science and engineering. We show that an optimal vector in a least-squares sense can be found by solving a convex problem, thus establishing a hidden convexity in Fourier phase retrieval. We then show that the standard semidefinite relaxation approach yields the optimal cost function value (albeit not necessarily an optimal solution). A method is then derived to retrieve an optimal minimum phase solution in polynomial time. Using these results, a new measuring technique is proposed which guarantees uniqueness of the solution, along with an efficient algorithm that can solve large-scale Fourier phase retrieval problems with uniqueness and optimality guarantees.

Keywords: phase retrieval, over-sampled Fourier measurements, minimum phase, auto-correlation retrieval, semi-definite programming, alternating direction method of multipliers, holography.

1 Introduction

Phase retrieval seeks to recover a signal from the magnitudes of linear measurements [2]. This problem arises in various fields, including crystallography [3], microscopy, and optical imaging [4], due to the limitations of the detectors used in these applications. Different types of measurement systems have been proposed, e.g., over-sampled Fourier measurements, short-time Fourier measurements, and random Gaussian, to name just a few (see [5,6] for contemporary reviews), but over-sampled Fourier measurements are most common in practice. Two fundamental questions in phase retrieval are: i) is the signal uniquely determined by the noiseless magnitude measurements (up to inherent ambiguities such as global phase); and ii) is there an efficient algorithm that can provably compute an optimal estimate of the signal according to a suitable criterion?

This paper considers the 1D phase retrieval problem from over-sampled Fourier measurements. It is well known that there is no uniqueness in 1D Fourier phase retrieval, i.e. there are multiple 1D signals with the same Fourier magnitude. This is true even when we ignore trivial ambiguities which include a global phase shift, conjugate inversion and spatial shift. Nonuniqueness also holds when the support of the signal is bounded within a known range [7]. In addition, since the phase retrieval problem is nonconvex, there are no known algorithms that provably minimize the least-squares error in recovering the underlying signal [8].

---

arXiv:1603.05215v2 [math.OC] 7 Jul 2016

Existing methods that attempt to resolve the identifiability issue include introducing sparsity assumptions on the input signal [9–12], taking multiple masked Fourier measurements [13], or using the short-time Fourier transform [14, 15]. Algorithmically, the most popular techniques for phase retrieval are based on alternating projections [8, 16, 17], pioneered by Gerchberg and Saxton [16] and extended by Fienup [17]. More recently, phase retrieval has been treated using semidefinite programming (SDP) and low-rank matrix recovery ideas [18, 19]. Several greedy approaches to phase retrieval have also been introduced such as GESPAR [20, 21]. Semidefinite relaxation for phase retrieval, referred to as PhaseLift [18], is known to recover the true underlying signal when the measurement vectors are Gaussian. However, in the case of Fourier measurements, as we are considering here, there are no known optimality results for this approach.

Despite the apparent difficulty in 1D Fourier phase retrieval, we establish in this paper that under certain conditions this problem can be solved optimally in polynomial time. In particular, we first show that the least-squares formulation of this problem can always be optimally solved using semidefinite relaxation. Namely, methods such as PhaseLift do in fact optimize the nonconvex cost in the case of Fourier measurements. By slightly modifying the basic PhaseLift program we propose a new semidefinite relaxation whose solution is proven to be rank-one and to minimize the error. However, due to the nonuniqueness in the problem, the solution is not in general equal to the true underlying vector. To resolve this ambiguity, we propose a semidefinite relaxation that will always return the minimum phase solution which optimizes the least-squares cost.

Based on this result, we suggest a new approach to measure 1D signals which transforms an arbitrary signal into a minimum phase counterpart by simply adding an impulse, so that identifiability is restored. We then recover the input using the proposed SDP. This measurement strategy resembles classic holography used commonly in optics [22]. Finally, we propose an efficient iterative algorithm to recover the underlying signal without the need to resort to lifting techniques as in the SDP approach. Our method, referred to as auto-correlation retrieval – Kolmogorov factorization (CoRK), is able to solve very large scale problems efficiently. Comparing to existing methods that provide identifiability, the proposed approach is easy to implement (both conceptually and practically), requires a minimal number of measurements, and can always be solved to global optimality using efficient computational methods.

The rest of the paper is organized as follows. We introduce the problem in more detail in Section 2. In Section 3 we reformulate phase retrieval by changing the variable from the signal itself to the correlation of the signal, and discuss two ways to characterize a correlation sequence, leading to convex optimization problems. Extracting signals from a given correlation sequence, known as *spectral factorization*, is briefly reviewed in Section 4, where we also review the notion of *minimum phase*. We then propose a new measurement system based on transforming an arbitrary signal into a minimum phase one, via a very simple impulse adding technique, in Section 5. The highly scalable CoRK algorithm is developed in Section 6 to recover the signal, providing the ability to optimally solve large-scale 1D Fourier phase retrieval problems very efficiently. We summarize the proposed method in Section 7, with some discussion on how to modify the approach if we know that the sought signal is real. Computer simulations are provided in Section 8 and show superiority of our proposed techniques both in terms of accuracy and efficiency. We finally conclude the paper in Section 9.

Throughout the paper, indices for vectors and matrices start at 0, so that the first entry of a vector $\boldsymbol{x}$ is x_0, and the upper-left-most entry of a matrix $\boldsymbol{X}$ is X_{00}. The superscript $(\cdot)^*$ denotes element-wise conjugate (without transpose), and $(\cdot)^H$ denotes Hermitian transpose of a vector or

a matrix.

2 The Phase Retrieval Problem

In the phase retrieval problem, we are interested in estimating a signal $\boldsymbol{x} \in \mathbb{C}^N$ from the squared magnitude of its Fourier transform. The *discrete-time Fourier transform* (DTFT) of a vector $\boldsymbol{x} \in \mathbb{C}^N$ is a trigonometric polynomial in ω defined as

$$X(e^{j\omega}) = \sum_{n=0}^{N-1} x_n e^{-j\omega n},$$

and is periodic with period 2π. In practice it is easier to obtain samples from the continuous function $X(e^{j\omega})$, which leads to the M-point *discrete Fourier transform* (DFT) of $\boldsymbol{x}$, where $M \geq N$, if we sample at points

$$\omega = 0, \frac{2\pi}{M}, ..., \frac{2\pi(M-1)}{M}.$$

This operation is equivalent to the matrix-vector multiplication $\boldsymbol{F}_M \boldsymbol{x}$, where $\boldsymbol{F}_M$ is the first N columns of the M-point DFT matrix, i.e.,

$$\boldsymbol{F}_M = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \phi & \phi^2 & \cdots & \phi^{N-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \phi^{M-1} & \phi^{2(M-1)} & \cdots & \phi^{(N-1)(M-1)} \end{bmatrix},$$

with $\phi = e^{-j2\pi/M}$. The matrix-vector multiplication $\boldsymbol{F}_M \boldsymbol{x}$ can be carried out efficiently via the fast Fourier transform (FFT) algorithm with complexity $O(M \log M)$, unlike the general case which takes $O(MN)$ flops to compute.

With this notation, our measurements are given by

$$\boldsymbol{b} = |\boldsymbol{F}_M \boldsymbol{x}|^2 + \boldsymbol{w}, \tag{1}$$

where $\boldsymbol{w}$ is a noise vector and $|\cdot|^2$ is taken element-wise. To recover $\boldsymbol{x}$ from $\boldsymbol{b}$ we consider a least-squares cost (which coincides with the maximum likelihood criterion assuming Gaussian noise):

$$\underset{\boldsymbol{x} \in \mathbb{C}^N}{\text{minimize}} \quad \left\| \boldsymbol{b} - |\boldsymbol{F}_M \boldsymbol{x}|^2 \right\|^2. \tag{2}$$

The most popular methods for solving (2) are the Gershburg-Saxton (GS) and Fienup's algorithms. Both techniques start with the noiseless scenario $\boldsymbol{b} = |\boldsymbol{F}_M \boldsymbol{x}|^2$, and reformulate it as the following feasibility problem by increasing the dimension of $\boldsymbol{x}$ to M and then imposing an additional *compact support* constraint:

$$\begin{aligned} \text{find} \quad & \boldsymbol{x} \in \mathbb{C}^M \\ \text{such that} \quad & \boldsymbol{b} = |\boldsymbol{F}_M \boldsymbol{x}|^2 \\ & x_n = 0, \quad n = N, N+1, ..., M-1. \end{aligned}$$

It is easy to derive projections onto the two individual sets of equality constraints, but not both. Therefore, one can apply alternating projections, which leads to the GS algorithm, or Dykstra's alternating projections, which results in Fienup's algorithm when the step size is set to one [23].

Due to the non-convexity of the quadratic equations, neither algorithm is guaranteed to find a solution; nonetheless, Fienup's approach has been observed to work successfully in converging to a point that satisfies both set of constraints. When the measurements are corrupted by noise, Fienup's algorithm does not in general converge. An alternative interpretation of the GS algorithm shows that it monotonically decreases the cost function of the following optimization problem (which is different from (2)),

$$\underset{\boldsymbol{x}\in\mathbb{C}^N,\boldsymbol{\psi}\in\mathbb{C}^M}{\text{minimize}} \quad \left\|\mathsf{diag}\left\{\sqrt{\boldsymbol{b}}\right\}\boldsymbol{\psi}-\boldsymbol{F}_M\boldsymbol{x}\right\|^2 \tag{3}$$
$$\text{subject to} \quad |\psi_m|=1, \quad m=0,1,...,M-1.$$

More recently, the general phase retrieval problem has been recognized as a non-convex *quadratically constrained quadratic program* (QCQP), for which the prevailing approach is to use semidefinite relaxation [24] to obtain a lower bound on the optimal value of (2). In the field of phase retrieval, this procedure is known as *PhaseLift* [18]. Specifically, under a Gaussian noise setting, PhaseLift solves

$$\underset{\boldsymbol{X}\in\mathbb{H}_+^N}{\text{minimize}} \quad \sum_{m=0}^{M-1}\left(b_m-\mathrm{tr}\left\{\boldsymbol{f}_m\boldsymbol{f}_m^H\boldsymbol{X}\right\}\right)^2+\lambda\mathrm{tr}\left\{\boldsymbol{X}\right\}, \tag{4}$$

where $\mathbb{H}_+^N$ denotes the set of Hermitian positive semidefinite matrices of size $N\times N$, $\boldsymbol{f}_m^H$ is the mth row of $\boldsymbol{F}_M$, and the term $\lambda\mathrm{tr}\left\{\boldsymbol{X}\right\}$ is used to encourage the solution to be low-rank. Problem (4) can be cast as an SDP and solved in polynomial time. If the solution of (4), denoted as $\boldsymbol{X}_\star$, turns out to be rank one, then we also obtain the optimal solution of the original problem (2) by extracting the rank one component of $\boldsymbol{X}_\star$. However, for general measurement vectors PhaseLift is not guaranteed to yield a rank one solution, especially when the measurement $\boldsymbol{b}$ is noisy. In that case PhaseLift resorts to sub-optimal solutions, for example by taking the first principal component of $\boldsymbol{X}_\star$, possibly refined by a traditional method like the GS algorithm. An SDP relaxation for the alternative formulation (3) is proposed in [25], and referred to as *PhaseCut*.

