ON CONVEXITY AND IDENTIFIABILITY IN 1-D FOURIER PHASE RETRIEVAL

Kejun Huang[*] *Yonina C. Eldar*[†] *Nicholas D. Sidiropoulos*[*]

[*]University of Minnesota, Minneapolis, MN, 55455, USA
[†]Technion – Israel Institute of Technology, Haifa, 32000, Israel

ABSTRACT

This paper considers phase retrieval from the magnitude of 1-D over-sampled Fourier measurements. We first revisit the well-known lack of identifiability in this case, and point out that there always exists a solution that is minimum phase, even though the desired signal is not. Next, we explain how the least-squares formulation of this problem can be optimally solved via PhaseLift followed by spectral factorization, and this solution is always minimum phase. A simple approach is then proposed to circumvent non-identifiability: adding an impulse to an arbitrary complex signal (offset to the Fourier transform) before taking the quadratic measurements, so that a minimum phase signal is constructed and thus can be uniquely estimated. Simulations with synthetic data show the effectiveness of the proposed method.

Index Terms— phase retrieval, over-sampled Fourier measurements, minimum phase, auto-correlation estimation, semi-definite programming.

1. INTRODUCTION

Phase retrieval seeks to find a signal from the magnitudes of linear measurements. This problem arises in various applications, including crystallography, microscopy, and optical imaging, due to the limitations of the detectors used in those applications. Different types of measurement systems have been proposed and used in practice, e.g., over-sampled Fourier measurements, short-time Fourier measurements, random Gaussian, to name just a few (see [1] for a contemporary review). Two of the fundamental questions regarding this problem are: i) Is the signal uniquely determined by the (noiseless) magnitude measurements (up to inherent and usually inconsequential ambiguities like a global phase); and ii) Is there an efficient algorithm that can provably compute an optimal estimate of the signal according to a suitable criterion (like maximum likelihood)?

This paper considers the phase retrieval problem with a specific but widely used measurement system: 1-D over-sampled Fourier measurements. We will answer both the aforementioned questions positively. We first show that the least-squares formulation of this problem can be optimally solved in polynomial time. Due to the fact that all solutions obtained from our algorithm have the minimum phase property, we then propose an approach that first transforms an arbitrary signal into a minimum phase signal by simply adding an impulse, so that identifiability is restored. Compared to the existing methods that provide identifiability with 1-D masked or short-time Fourier measurements [2, 3], the proposed approach is much easier conceptually, requires minimal number of measurements, and can always be solved to global optimality, as we will see.

As a sneak preview and roadmap of our approach, let us summarize the proposed procedure to recover an arbitrary signal $\boldsymbol{s} = [\, s_1 \; s_2 \; ... \; s_{N-1} \,]^T$ from the magnitude of 1-D over-sampled Fourier measurements.

1. Construct $\boldsymbol{s}_{\min}$ by inserting s_0 in front of $\boldsymbol{s}$, i.e., $\boldsymbol{s}_{\min} = [\, s_0 \; s_1 \; ... \; s_{N-1} \,]^T$, such that $|s_0| \geq \|\boldsymbol{s}\|_1$;
2. Take the M-point discrete Fourier transform of $\boldsymbol{s}_{\min}$, where $M \geq 2N$, and measure its squared magnitude $\{b_1, ..., b_M\}$;
3. Solve the PhaseLift problem (4) with $\{b_1, ..., b_M\}$, denote the solution as $\boldsymbol{X}_\star$, and construct an auto-correlation sequence $\boldsymbol{r}$ using (5b);
4. Perform spectral factorization on $\boldsymbol{r}$ to obtain $\boldsymbol{x}_\star$. This $\boldsymbol{x}_\star$ is the optimal solution for the least-squares phase retrieval problem (2);
5. Delete the first element of $\boldsymbol{x}_\star$ to obtain an estimate of $\boldsymbol{s}$, denoted as $\hat{\boldsymbol{s}}$.

Notice that we use $\boldsymbol{s}$ and its variations to denote the signal to be estimated, and $\boldsymbol{x}$ to denote optimization variables.

2. NON-IDENTIFIABILITY AND HIDDEN CONVEXITY

In this section, we first provide insight on signal recovery from the magnitude of 1-D Fourier measurements. First we revisit the well-known fact that an arbitrary $\boldsymbol{s} \in \mathcal{C}^N$ cannot be uniquely identified from $b_m = |\boldsymbol{f}_m^H \boldsymbol{s}|^2$, for $m = 1, ..., M$, where $\boldsymbol{f}_m^H$ is the m-th row of the M-point discrete Fourier transform (DFT) matrix, truncated at length N. Explicitly,

$$\boldsymbol{f}_m^H = \begin{bmatrix} 1 & \phi^{(m-1)} & \phi^{2(m-1)} & \cdots & \phi^{(N-1)(m-1)} \end{bmatrix}, \qquad (1)$$

and $\phi = \exp(\frac{-2\pi\sqrt{-1}}{M})$. The least-squares formulation of phase retrieval, given the squared magnitudes of 1-D Fourier measurements $\{b_1, ..., b_M\}$, is

$$\underset{\boldsymbol{x} \in \mathcal{C}^N}{\text{minimize}} \quad \sum_{m=1}^{M} (b_m - |\boldsymbol{f}_m^H \boldsymbol{x}|^2)^2. \qquad (2)$$

We will explain how this seemingly non-convex problem can be optimally solved via the popular PhaseLift followed by spectral factorization (SF).

[*]Email: (huang663,nikos)@umn.edu. N. Sidiropoulos was supported by NSF CIF-1525194. K. Huang was supported by a UMII dissertation fellowship.