In the next section we show that despite the nonconvexity of (2), we can find an optimal solution in polynomial time using semidefinite relaxation. Since there is no uniqueness in 1D phase retrieval, there are many possible solutions even in the noise free setting. Among all solutions, we extract the minimum phase vector that minimizes the least-squares error. We will also suggest an alternative to the SDP formulation, that avoids squaring the number of variables like PhaseLift. Next we will show how any vector can be modified to be minimum phase by adding a sufficiently large impulse to it. This then paves the way to recovery of arbitrary 1D signals efficiently and provably optimally from their Fourier magnitude.

3   Convex reformulation

In this section we show how (2) can be optimally solved in polynomial time, despite its non-convex formulation. Similar results have been shown in the application of multicast beamforming under far-field line-of-sight propagation conditions [26], and more generally for non-convex QCQPs with Toeplitz quadratics [27, 28].

3.1 Phase retrieval using auto-correlation

We begin by reformulating the Fourier phase retrieval problem in terms of the auto-correlation function. Consider the mth entry of $|\boldsymbol{F}_M\boldsymbol{x}|^2$:

$$
\begin{aligned}
|\boldsymbol{f}_m^H\boldsymbol{x}|^2 &= \sum_{n=0}^{N-1}\phi^{nm}x_n\sum_{\nu=0}^{N-1}\phi^{-\nu m}x_\nu^* \\
&= \sum_{n=0}^{N-1}\sum_{\nu=0}^{N-1}x_n x_\nu^*\phi^{(n-\nu)m} \\
&= \sum_{n=0}^{N-1}\sum_{k=n+1-N}^{n}x_n x_{n-k}^*\phi^{km} \\
&= \sum_{k=1-N}^{N-1}\phi^{km}\sum_{n=\max(k,0)}^{\min(N-1+k,N-1)}x_n x_{n-k}^* \\
&= \sum_{k=1-N}^{N-1}\phi^{km}r_k,
\end{aligned}
$$

where

$$
r_k = \sum_{n=\max(k,0)}^{\min(N-1+k,N-1)}x_n x_{n-k}^*, \tag{5}
$$

$$
k = 1-N, ..., -1, 0, 1, ..., N-1,
$$

is the k-lag auto-correlation of $\boldsymbol{x}$. Let us define

$$
\tilde{\boldsymbol{r}} = [\ r_{1-N}\ ...\ r_{-1}\ r_0\ r_1\ ...\ r_{N-1}\ ]^T.
$$

We first observe that $|\boldsymbol{f}_m^H\boldsymbol{x}|^2$, originally quadratic with respect to $\boldsymbol{x}$, is now linear in $\tilde{\boldsymbol{r}}$. Moreover, by definition $r_k = r_{-k}^*$. Removing the redundancy, we let

$$
\boldsymbol{r} = [\ r_0\ r_1\ ...\ r_{N-1}\ ]^T,
$$

and write

$$
|\boldsymbol{F}_M\boldsymbol{x}|^2 = \Re\left\{\boldsymbol{F}_M\tilde{\boldsymbol{I}}\boldsymbol{r}\right\},
$$

where $\tilde{\boldsymbol{I}} = \mathsf{diag}\left\{[\ 1\ 2\ 2\ ...\ 2\ ]\right\}$, and $\Re\{\cdot\}$ takes the real part of its argument. We can then rewrite (2) as

$$
\begin{aligned}
\underset{\boldsymbol{r}\in\mathbb{C}^N}{\text{minimize}}\quad &\left\|\boldsymbol{b} - \Re\left\{\boldsymbol{F}_M\tilde{\boldsymbol{I}}\boldsymbol{r}\right\}\right\|^2 \\
\end{aligned} \tag{6}
$$
$$
\text{subject to}\quad \boldsymbol{r}\text{ is a finite auto-correlation sequence.}
$$

The abstract constraint imposed on $\boldsymbol{r}$ in (6) is to ensure that there exists a vector $\boldsymbol{x}\in\mathbb{C}^N$ such that (5) holds.

The representation in (5) of the auto-correlation sequence is nonconvex (in $\boldsymbol{r}$ and $\boldsymbol{x}$ jointly). In the next subsection we consider convex reformulations of this constraint based on [29]. In Section 4 we discuss how to obtain $\boldsymbol{x}$ from $\boldsymbol{r}$ in a computationally efficient way.

3.2 Approximate characterization of a finite auto-correlation sequence

Denote the DTFT of $\tilde{r}$ by

$$R(e^{j\omega}) = \sum_{k=1-N}^{N-1} r_k e^{-j\omega k}.$$

Since $\tilde{r}$ is conjugate symmetric, $R(e^{j\omega})$ is real for all $\omega \in [0, 2\pi]$. The sequence $\tilde{r}$ (or equivalently r) is a finite auto-correlation sequence if and only if [30]

$$R(e^{j\omega}) \geq 0, \quad \omega \in [0, 2\pi], \tag{7}$$

which is a collection of infinitely many linear inequalities. A natural way to approximately satisfy (7) is to sample this infinite set of inequalities, and replace (7) by a finite (but large) set of L linear inequalities

$$R(e^{j2\pi l/L}) \geq 0, \quad l = 0, 1, ..., L - 1.$$

In matrix form, these constraints can be written as (similar to the expression for $|\boldsymbol{F}_M \boldsymbol{x}|^2$ that we derived before)

$$\Re\left\{\boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r}\right\} \geq 0,$$

where $\boldsymbol{F}_L$ is the first N columns of the L-point DFT matrix.

Using this result, we can approximately formulate problem (6) as

$$\begin{aligned}
\underset{\boldsymbol{r} \in \mathbb{C}^N}{\text{minimize}} \quad & \left\| \boldsymbol{b} - \Re\left\{\boldsymbol{F}_M \tilde{\boldsymbol{I}} \boldsymbol{r}\right\} \right\|^2 \\
\text{subject to} \quad & \Re\left\{\boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r}\right\} \geq 0,
\end{aligned} \tag{8}$$

with a sufficiently large L. In practice, the approximate formulation works very well for $L \geq 20N$, a modest (and typically a constant times) increase in the signal dimension. However, it is not exact—one can satisfy the constraints in (8), yet still have $R(e^{j\omega}) < 0$ for some ω.

3.3 Exact parameterization of a finite auto-correlation sequence

It turns out that it is possible to characterize the infinite set of inequalities (7) through a finite representation, via an auxiliary positive semidefinite matrix. One way is to use a $N \times N$ positive semidefinite matrix to parameterize $\boldsymbol{r}$ [28]

$$\begin{aligned}
r_k &= \text{tr}\left\{\boldsymbol{T}_k \boldsymbol{X}\right\}, \quad k = 0, 1, ..., N - 1, \\
\boldsymbol{X} &\succeq 0,
\end{aligned}$$

where $\boldsymbol{T}_k$ is the kth elementary Toeplitz matrix of appropriate size, with ones on the kth sub-diagonal, and zeros elsewhere, and $\boldsymbol{T}_0 = \boldsymbol{I}$. Another way, which can be derived from the Kalman-Yakubovich-Popov (KYP) lemma [31], uses a $(N-1) \times (N-1)$ matrix $\boldsymbol{P}$ to characterize $\boldsymbol{r}$ as a finite correlation sequence

$$\begin{bmatrix} r_0 & \boldsymbol{r}_{1:N-1}^H \\ \boldsymbol{r}_{1:N-1} & \boldsymbol{P} \end{bmatrix} - \begin{bmatrix} \boldsymbol{P} & \boldsymbol{0} \\ \boldsymbol{0} & 0 \end{bmatrix} \succeq 0,$$

where $\boldsymbol{r}_{1:N-1} = [\, r_1 \ r_2 \ ... \ r_{N-1} \,]^T$. It can be shown that the two formulations are equivalent [32].

6

Adopting the first characterization, we may rewrite (6) as

$$\underset{\boldsymbol{r}\in\mathbb{C}^N,\boldsymbol{X}\in\mathbb{H}_+^N}{\text{minimize}} \quad \left\|\boldsymbol{b} - \Re\left\{\boldsymbol{F}_M\tilde{\boldsymbol{I}}\boldsymbol{r}\right\}\right\|^2$$

$$\text{subject to} \quad r_k = \text{tr}\left\{\boldsymbol{T}_k\boldsymbol{X}\right\}, \quad k = 0, 1, ..., N-1. \tag{9}$$

It is easy to see that we can actually eliminate the variable $\boldsymbol{r}$ from (9), ending up with the PhaseLift formulation (4) without the trace regularization. This result is proven in Appendix A.

The implication behind the above analysis is that, even though PhaseLift does not produce a rank one solution in general[1], there always exists a point in $\mathbb{C}^N$ that attains the cost provided by the SDP relaxation. In other words, the seemingly non-convex problem (2) can be solved in polynomial time.

The trace parameterization based formulation (9) exactly characterizes (6), with the price that now the problem dimension is $O(N^2)$, a significant increase comparing to that of (8). In Section 6 we will derive an efficient alternative based on ADMM. In the next section we show how to extract a vector $\boldsymbol{x}$ from the auto-correlation $\boldsymbol{r}$ that solves (9).

4   Spectral factorization

After we solve (6), either approximately via (8) or exactly by (9), we obtain the auto-correlation of the optimal solution of the original problem (2). The remaining question is how to find a vector $\boldsymbol{x}$ that generates such an auto-correlation, as defined in (5). This is a classical problem known as *spectral factorization* (SF) in signal processing and control. There exist extensive surveys on methods and algorithms that solve this problem, e.g., [33], [34, Appendix], and [28, Appendix B]. In this section, we first derive an algebraic solution for SF, which also explains the non-uniqueness of 1D Fourier phase retrieval, and reviews the notion of *minimum phase*. We then survey two practical methods for SF, which we will rely on when developing CoRK, a highly scalable phase-retrieval algorithm.

4.1   Algebraic solution

The DFT of an arbitrary signal $\boldsymbol{x} \in \mathbb{C}^N$ can be obtained by sampling its z-transform on the unit circle $|z| = 1$. Let $X(z)$ be the z-transform of $\boldsymbol{x}$, which is a polynomial of order $N-1$. It can be written in factored form as

$$X(z) = \sum_{n=0}^{N-1} x_n z^{-n} = x_0 \prod_{n=1}^{N-1}(1 - \xi_n z^{-1}), \tag{10}$$

where $\xi_1, ..., \xi_{N-1}$ are the zeros (roots) of the polynomial $X(z)$. The quadratic measurements can similarly be interpreted as sampled from the z-transform of $\tilde{\boldsymbol{r}}$ defined as

$$R(z) = |X(z)|^2 = X(z)X^*(1/z^*)$$

$$= |x_0|^2 \prod_{n=1}^{N-1}(1 - \xi_n z^{-1})(1 - \xi_n^* z).$$

---

[1]In fact, it has been shown that if the interior-point algorithm is used to solve an SDP, it will always generate a solution that is of maximal rank [24].

As we can see, the zeros of $R(z)$ always come in conjugate reciprocal pairs. Therefore, given $R(z)$, we cannot determine whether a zero ξ or its conjugate reciprocal $(\xi^*)^{-1}$ is a root of $X(z)$, which is the reason that $\boldsymbol{x}$ cannot be reconstructed from $R(z)$. In other words, for a given signal $\boldsymbol{x}$, we can find the zeros of its z-transform as in (10), take the conjugate reciprocal of some of them, and then take the inverse z-transform to obtain another signal $\boldsymbol{y}$ [35]. If we re-scale $\boldsymbol{y}$ to have the same ℓ_2 norm as $\boldsymbol{x}$, then it is easy to verify that

$$|\boldsymbol{F}_M\boldsymbol{x}|^2 = |\boldsymbol{F}_M\boldsymbol{y}|^2,$$

no matter how large M is, even though clearly $\boldsymbol{x}$ and $\boldsymbol{y}$ are not equal.

Traditionally, this problem is often seen in design problems where uniqueness is not important, e.g., FIR filter design [34] and far-field multicast beamforming [26]. There, it is natural (from the maximal energy dissipation point of view) to pick the zeros to lie within the unit circle[2], yielding a so-called *minimum phase* signal.

The above analysis provides a direct method for SF. For a given auto-correlation sequence $\boldsymbol{r}$, we first calculate the roots of the polynomial

$$R(z) = r_0 + \sum_{k=1}^{N-1} \left(r_k z^{-k} + r_k^* z^k \right).$$

Because $\boldsymbol{r}$ is a valid correlation sequence, the roots come in conjugate reciprocal pairs. Therefore, we can pick the $N-1$ roots that are inside the unit circle, expand the expression, and then scale it to have ℓ_2 norm equal to $\sqrt{r_0}$.

Numerically, the roots of $R(z)$ can be found by calculating the eigenvalues of the following companion matrix:

$$\begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & & \vdots \\ \vdots & \vdots & & \ddots & 0 \\ 0 & 0 & \cdots & 0 & 1 \\ -\frac{r_{N-1}}{r_{N-1}^*} & \cdots & -\frac{r_0}{r_{N-1}^*} & \cdots & -\frac{r_{N-2}^*}{r_{N-1}^*} \end{bmatrix}.$$

However, when expanding the factored form to get the coefficients of the polynomial, the procedure is very sensitive to roundoff error, which quickly becomes significant as N approaches 64.

4.2 SDP-based method for SF

To obtain a more stable method numerically, we can use an SDP approach to SF.

For a valid correlation $\boldsymbol{r}$, it is shown in [28, Chapter 2.6.1] that the solution of the following SDP

$$\begin{aligned} \underset{\boldsymbol{X} \in \mathbb{H}_+^N}{\text{maximize}} \quad & X_{00} \\ \text{subject to} \quad & r_k = \text{tr}\left\{\boldsymbol{T}_k \boldsymbol{X}\right\}, \\ & k = 0, 1, ..., N-1, \end{aligned} \tag{11}$$

---

[2]If a zero ξ lie exactly on the unit circle, then $(\xi^*)^{-1} = \xi$, meaning $R(z)$ has a double root at that point, so we simply pick one as the zero for the signal.

is always rank one. In addition, its rank one component generates the given correlation and is minimum phase. Algorithms for SDP are numerically stable, although the complexity could be high if we use a general-purpose SDP solver.

The constraints in (11) are the same as the convex reformulation (9), which according to Appendix A is equivalent to PhaseLift without the trace regularization. Therefore, we may consider instead solving the problem

$$\underset{\boldsymbol{X} \in \mathbb{H}_+^N}{\text{minimize}} \quad \sum_{m=0}^{M-1} \left(b_m - \text{tr}\left\{\boldsymbol{f}_m \boldsymbol{f}_m^H \boldsymbol{X}\right\}\right)^2 - \lambda X_{00}. \tag{12}$$

In Appendix B we show that for a suitable value of λ, the solution of (12) is guaranteed to be rank one, and it attains the optimal least-squares error.

The formulation (12) allows to use customized SDP solvers to obtain a solution efficiently. For example, PhaseLift was originally solved using TFOCS [36], which applies various modern first-order methods to minimize convex functions composed of a smooth part and a non-smooth part that has a simple proximity operator. For problem (12), the cost function is the smooth part, and the non-smooth part is the indicator function for the cone of semidefinite matrices.

It is now widely accepted that the trace of a semidefinite matrix, or nuclear norm for a general matrix, encourages the solution to be low rank [37]. However, interestingly, in our setting, the correct regularization is in fact $-\lambda X_{00}$ which guarantees the solution to be rank one for an appropriate choice of λ.

4.3   Kolmogorov's method for SF

Kolmogorov proposed the following method for SF that is both efficient and numerically stable [31]. In what follows, we assume that $R(z)$ contains no zeros on the unit circle.

Consider taking the logarithm of the z-transform of $\boldsymbol{x}$. Every zero (and pole, which we do not have since $\boldsymbol{x}$ is finite length) of $X(z)$ then becomes a pole of $\log X(z)$. Assuming the roots of $X(z)$ lie strictly inside the unit circle, this implies that there exists a region of convergence (ROC) for $\log X(z)$ that contains the unit circle $|z| = 1$ and infinity $|z| = \infty$. This in turn means that $\log X(z)$ is unilateral

$$\log X(z) = \sum_{n=0}^{\infty} \alpha_n z^{-n}.$$

Evaluating $\log X(z)$ at $z = e^{j\omega}$, yields

$$\Re\left\{\log X(e^{j\omega})\right\} = \sum_{n=0}^{\infty} \alpha_n \cos \omega n,$$

$$\Im\left\{\log X(e^{j\omega})\right\} = -\sum_{n=0}^{\infty} \alpha_n \sin \omega n,$$

implying that $\Re\left\{\log X(e^{j\omega})\right\}$ and $\Im\left\{\log X(e^{j\omega})\right\}$ are Hilbert transform pairs. Moreover, since we are given $R(z) = |X(z)|^2$, we have that

$$\Re\left\{\log X(e^{j\omega})\right\} = \frac{1}{2} \log R(e^{j\omega}).$$

9

We can therefore calculate $\Im\left\{\log X(e^{j\omega})\right\}$ from the Hilbert transform of $\Re\left\{\log X(e^{j\omega})\right\}$, and reversely reconstruct a signal $\boldsymbol{x}$ that generates the given correlation $\boldsymbol{r}$ and is minimum phase.

In practice, all of the aforementioned transforms may be well approximated by a DFT with sufficiently large length L. The detailed procedure of Kolmogorov's method then becomes:

1. compute the real part

$$\boldsymbol{\gamma} = \frac{1}{2}\log \Re\left\{\boldsymbol{F}_L \tilde{\boldsymbol{I}}\boldsymbol{r}\right\};$$

2. compute the imaginary part by taking the Hilbert transform of $\boldsymbol{\gamma}$, approximated by a DFT

$$\boldsymbol{\phi} = \boldsymbol{F}_L \boldsymbol{\gamma},$$

$$\varphi_n = \begin{cases} 0, & n = 0, L/2, \\ -j\phi_n, & n = 1, 2, ..., L/2 - 1, \\ j\phi_n, & n = L/2 + 1, ..., L - 1, \end{cases}$$

$$\boldsymbol{\eta} = \frac{1}{L}\boldsymbol{F}_L^H \boldsymbol{\varphi};$$

3. compute $(1/L)\boldsymbol{F}_L^H \exp(\boldsymbol{\gamma} - j\boldsymbol{\eta})$, and take the first N elements as the output.

Kolmogorov's method generates a valid output as long as in the first step we have that $\Re\left\{\boldsymbol{F}_L \tilde{\boldsymbol{I}}\boldsymbol{r}\right\} \geq 0$, so that $\boldsymbol{\gamma}$ is real. This fits well with the approximate formulation (8), since if we choose the same L for both (8) and Kolmogorov's method, then it is guaranteed to be able to run without encountering a syntax error. Obviously, to obtain a solution with high accuracy, we need L to be sufficiently large; empirically we found $L \geq 20N$ to be good enough in practice. Computationally, this approach requires computing four L-point (inverse-) FFTs with complexity $O(L \log L)$, a very light computation burden even for large L.

5   A new measurement system

So far we have seen that an arbitrary signal cannot be uniquely determined from the magnitude of its over-sampled 1D Fourier measurements, because there always exists a minimum phase signal that yields the same measurements. Nonetheless, a least-squares estimate can be efficiently calculated by solving either (8) or (9) followed by spectral factorization, leading to an optimal solution that is minimum phase. If the true signal $\boldsymbol{x}$ is indeed minimum phase, then we can optimally estimate it in polynomial-time. However, the minimum phase property is not a natural assumption to impose on a signal in general.

We propose to resolve this ambiguity by deliberately making the signal minimum phase before taking the quadratic measurements, so that the augmented signal is uniquely identified in polynomial time. The true signal is then recovered easily by reversing this operation.

For an arbitrary complex signal $\boldsymbol{s}$, we suggest adding δ in front of $\boldsymbol{s}$ before taking measurements, where δ satisfies that $|\delta| \geq \|\boldsymbol{s}\|_1$. Denote the augmented signal as $\boldsymbol{s}_{\min}$. The following proposition shows that $\boldsymbol{s}_{\min}$ is minimum phase.

Proposition 1. *Consider an arbitrary complex signal*

$$\boldsymbol{s} = \begin{bmatrix} s_0 & s_1 & ... & s_{N-1} \end{bmatrix}^T.$$

10

Then the augmented signal

$$\boldsymbol{s}_{min} = [\; \delta \; s_0 \; ... \; s_{N-1} \;]^T, \tag{13}$$

where $|\delta| \geq \|\boldsymbol{s}\|_1$, *is minimum phase.*

Proof. To establish the result we need to show that the zeros of the z-transform of $\boldsymbol{s}_{\min}$

$$\delta + s_0 z^{-1} + ... + s_{N-1} z^{-N},$$

or equivalently the roots of the polynomial

$$V(z) = z^N + \frac{s_0}{\delta} z^{N-1} + ... + \frac{s_{N-1}}{\delta},$$

all lie inside the unit circle. To this end, we rely on the following lemma.

Lemma 1. *[38, Theorem 1] Let ζ be a zero of the polynomial*

$$z^N + c_{N-1} z^{N-1} + ... + c_1 z + c_0,$$

where $c_0, ..., c_{N-1} \in \mathbb{C}$ *and* N *is a positive integer. Then*

$$|\zeta| \leq \max \left\{ 1, \sum_{n=0}^{N-1} |c_n| \right\}.$$

Substituting the coefficients of $V(z)$ into the inequality in Lemma 1 establishes the result. $\square$

Conceptually, the approach we propose is very simple: all we need is a way to over-estimate the ℓ_1 norm of the target signal, and a mechanism to insert an impulse in front of the signal before taking quadratic measurements. For example, if we assume each element in $\boldsymbol{s}$ comes from a complex Gaussian distribution with variance σ^2, then we know that the probability that the magnitude of one element exceeds 3σ is almost negligible; therefore, we can simply construct $\boldsymbol{s}_{\min}$ by setting $\delta = 3\sigma N$, resulting in $\boldsymbol{s}_{\min}$ being minimum phase with very high probability.

Our approach can be used with a number of measurements M, as small as $2N$. Indeed, consider the equivalent reformulation (6), in which the measurements $\boldsymbol{b}$ are linear with respect to $\boldsymbol{r}$. From elementary linear algebra, we know that N complex numbers can be uniquely determined by as few as $2N$ real linearly independent measurements, even without the specification that $\boldsymbol{r}$ is a finite correlation sequence. From a unique $\boldsymbol{r}$, a unique minimum phase $\boldsymbol{s}_{\min}$ can be determined using SF.

The impulse may also be appended at the end of the signal $\boldsymbol{s}$, resulting in a *maximum phase* signal, meaning all the zeros of its z-transform are outside the unit circle.

Proposition 2. *For an arbitrary complex signal*

$$\boldsymbol{s} = [\; s_0 \; s_1 \; ... \; s_{N-1} \;]^T,$$

the augmented signal

$$\boldsymbol{s}_{max} = [\; s_0 \; ... \; s_{N-1} \; \delta \;]^T,$$

where $|\delta| \geq \|\boldsymbol{s}\|_1$, *is maximum phase.*

Proof. The z-transform of $\boldsymbol{s}_{\max}$ is

$$S(z) = s_0 + s_1 z^{-1} + ... + s_{N-1} z^{1-N} + \delta z^{-N}.$$

Let $\tilde{z} = z^{-1}$. Then

$$S(\tilde{z}) = s_0 + s_1 \tilde{z} + ... + s_{N-1} \tilde{z}^{N-1} + \delta \tilde{z}^N,$$

which is a polynomial in $\tilde{z}$ and, according to Lemma 1, has all its roots inside the unit circle. Taking the reciprocal, this means that all the zeros of the z-transform of $\boldsymbol{s}_{\max}$ lie outside the unit circle. $\square$

It is easy to show that for a maximum phase signal

$$\boldsymbol{s}_{\max} = [\ s_0\ ...\ s_{N-1}\ \delta\ ]^T,$$

its equivalent minimum phase signal is

$$[\ \delta^*\ s_{N-1}^*\ s_{N-2}^*\ ...\ s_0^*\ ]^T.$$

This means that if we measure the intensity of the Fourier transform of $\boldsymbol{s}_{\max}$ instead, we can still uniquely recover $\boldsymbol{s}$ via solving (6) followed by SF, take the conjugate reversal of the solution, and then delete the first element to obtain an estimate.

Furthermore, from the analysis above, we can extend our measuring technique so that the impulse is added away from the signal:

$$[\ \delta\ 0\ ...\ 0\ s_0\ ...\ s_{N-1}\ ]^T, \tag{14}$$

or

$$[\ s_0\ ...\ s_{N-1}\ 0\ ...\ 0\ \delta\ ]^T, \tag{15}$$

with an arbitrary number of zeros between δ and $\boldsymbol{s}$. The resulting augmented signal is still minimum/maximum phase, thus can be uniquely identified from the intensity of its Fourier transform. However, in this case more measurements are required for identifiability.

It was briefly mentioned in [39] without proof that if the signal has a strong first component, then it is minimum phase. Here we propose *adding* an impulse as a means of restoring identifiability for *arbitrary* signals. This concept is very similar to *holography*, which was invented by Gabor in 1948 [22], and was awarded the Nobel Prize in Physics in 1971. One form of holography, spectral interferometry, inserts an impulse exactly in the form of (14) or (15), and then measures the squared magnitude of the Fourier transform of the combined signal [40, 41]. Despite the similarities in obtaining the measurements, the theory behind our approach and holography is very different, in several ways:

1. Holography constrains the length of the zeros between $\boldsymbol{s}$ and δ to be at least N, the length of the signal, which is not required in our method; in fact, in terms of the number of measurements needed, the shorter the better.

2. We require $|\delta| \geq \|\boldsymbol{s}\|_1$, whereas holography does not have any restriction on the energy of the impulse.

3. We provide an optimization based framework to recover the underlying signal directly in a robust fashion by exploiting the minimum phase property which is not part of the framework in holography.

6 Scalable algorithms

In this section, we briefly describe how to directly use existing PhaseLift solvers to obtain guaranteed rank one solutions efficiently. Due to the extra memory required using SDP-relaxation methods for large-scale problems, we then design a new algorithm called CoRK that solves the approximate formulation (8) with high efficiency, in terms of both time and memory.

6.1 SDP based algorithms

In the original PhaseLift paper [18], the authors used TFOCS [36] to solve the SDP (4), which is a relatively efficient way to solve large scale SDPs. As we discussed in Section 4.2, we can ensure a rank one minimum phase solution by solving (12) with an appropriate choice of λ. We find such a λ via bisection: first we solve (12) with $\lambda = 0$ to get the optimal value of the fitting term; next we solve it with a λ large enough that the solution is guaranteed to be rank one, but with a possibly larger fitting error; finally we bisect λ until both criteria are met. Since TFOCS allows to specify initializations, after we solve (12) for the first time, subsequent evaluations can be performed much more rapidly via warm start. Moreover, there is a wide range of λ that result in both a rank one solution and an optimal fitting error, so that the overall increase in time required is small compared to that of solving a single PhaseLift problem.

The disadvantage of SDP based algorithms is that one cannot avoid lifting the problem dimension to $O(N^2)$. If N is moderately large, on the order of a few thousand, then it is very difficult to apply SDP based methods in practice.

6.2 CoRK: An efficient ADMM method

We now propose a new algorithm for solving the approximate formulation (8), which avoids the dimension lifting, and easily handles problem sizes up to millions. Since (8) is a convex quadratic program, there are plenty of algorithms to solve it reliably. For example, a general purpose interior-point method solves it with worst case complexity $O(N^{3.5})$ without taking any problem structure into account. Nevertheless, noticing that all of the linear operations in (8) involve DFT matrices, we aim to further recude the complexity by exploiting the FFT operator.

We propose solving (8) using the alternating direction method of multipliers (ADMM) [42]. To apply ADMM, we first rewrite (8) by introducing an auxiliary variable $\boldsymbol{z} \in \mathbb{R}^L$:

$$\begin{aligned} \underset{\boldsymbol{r} \in \mathbb{C}^N, \boldsymbol{z} \in \mathbb{R}^L}{\text{minimize}} \quad & \left\| \boldsymbol{b} - \Re\left\{ \boldsymbol{F}_M \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} \right\|^2 + \mathcal{I}_+(\boldsymbol{z}) \\ \text{subject to} \quad & \Re\left\{ \boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} = \boldsymbol{z}, \end{aligned} \tag{16}$$

where

$$\mathcal{I}_+(\boldsymbol{z}) = \begin{cases} 0, & \boldsymbol{z} \geq 0, \\ +\infty, & \text{otherwise}, \end{cases}$$

is the indicator function of the non-negative orthant in $\mathbb{R}^L$. Treating $\boldsymbol{r}$ as the first block of variables, $\boldsymbol{z}$ as the second, and $\Re\{\boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r}\} = \boldsymbol{z}$ as the coupling linear equality constraint with scaled Lagrange

multiplier $\boldsymbol{u}$, we arrive at the following iterative updates:

$$
\boxed{
\begin{aligned}
\boldsymbol{r} &\leftarrow \frac{1}{M+\rho L}\left(\boldsymbol{F}_M^H \boldsymbol{b} + \rho \boldsymbol{F}_L^H (\boldsymbol{z}-\boldsymbol{u})\right), \\
\boldsymbol{z} &\leftarrow \max\left(0, \Re\left\{\boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r}\right\} + \boldsymbol{u}\right), \\
\boldsymbol{u} &\leftarrow \boldsymbol{u} + \Re\left\{\boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r}\right\} - \boldsymbol{z},
\end{aligned}
}
\tag{17}
$$

where ρ is a constant described below. The derivation of the algorithm is given in Appendix C.

All the operations in (17) are taken element-wise or involve FFT computations of length L. This leads to a complexity $O(L \log L)$, and effectively $O(N \log N)$ if L is chosen as $O(N)$. In contrast, a naive implementation of ADMM for quadratic programming requires solving a least-squares problem in every iteration [43], leading to an $O(MN^2)$ per-iteration complexity, which is significantly higher.

In terms of convergence rate, it is shown in [43, Theorem 3] that our ADMM approach will converge linearly, for all values of $\rho > 0$. The same reference provides an optimal choice of ρ that results in the fastest convergence rate, for the case when there are fewer inequality constraints than variables in the first block. However, this requirement unfortunately is not fulfilled for problem (8). Inspired by the choice of ρ in [43], we found empirically that by setting $\rho = M/L$, the proposed iterates (17) converge very fast, and at a rate effectively independent of the problem dimension. For practical purposes, the convergence rate we have achieved is good enough (typically in less than 100 iterations), but there exists pre-conditioning methods to further accelerate it [44].

As explained in Section 4.3, this formulation blends well with Kolmogorov's method for spectral factorization. We thus refer to our approach for solving the 1D Fourier phase retrieval problem (2) as the *auto-<u>cor</u>relation <u>r</u>etrieval – <u>K</u>olmogorov factorization* (**CoRK**) algorithm.

7   Summary of the proposed method

Throughout the paper, we divided 1D Fourier phase retrieval into several steps (measurement, formulation, algorithms, etc.), and for each step discussed several solution methods. For clarity and practical purposes, we summarize our approach below based on specific choices we found to be efficient.

Given an arbitrary vector $\boldsymbol{s}$, we propose the following steps if it is possible to insert an impulse to the signal before measuring its Fourier intensity; otherwise, only step 2 will be invoked, in which case we are guaranteed to solve the maximum likelihood formulation optimally, but the estimation error may still be large.

1. Construct $\boldsymbol{s}_{\min}$ by inserting δ in front of $\boldsymbol{s}$, i.e.,

$$
\boldsymbol{s}_{\min} = [\ \delta\ s_0\ s_1\ ...\ s_{N-1}\ ]^T,
$$

such that $|\delta| > \|\boldsymbol{s}\|_1$. Take the M-point DFT of $\boldsymbol{s}_{\min}$, where $M > 2N$, and measure its squared magnitude $\boldsymbol{b} = |\boldsymbol{F}_M \boldsymbol{s}_{\min}|^2$.

2. Apply CoRK as follows:

(a) Formulate the least-squares problem with respect to the auto-correlation of $\boldsymbol{s}_{\min}$ as in (8),

$$\underset{\boldsymbol{r}\in\mathbb{C}^{N+1}}{\text{minimize}} \quad \left\| \boldsymbol{b} - \Re\left\{ \boldsymbol{F}_M \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} \right\|^2$$

$$\text{subject to} \quad \Re\left\{ \boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} \geq 0,$$

by picking L as the smallest power of 2 that is greater than $32N$. Solve this problem using the iterative algorithm (17), repeated here by setting $\rho = M/L$:

$$\boldsymbol{r} \leftarrow \frac{1}{2}\left(\frac{1}{M}\boldsymbol{F}_M^H \boldsymbol{b} + \frac{1}{L}\boldsymbol{F}_L^H(\boldsymbol{z} - \boldsymbol{u}) \right),$$

$$\boldsymbol{z} \leftarrow \max\left(0, \Re\left\{ \boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} + \boldsymbol{u} \right),$$

$$\boldsymbol{u} \leftarrow \boldsymbol{u} + \Re\left\{ \boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} - \boldsymbol{z}.$$

(b) Extract the minimum phase signal that generates the correlation $\boldsymbol{r}$ by Kolmogorov's method, using the same L, as follows (also given in Section 4.3):

$$\boldsymbol{\gamma} = \frac{1}{2}\log \Re\left\{ \boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r} \right\},$$

$$\boldsymbol{\phi} = \boldsymbol{F}_L \boldsymbol{\gamma},$$

$$\varphi_n = \begin{cases} 0, & n = 0, L/2, \\ -j\phi_n, & n = 1, 2, ..., L/2 - 1, \\ j\phi_n, & n = L/2 + 1, ..., L - 1, \end{cases}$$

$$\boldsymbol{\eta} = \frac{1}{L}\boldsymbol{F}_L^H \boldsymbol{\varphi},$$

$$\boldsymbol{x} = (1/L)\boldsymbol{F}_L^H \exp(\boldsymbol{\gamma} - j\boldsymbol{\eta}).$$

3. Obtain an estimate $\hat{\boldsymbol{s}}$ via

$$\hat{s}_n = x_{n+1}, \quad n = 0, 1, ..., N - 1.$$

Often, we have prior information that the signal of interest $\boldsymbol{s}$ is real. This implies that the auto-correlation is also real so that we need to restrict the domain of $\boldsymbol{r}$ and/or $\boldsymbol{X}$ in problem (8) or (12) to be real. For the z-transform of a real signal, the zeros are either real or come in conjugate pairs, which does not help the non-uniqueness of the solution if we directly measure the 1D Fourier magnitude. In Kolmogorov's SF method, if the input $\boldsymbol{r}$ is real and a valid correlation, then $\boldsymbol{\gamma} + j\boldsymbol{\eta}$ is conjugate symmetric, so that it is guaranteed to output a real valued signal. As for the new measurement system, all the claims made for constructing minimum/maximum phase signals still hold, when the signal is restricted to be real. In Appendix C we show how to modify the ADMM method (17) to accommodate real signals by adding an additional projection to the real domain in the $\boldsymbol{r}$ update.

When $\boldsymbol{x}$ is real, the DFT has conjugate symmetry

$$X(e^{-j2\pi m/M}) = X^*(e^{-j2\pi(M-m)/M}),$$

meaning for the squared magnitude, $b_m = b_{M-m}$. Therefore, if we take M samples between $[0, 2\pi]$, then only the first $M/2$ measurements provide useful information. Consequently, we still need $M > 2N$ measurements sampled between $[0, 2\pi]$ to ensure identifiability.

8 Simulations

We now demonstrate the algorithms and the new measurement system we proposed via simulations. We first show that both algorithms, iteratively solving (12) while bisecting λ and solving (8) followed by Kolmogorov's method, are able to solve the original problem (2) optimally, meaning the cost attains the lower bound given by PhaseLift. Then, we show that applying the new measurement system proposed in Section 5 we can uniquely identify an arbitrary 1D signal, whereas directly measuring the over-sampled Fourier intensity does not recover the original input, regardless of the algorithms being used. All simulations are performed in MATLAB on a Linux machine.

8.1 Minimizing the least-squares error

We first test the effectiveness of the proposed algorithms on random problem instances. Fixing $N = 128$, we randomly set M as an integer between $[2N, 8N]$, and generate $\boldsymbol{b}$ from an i.i.d. uniform distribution between $[0, 1]$. We compare the following algorithms:

- **PhaseLift.** Using TFOCS [36] to solve problem (4) with $\lambda = 0$. This in general does not give a rank one solution, therefore only serves as a theoretical lower bound on the minimal least-squares error (2).

- **PhaseLift-PC.** The leading principal component of the plain PhaseLift solution.

- **PhaseLift-SF.** Iteratively solving (12) while bisecting λ until an equivalent rank one solution for (4) is found, again using TFOCS. Each time TFOCS is initialized with the solution from the previous iteration for faster convergence. As we proved, this method is guaranteed to achieve the lower bound given by PhaseLift.

- **CoRK.** Proposed method summarized in the second step of Section 8. For both steps, L is set to be the smallest power of 2 that is greater than $32N$.

- **Fienup-GS.** Fienup's algorithm [45] with 1000 iterations, then refined by the GS algorithm until the error defined in (3) converges.

The optimality gaps between the minimal error in (2) obtained by the aforementioned methods and the theoretical lower bound given by PhaseLift are shown in Fig. 1; The running time of the different methods is provided in Fig. 2, for the 100 Monte-Carlo trials we tested. As we can see, PhaseLift-SF provides an optimal rank one solution, although it takes more time compared to solving one single PhaseLift problem. Since PhaseLift-SF only provides a solution that is "numerically" rank one, when we take its rank one component and evaluate the actual fitting error (2), it is still a little bit away from the theoretical PhaseLift lower bound, as shown in the red circles in Fig. 1. On the other hand, the proposed CoRK method is able to approximately solve the problem with high accuracy (a lot of times even better than PhaseLift-SF) in a very small amount of time (shorter than the standard Fienup-GS algorithm). The conventional Fienup-GS method and the leading principal component of the PhaseLift solution (in general not close to rank one) do not come close to the PhaseLift lower bound.

Figure 1: Optimality gaps from the PhaseLift lower bound in each Monte-Carlo trial, cf. Section 8.1.

Figure 2: Time consumed in each Monte-Carlo trial, cf. Section 8.1.

8.2 Estimation performance

We now verify that our proposed measurement technique, described in Section 5, is able to recover a signal s up to global phase ambiguity. As we have shown in the previous simulation, the CoRK method performs similar to those based on PhaseLift, but with much more efficiency. Therefore, we only compare CoRK and Fienup's algorithm as baselines. In each Monte-Carlo trial, we first set N as a random integer between $[1, 1024]$, and then choose M as the smallest power of 2 that is larger than $4N$. A random signal $s \in \mathbb{C}^N$ is then generated with elements drawn from i.i.d. $\mathcal{CN}(0,1)$. Two kinds of measurements are collected for performance comparison:

- **Direct:** directly measure $|\boldsymbol{F}_M \boldsymbol{s}|^2$;

- **Minimum Phase:** Construct a minimum phase signal $\boldsymbol{s}_{\min}$ as in (13) with $\delta = 3N$, and then measure $|\boldsymbol{F}_M \boldsymbol{s}_{\min}|^2$.

For the minimum phase measurements, when Fienup's algorithm is used, we add an additional step in which we use the solution to generate an auto-correlation sequence and then apply spectral factorization (Kolmogorov's method) to obtain a minimum phase signal with the same fitting error. We denote this approach by **Fienup-SF**. We employ a simple prior that the added impulse δ is real

17

Figure 3: Estimation error in each Monte-Carlo trial for the first simulation in Section 8.2.

Figure 4: Computation time in each Monte-Carlo trial (with the new measurement system) for the first simulation in Section 8.2.

and positive to resolve the global phase ambiguity: after obtaining the minimum phase solution, the result is first rotated so that the first entry is real and positive, and then this entry is deleted to obtain an estimate $\hat{s}$. For direct measurements, the estimation error is defined as

$$\min_{|\psi|=1} \|s - \psi\hat{s}\|^2.$$

The estimation error in each of the 100 Monte-Carlo trials is shown in Fig. 3. We obtain perfect signal recovery when the new measurement system is used together with the CoRK recovery algorithm. This is never the case for direct measurements, even though the fitting error $\|b - |F_M^H x|^2\|^2$ is always close to zero. For the new measuring system, CoRK obtains a solution with much higher accuracy, and lower computation time (shown in Fig. 4), compared to the widely used Fienup's algorithm.

The new measuring system, together with the proposed CoRK method or Fienup's algorithm followed by spectral factorization, is also robust to noise, in the sense that the mean squared error (MSE) $\mathbf{E}\|s - \hat{s}\|^2$ can essentially attain the Cramér-Rao bound (CRB). The CRB for the phase-retrieval problem is derived in [46], and is used here to benchmark the performance of our new technique, which is the only method that can guarantee perfect signal reconstruction in the

18

Figure 5: Normalized MSE $\|\boldsymbol{s} - \hat{\boldsymbol{s}}\|^2/\|\boldsymbol{s}\|^2$ vs. normalized CRB using the new measurement system, where we increased the number of measurements on the left, and SNR on the right.

noiseless case. The CRB results in [46] are with respect to the real and imaginary part of the signal being measured, in our case $\boldsymbol{s}_{\min}$. We therefore sum over the diagonals of the pseudo-inverse of the Fisher information matrix except for the 1st and $n + 1$st entries, which corresponds to the real and imaginary parts of δ, and define that as the CRB for $\mathbf{E}\|\boldsymbol{s} - \hat{\boldsymbol{s}}\|^2$. Furthermore, since the CRB is dependent on the true value of $\boldsymbol{s}_{\min}$, we show the results with respect to a fixed signal $\boldsymbol{s}$ in order to keep the CRB curve consistent with how we change one parameter setting of the simulation.

We set $N = 1024$ and generate a fixed signal $\boldsymbol{s} \in \mathbb{C}^N$ from i.i.d. $\mathcal{CN}(0,1)$. Similar to the previous simulation, $\boldsymbol{s}_{\min}$ is constructed according to (13) with $\delta = 3N$, so that it is minimum phase with very high probability. White Gaussian noise with variance σ^2 is added to the squared magnitude of the Fourier transform of $\boldsymbol{s}_{\min}$. The signal-to-noise ratio (SNR) is defined as $10 \log_{10}(\||\boldsymbol{F}_M \boldsymbol{s}_{\min}|^2\|^2/M\sigma^2)$. The performance is plotted in Fig. 4, where we show the normalized error $\|\boldsymbol{s} - \hat{\boldsymbol{s}}\|^2/\|\boldsymbol{s}\|^2$ versus the normalized CRB (original CRB divided by $\|\boldsymbol{s}\|^2$), averaged over 100 Monte-Carlo trials. On the left, we fix SNR= 40dB, and increase the number of measurements M from $2N$ to $16N$. On the right, we fix $M = 8N$, and increase the SNR from 30dB to 60dB. The SNR may seem high here, but notice that most of the signal power is actually concentrated in the artificially added impulse δ, so that the actual noise power is much higher comparing to that of $\boldsymbol{s}$ *per se*. In all cases the MSE obtained from our proposed method is able to attain the CRB sharply, even for as few as $M = 2N$ measurements.

9 Conclusion

We studied the phase retrieval problem with 1D Fourier measurements, a classical estimation problem that has challenged researchers for decades. Our contributions to this challenging problem are as follows:

- *Convexity.* We showed that the 1D Fourier phase retrieval problem can be solved by a convex SDP.

- *Uniqueness.* We proposed a simple measurement technique that adds an impulse to the signal before measuring the Fourier intensities, so that any signal can be uniquely identified.

- *Algorithm.* We developed CoRK – a highly scalable algorithm to solve this problem, which was shown in our simulations to outperform all prior art.

In terms of future research directions, it is interesting to investigate how to incorporate constraints such as sparsity, non-negativity, or bandlimitedness into the auto-correlation parametrization. It is also tempting to consider extending our hidden convexity result and fast algorithm from 1D to other Fourier-based phase retrieval problems, for example the 2D case, which is identifiable, but difficult to solve.

A   Equivalence of PhaseLift and Problem (9)

We show that the exact convex reformulation (9) is in fact equivalent to PhaseLift (4) without the trace regularization if we eliminate the variable r in (9). A similar claim is given in [27] for the more general Toeplitz QCQP, but we show this again here in our context for completeness.

Consider the mth entry in the vector $\Re\{F_M \tilde{I} r\}$, and replace r by its trace parameterization

$$r_k = \mathrm{tr}\left\{T_k X\right\}, \ \ k = 0, 1, ..., N-1.$$

We then have that

$$
\begin{aligned}
\Re\left\{f_m^H \tilde{I} r\right\} &= r_0 + 2\sum_{k=1}^{N-1} \Re\left\{\phi^{km} r_k\right\} \\
&= r_0 + \sum_{k=1}^{N-1} \left(\phi^{km} r_k + \phi^{-km} r_k^*\right) \\
&= \mathrm{tr}\left\{T_0 X\right\} + \sum_{k=1}^{N-1} \left(\phi^{km}\mathrm{tr}\left\{T_k X\right\} + \phi^{-km}\mathrm{tr}\left\{T_k^T X\right\}\right) \\
&= \mathrm{tr}\left\{f_m f_m^H X\right\},
\end{aligned}
\tag{18}
$$

where the last step is a result of the fact that

$$T_0 + \sum_{k=1}^{N-1} \left(\phi^{km} T_k + \phi^{-km} T_k^T\right) = f_m f_m^H.$$

We now eliminate r in (9) as follows:

$$
\begin{aligned}
\left\|b - \Re\left\{F_M \tilde{I} r\right\}\right\|^2 &= \sum_{m=0}^{M-1} \left(b_m - \Re\left\{f_m^H \tilde{I} r\right\}\right)^2 \\
&= \sum_{m=0}^{M-1} \left(b_m - \mathrm{tr}\left\{f_m f_m^H X\right\}\right)^2,
\end{aligned}
\tag{19}
$$

which is exactly the cost function for PhaseLift (4) without the trace regularization.

B Combining PhaseLift with spectral factorization

As we explained in the paper, 1D Fourier phase retrieval can be solved exactly using SDP in the following two steps. First, solve problem (9) (repeated here),

$$\underset{\boldsymbol{r}\in\mathbb{C}^N, \boldsymbol{X}\in\mathbb{H}^N_+}{\text{minimize}} \quad \left\| \boldsymbol{b} - \Re\left\{ \boldsymbol{F}_M \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} \right\|^2 \tag{20}$$
$$\text{subject to} \quad r_k = \text{tr}\left\{ \boldsymbol{T}_k \boldsymbol{X} \right\}, \quad k = 0, 1, ..., N-1.$$

Denote $\boldsymbol{r}_\star$ as the optimal solution of (20), which is the unique solution since the cost of (20) is strongly convex with respect to $\boldsymbol{r}$. Next, perform an SDP-based spectral factorization by computing the solution to

$$\underset{\boldsymbol{X}\in\mathbb{H}^N_+}{\text{maximize}} \quad X_{00} \tag{21}$$
$$\text{subject to} \quad r_{\star k} = \text{tr}\left\{ \boldsymbol{T}_k \boldsymbol{X} \right\}, k = 0, 1, ..., N-1,$$

and let the optimal solution of (21) be $\boldsymbol{X}_\star$. Note that $\boldsymbol{X}_\star$ is the unique solution for (21), and it is rank one [28]. Here we want to show that these two steps can actually be combined into one, as shown in the following proposition.

Proposition 3. *Consider the following SDP*

$$\underset{\boldsymbol{X}\in\mathbb{H}^N_+}{\text{minimize}} \quad \sum_{m=0}^{M-1} \left(b_m - \text{tr}\left\{ \boldsymbol{f}_m \boldsymbol{f}_m^H \boldsymbol{X} \right\} \right)^2 - \lambda X_{00}. \tag{22}$$

There exists some positive λ such that the solution of (22) is guaranteed to be $\boldsymbol{X}_\star$, which is also the solution of (21) with $\boldsymbol{r}_\star$ being the solution of (20).

Proof. Notice that $(\boldsymbol{r}_\star, \boldsymbol{X}_\star)$ is a feasible solution for (20), and in fact $\boldsymbol{r}_\star$ is the unique solution. On the other hand, $\boldsymbol{X}_\star$ is only one solution among a set of solutions for (20), but is the unique solution for problem (21). Now consider the following problem

$$\underset{\boldsymbol{r}\in\mathbb{C}^N, \boldsymbol{X}\in\mathbb{H}^N_+}{\text{minimize}} \quad \left\| \boldsymbol{b} - \Re\left\{ \boldsymbol{F}_M \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} \right\|^2$$
$$\text{subject to} \quad r_k = \text{tr}\left\{ \boldsymbol{T}_k \boldsymbol{X} \right\}, \quad k = 0, 1, ..., N-1, \tag{23}$$
$$X_{00} \geq X_{\star 00}.$$

It is easy to see that $(\boldsymbol{r}_\star, \boldsymbol{X}_\star)$ is its unique solution. From Lagrange duality, the above problem has the same solution as

$$\underset{\boldsymbol{r}\in\mathbb{C}^N, \boldsymbol{X}\in\mathbb{H}^N_+}{\text{minimize}} \quad \left\| \boldsymbol{b} - \Re\left\{ \boldsymbol{F}_M \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} \right\|^2 + \lambda(X_{\star 00} - X_{00})$$
$$\text{subject to} \quad r_k = \text{tr}\left\{ \boldsymbol{T}_k \boldsymbol{X} \right\}, \quad k = 0, 1, ..., N-1, \tag{24}$$

where $\lambda > 0$ is the optimal Lagrangian multiplier with respect to the equality constraint $X_{00} = X_{\star 00}$ in (23).

Finally, by dropping the constant $\lambda X_{\star 00}$ in the cost of (24) and eliminating the variable $\boldsymbol{r}$ as we described in Appendix A, we end up with the formulation (22), or (12), leading to our conclusion that $\boldsymbol{X}_\star$ is the unique solution to (22). $\qquad\square$

Proposition 3 allows to solve 1D Fourier phase retrieval via solving (12), to which all methods provided by TFOCS can still be applied, and by tuning λ we obtain an optimal solution that is also rank one.

C   Derivation of Algorithm (17)

We first review the alternating direction method of multipliers (ADMM), which is the tool used to derive algorithm (17).

Consider the following optimization problem:

$$\underset{\boldsymbol{x}, \boldsymbol{z}}{\text{minimize}} \quad f(\boldsymbol{x}) + g(\boldsymbol{z}),$$

$$\text{subject to} \quad \boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{z} = \boldsymbol{c}.$$

ADMM solves this problem using the following updates

$$\boldsymbol{x} \leftarrow \arg\min_{\boldsymbol{x}} f(\boldsymbol{x}) + \rho \|\boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{z} - \boldsymbol{c} + \boldsymbol{u}\|^2,$$

$$\boldsymbol{z} \leftarrow \arg\min_{\boldsymbol{z}} g(\boldsymbol{z}) + \rho \|\boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{z} - \boldsymbol{c} + \boldsymbol{u}\|^2,$$

$$\boldsymbol{u} \leftarrow \boldsymbol{u} + \boldsymbol{A}\boldsymbol{x} + \boldsymbol{B}\boldsymbol{z} - \boldsymbol{c},$$

where $\rho > 0$ is the step size. ADMM converges to an optimal solution as long as the problem is closed, convex, and proper. The flexible two-block structure of the algorithm often leads to very efficient and/or parallel algorithms that work well in practice.

We now apply ADMM to problem (16), which leads to the following steps:

$$\boldsymbol{r} \leftarrow \arg\min_{\boldsymbol{r}} \left\| \boldsymbol{b} - \Re\left\{ \boldsymbol{F}_M \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} \right\|^2 + \rho \left\| \Re\left\{ \boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} - \boldsymbol{z} + \boldsymbol{u} \right\|^2$$

$$\boldsymbol{z} \leftarrow \arg\min_{\boldsymbol{z}} \mathcal{I}_+(\boldsymbol{z}) + \rho \left\| \Re\left\{ \boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} - \boldsymbol{z} + \boldsymbol{u} \right\|^2$$

$$\boldsymbol{u} \leftarrow \boldsymbol{u} + \Re\left\{ \boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} - \boldsymbol{z}$$

The explicit update for $\boldsymbol{z}$ is very straight forward—it is simply a projection of the point $\Re\{\boldsymbol{F}_L \tilde{\boldsymbol{I}} \boldsymbol{r}\} + \boldsymbol{u}$ onto the non-negative orthant, leading to the update of $\boldsymbol{z}$ as shown in (17).

We next derive the update for $\boldsymbol{r}$ in detail. First note that $\boldsymbol{r}$ is in general complex, so that the derivative needs to be taken with care. To this end we treat $\boldsymbol{r}$ and $\boldsymbol{r}^*$ as independent variables and take derivatives with respect to them separately. This approach is referred to as the *Wirtinger derivative* [47]. Focusing on the first term of the $\boldsymbol{r}$ sub-problem, we can rewrite it as

$$\left\| \boldsymbol{b} - \frac{1}{2} \boldsymbol{F}_M \tilde{\boldsymbol{I}} \boldsymbol{r} - \frac{1}{2} \boldsymbol{F}_M^* \tilde{\boldsymbol{I}} \boldsymbol{r}^* \right\|^2. \tag{25}$$

Taking the derivative with respect to $\boldsymbol{r}$ while treating $\boldsymbol{r}^*$ as an independent variable, we obtain

$$\frac{1}{2} \tilde{\boldsymbol{I}} \boldsymbol{F}_M^H \boldsymbol{F}_M \tilde{\boldsymbol{I}} \boldsymbol{r} + \frac{1}{2} \tilde{\boldsymbol{I}} \boldsymbol{F}_M^H \boldsymbol{F}_M^* \tilde{\boldsymbol{I}} \boldsymbol{r}^* - \tilde{\boldsymbol{I}} \boldsymbol{F}_M^H \boldsymbol{b}. \tag{26}$$

Recall that $\boldsymbol{F}_M$ represents the first N columns of the DFT matrix, which are orthogonal to each other, so that $\boldsymbol{F}_M^H \boldsymbol{F}_M = M\boldsymbol{I}$. On the other hand, the columns of $\boldsymbol{F}_M^*$, except for the first one, come

from the last $N-1$ columns of the same DFT matrix; if we assume $M \geq 2N$, then we have that all columns of $\boldsymbol{F}_M^*$ are orthogonal to columns of $\boldsymbol{F}_M$, except for the first one, which is equal to $\mathbf{1}$. Therefore, $\boldsymbol{F}_M^H \boldsymbol{F}_M^* = M\boldsymbol{E}_{00}$, where $\boldsymbol{E}_{00}$ has only one entry in the upper-left corner that is equal to one, and zeros elsewhere. Since $\tilde{\boldsymbol{I}} = \mathsf{diag}\{[\ 1\ 2\ 2\ ...\ 2\ ]\}$, (26) simplifies to

$$\frac{M}{2}\tilde{\boldsymbol{I}}\tilde{\boldsymbol{I}}\boldsymbol{r} + \frac{1}{2}\tilde{\boldsymbol{I}}\boldsymbol{E}_{00}\tilde{\boldsymbol{I}}\boldsymbol{r}^* - \tilde{\boldsymbol{I}}\boldsymbol{F}_M^H\boldsymbol{b} = M\tilde{\boldsymbol{I}}\boldsymbol{r} - \tilde{\boldsymbol{I}}\boldsymbol{F}_M^H\boldsymbol{b},$$

where we used the fact that r_0 is real. Similar expressions apply to the second term of the $\boldsymbol{r}$ sub-problem. By setting the gradient of the $\boldsymbol{r}$ sub-problem equal to zero, we have that

$$M\tilde{\boldsymbol{I}}\boldsymbol{r} - \tilde{\boldsymbol{I}}\boldsymbol{F}_M^H\boldsymbol{b} + \rho\left(L\tilde{\boldsymbol{I}}\boldsymbol{r} - \tilde{\boldsymbol{I}}\boldsymbol{F}_L^H(\boldsymbol{z} - \boldsymbol{u})\right) = 0,$$

which leads to the update for $\boldsymbol{r}$ given in (17). The gradient with respect to $\boldsymbol{r}^*$ can be shown to be exactly the conjugate of the gradient of $\boldsymbol{r}$. Therefore, equating it to zero gives the same result.

Finally, if we restrict $\boldsymbol{r}$ to be real, then the term (25) can be written as

$$\left\| \boldsymbol{b} - \frac{1}{2}\boldsymbol{F}_M\tilde{\boldsymbol{I}}\boldsymbol{r} - \frac{1}{2}\boldsymbol{F}_M^*\tilde{\boldsymbol{I}}\boldsymbol{r} \right\|^2.$$

Its gradient with respect to $\boldsymbol{r}$ is then

$$\begin{aligned}
&\frac{1}{2}\tilde{\boldsymbol{I}}\boldsymbol{F}_M^H\boldsymbol{F}_M\tilde{\boldsymbol{I}}\boldsymbol{r} + \frac{1}{2}\tilde{\boldsymbol{I}}\boldsymbol{F}_M^H\boldsymbol{F}_M^*\tilde{\boldsymbol{I}}\boldsymbol{r} - \tilde{\boldsymbol{I}}\boldsymbol{F}_M^H\boldsymbol{b} \\
&+ \frac{1}{2}\tilde{\boldsymbol{I}}\boldsymbol{F}_M^T\boldsymbol{F}_M\tilde{\boldsymbol{I}}\boldsymbol{r} + \frac{1}{2}\tilde{\boldsymbol{I}}\boldsymbol{F}_M^T\boldsymbol{F}_M^*\tilde{\boldsymbol{I}}\boldsymbol{r} - \tilde{\boldsymbol{I}}\boldsymbol{F}_M^T\boldsymbol{b} \\
= \ &\frac{M}{2}\tilde{\boldsymbol{I}}\tilde{\boldsymbol{I}}\boldsymbol{r} + \frac{M}{2}\tilde{\boldsymbol{I}}\boldsymbol{E}_{00}\tilde{\boldsymbol{I}}\boldsymbol{r} - \tilde{\boldsymbol{I}}\boldsymbol{F}_M^H\boldsymbol{b} \\
&+ \frac{M}{2}\tilde{\boldsymbol{I}}\boldsymbol{E}_{00}\tilde{\boldsymbol{I}}\boldsymbol{r} + \frac{M}{2}\tilde{\boldsymbol{I}}\tilde{\boldsymbol{I}}\boldsymbol{r} - \tilde{\boldsymbol{I}}\boldsymbol{F}_M^T\boldsymbol{b} \\
= \ &2M\tilde{\boldsymbol{I}}\boldsymbol{r} - 2\Re\left\{\tilde{\boldsymbol{I}}\boldsymbol{F}_M^H\boldsymbol{b}\right\}.
\end{aligned}$$

Therefore, it is easy to see that the update of $\boldsymbol{r}$ that is real is

$$\boldsymbol{r} = \Re\left\{\frac{1}{M + \rho L}\left(\boldsymbol{F}_M^H\boldsymbol{b} + \rho\boldsymbol{F}_L^H(\boldsymbol{z} - \boldsymbol{u})\right)\right\},$$

which is simply a projection of the original complex update onto the real domain.

Acknowledgment

References

[1] K. Huang, Y. C. Eldar, and N. D. Sidiropoulos. On convexity and identifiability in 1-D Fourier phase retrieval. In *IEEE International Conference on Acoustic, Speech, and Signal Processing (ICASSP)*, March 20–25, Shanghai, China, 2016 (to appear).

[2] N. E. Hurt. *Phase Retrieval and Zero Crossings: Mathematical Methods in Image Reconstruction*, volume 52. Springer, 2001.

[3] R. W. Harrison. Phase problem in crystallography. *JOSA A*, 10(5):1046–1055, 1993.

[4] A. Walther. The question of phase retrieval in optics. *Journal of Modern Optics*, 10(1):41–49, 1963.

[5] 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 Processing Magazine*, 32(3):87–109, 2015.

[6] K. Jaganathan, Y. C. Eldar, and B. Hassibi. Phase retrieval: An overview of recent developments. In A. Stern, editor, *Optical Compressive Sensing*, 2016 (to appear).

[7] E. Hofstetter. Construction of time-limited functions with specified autocorrelation functions. *IEEE Transactions on Information Theory*, 10(2):119–126, 1964.

[8] H. H. Bauschke, P. L. Combettes, and D. R. Luke. Phase retrieval, error reduction algorithm, and Fienup variants: a view from convex optimization. *JOSA A*, 19(7):1334–1345, 2002.

[9] Y. M. Lu and M. Vetterli. Sparse spectral factorization: Unicity and reconstruction algorithms. In *IEEE International Conference on Acoustic, Speech, and Signal Processing (ICASSP)*, pages 5976–5979, 2011.

[10] H. Ohlsson and Y. C. Eldar. On conditions for uniqueness in sparse phase retrieval. In *IEEE International Conference on Acoustic, Speech, and Signal Processing (ICASSP)*, pages 1841–1845, 2014.

[11] K. Jaganathan, S. Oymak, and B. Hassibi. Recovery of sparse 1-D signals from the magnitudes of their Fourier transform. *CoRR*, abs/1206.1405, 2012.

[12] Y. C. Eldar and S. Mendelson. Phase retrieval: Stability and recovery guarantees. *Applied and Computational Harmonic Analysis*, 36(3):473–494, 2014.

[13] E. J. Candès, X. Li, and M. Soltanolkotabi. Phase retrieval from coded diffraction patterns. *Applied and Computational Harmonic Analysis*, 39(2):277–299, 2015.

[14] K. Jaganathan, Y. C. Eldar, and B. Hassibi. Recovering signals from the short-time Fourier transform magnitude. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 3277–3281, 2015.

[15] Y. C. Eldar, P. Sidorenko, D. G. Mixon, S. Barel, and O. Cohen. Sparse phase retrieval from short-time fourier measurements. *IEEE Signal Processing Letters*, 22(5):638–642, 2015.

[16] R. W. Gerchberg. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik*, 35:237, 1972.

[17] J. R. Fienup. Phase retrieval algorithms: a comparison. *Applied optics*, 21(15):2758–2769, 1982.

[18] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski. Phase retrieval via matrix completion. *SIAM Journal on Imaging Sciences*, 6(1):199–225, 2013.

[19] Y. Shechtman, Y. C. Eldar, A. Szameit, and M. Segev. Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing. *Optics Express*, 19(16):14807–14822, 2011.

[20] Y. Shechtman, A. Beck, and Y. C. Eldar. GESPAR: efficient phase retrieval of sparse signals. *IEEE Transactions on Signal Processing*, 62(4):928–938, Feb. 2014.

[21] A. Beck and Y. C. Eldar. Sparsity constrained nonlinear optimization: Optimality conditions and algorithms. *SIAM Optimization*, 23(3):1480–1509, Oct. 2013.

[22] D. Gabor. A new microscopic principle. *Nature*, 161(4098):777–778, 1948.

[23] V. Elser. Direct phasing of nanocrystal diffraction. *Acta Crystallographica Section A: Foundations of Crystallography*, 69(6):559–569, 2013.

[24] Z.-Q. Luo, W.-K. Ma, M.-C. So, Y. Ye, and S. Zhang. Semidefinite relaxation of quadratic optimization problems. *IEEE Signal Processing Magazine*, 27(3):20–34, 2010.

[25] I. Waldspurger, A. d'Aspremont, and S. Mallat. Phase recovery, MaxCut and complex semidefinite programming. *Mathematical Programming*, 149(1-2):47–81, 2015.

[26] E. Karipidis, N. D. Sidiropoulos, and Z.-Q. Luo. Far-field multicast beamforming for uniform linear antenna arrays. *IEEE Transactions on Signal Processing*, 55(10):4916–4927, 2007.

[27] A. Konar and N. D. Sidiropoulos. Hidden convexity in QCQP with Toeplitz-Hermitian quadratics. *IEEE Signal Processing Letters*, 22(10):1623–1627, 2015.

[28] B. Dumitrescu. *Positive Trigonometric Polynomials and Signal Processing Applications*. Springer, 2007.

[29] B. Alkire and L. Vandenberghe. Convex optimization problems involving finite autocorrelation sequences. *Mathematical Programming*, 93(3):331–359, 2002.

[30] M. G. Krein and A. A. Nudelman. *The Markov Moment Problem and Extremal Problems*, volume 50 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, Rhode Island, 1977.

[31] S.-P. Wu, S. P. Boyd, and L. Vandenberghe. FIR filter design via semidefinite programming and spectral factorization. In *IEEE Conference on Decision and Control*, volume 1, pages 271–276, 1996.

[32] B. Dumitrescu, I. Tăbuç, and P. Stoica. On the parameterization of positive real sequences and MA parameter estimation. *IEEE Transactions on Signal Processing*, 49(11):2630–2639, 2001.

[33] A. Sayed and T. Kailath. A survey of spectral factorization methods. *Numerical linear algebra with applications*, 8(6-7):467–496, 2001.

[34] S.-P. Wu, S. P. Boyd, and L. Vandenberghe. FIR filter design via spectral factorization and convex optimization. In *Applied and Computational Control, Signals and Circuits*, pages 215–245. 1999.

[35] R. Beinert and G. Plonka. Ambiguities in one-dimensional discrete phase retrieval from fourier magnitudes. *Journal of Fourier Analysis and Applications*, 21(6):1169–1198, 2015.

[36] S. R. Becker, E. J. Candès, and M. C. Grant. Templates for convex cone problems with applications to sparse signal recovery. *Mathematical Programming Computation*, 3(3):165–218, 2011.

[37] B. Recht, M. Fazel, and P. A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.

[38] H. P. Hirst and W. T. Macey. Bounding the roots of polynomials. *College Mathematics Journal*, pages 292–295, 1997.

[39] A. E. Yagle and A. E. Bell. One- and two-dimensional minimum and nonminimum phase retrieval by solving linear systems of equations. *IEEE Transactions on Signal Processing*, 47(11):2978–2989, 1999.

[40] M. Takeda, H. Ina, and S. Kobayashi. Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry. *Journal of the Optical Society of America*, 72(1):156–160, 1982.

[41] J.-C. M. Diels, J. J. Fontaine, I. C. McMichael, and F. Simoni. Control and measurement of ultrashort pulse shapes (in amplitude and phase) with femtosecond accuracy. *Applied Optics*, 24(9):1270–1282, 1985.

[42] S. P. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends®️ in Machine Learning*, 3(1):1–122, 2011.

[43] E. Ghadimi, A. Teixeira, I. Shames, and M. Johansson. Optimal parameter selection for the alternating direction method of multipliers (ADMM): quadratic problems. *IEEE Transactions on Automatic Control*, 60(3):644–658, 2015.

[44] P. Giselsson and S. P. Boyd. Linear convergence and metric selection in douglas-rachford splitting and admm. *IEEE Transactions on Automatic Control*, 2016 (to appear).

[45] J. R. Fienup. Reconstruction of an object from the modulus of its Fourier transform. *Optics Letters*, 3(1):27–29, 1978.

[46] 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 Transactions on Signal Processing*, 2016, submitted.

[47] W. Wirtinger. Zur formalen theorie der funktionen von mehr komplexen veränderlichen. *Mathematische Annalen*, 97(1):357–375, 1927.