[†]Email: yonina@ee.technion.ac.il. The work of Y. Eldar was funded by the European Unions Horizon 2020 research and innovation programme under grant agreement ERC-BNYQ, by the Israel Science Foundation under Grant no. 335/14, and by ICore: the Israeli Excellence Center Circle of Light.

2.1. Non-identifiability

It is well-known that phase retrieval from 1-D Fourier measurements is not identifiable without any additional prior information [4, 5], and here we briefly review where the non-identifiability comes from. From classical digital signal processing [6], we know that the DFT can be obtained by sampling the z-transform of the finite-length signal on the unit circle $|z|^2 = 1$. Let $S(z)$ be the z-transform of $\boldsymbol{s}$, which is a polynomial of order $N-1$, and can be written in the factored form:

$$S(z) = \sum_{n=1}^{N} s_n z^{-(n-1)} = s_1 \prod_{n=1}^{N-1} (1 - \xi_n z^{-1}), \tag{3}$$

where ξ_n's are the zeros (roots) of the polynomial $S(z)$. The quadratic measurements can be seen as sampled from another z-transform $|S(z)|^2$, which is

$$\begin{aligned} |S(z)|^2 &= S(z)S^*(1/z^*) \\ &= |s_1|^2 \prod_{n=1}^{N-1} (1 - \xi_n z^{-1})(1 - \xi_n^* z). \end{aligned}$$

As we can see, the zeros of $|S(z)|^2$ always come in conjugate reciprocal pairs, and even if we are given $|S(z)|^2$, we cannot determine whether a zero ξ or its conjugate reciprocal $(\xi^*)^{-1}$ is a root of $S(z)$, let alone reconstructing $\boldsymbol{s}$. In other words, for a given signal $\boldsymbol{s}$, we can find the zeros of its z-transform as in (3), take the conjugate reciprocal of some of them, and then take the inverse z-transform to obtain another signal $\boldsymbol{y}$. If we re-scale $\boldsymbol{y}$ to have the same ℓ_2 norm as $\boldsymbol{s}$, then it is easy to verify that

$$|\boldsymbol{f}_m^H \boldsymbol{s}|^2 = |\boldsymbol{f}_m^H \boldsymbol{y}|^2, \; \forall \, m = 1, ..., M,$$

no matter how large M is.

Traditionally, this problem is often seen in design problems where uniqueness is not important. There, it is natural (from the maximal energy dissipation point of view) to pick the zeros to lie within the unit circle, yielding a so-called *minimum phase* signal.

2.2. Hidden convexity

Ignoring the special structure of $\boldsymbol{f}_m$, it is easy to write (2) as a non-convex quadratically constrained quadratic programming (QCQP), for which the prevailing approach is to use semi-definite relaxation [7] to get a lower bound on the optimal value of (2). In the field of phase retrieval, this procedure is known as *PhaseLift* [8, 9]. Specifically, PhaseLift instead tries to solve the following problem (without the trace norm penalty)

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

where $\mathcal{H}_+^N$ denotes the set of Hermitian positive semi-definite matrices of size $N \times N$. Problem (4) can be cast as semi-definite programming (SDP) and solved in polynomial time, and 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 give a rank one solution, especially when the measurements $\{b_m\}$ are noisy. In that case one can only resort to sub-optimal solutions, for example by taking the first principal component of $\boldsymbol{X}_\star$, possibly refined by a traditional method like the Fienup algorithm [10].

What was not known in the phase retrieval community, to the best of our knowledge, is that for the particular 1-D Fourier measurements, recent advances in non-convex QCQPs [11, 12] can be used to show that problem (2) is equivalent to a convex optimization problem. Define the auto-correlation sequence of $\boldsymbol{x}$ as $r_k = \sum_{n=k+1}^{N} x_n x_{n-k}^*$, and stack it into a vector $\boldsymbol{r} \in \mathcal{C}^N$. Then it can be shown [11, 12] that (2) can be equivalently written as

$$\underset{\boldsymbol{r} \in \mathcal{C}^N, \boldsymbol{X} \in \mathcal{H}_+^N}{\text{minimize}} \quad \sum_{m=1}^{M} \left(b_m - \mathsf{Re}\left\{ \boldsymbol{f}_m^H \tilde{\boldsymbol{I}} \boldsymbol{r} \right\} \right)^2 \tag{5a}$$

$$\text{subject to} \quad r_k = \mathsf{Trace}\left\{ \boldsymbol{T}_k \boldsymbol{X} \right\}, \quad \forall \, k = 0, 1, ..., N-1, \tag{5b}$$

where $\tilde{\boldsymbol{I}} = \mathsf{diag}\left\{ [\, 1 \; 2 \; 2 \; ... \; 2 \,] \right\}$, $\boldsymbol{T}_k$ is a matrix with ones on its k-th sub-diagonal and zeros elsewhere, and $\boldsymbol{T}_0 = \boldsymbol{I}$. The cost function is linear least-squares with respect to $\boldsymbol{r}$, and the constraint (5b) is a semi-definite parametrization of $\boldsymbol{r}$ to ensure that it is a valid auto-correlation sequence [13]. In fact, in terms of optimization we can eliminate the variable $\boldsymbol{r}$, as in [11, 12], which then turns out to yield least-squares PhaseLift problem (4). The corresponding $\boldsymbol{r}$ can be obtained by plugging $\boldsymbol{X}_\star$ back into the formula (5b). However, the difference here from PhaseLift for general measurements is that the constraints (5b) guarantee that there always exists a rank one solution, as shown in [13, Appendix A.].

The question that remains is how to find a signal $\boldsymbol{x}$ that generates such an auto-correlation sequence $\boldsymbol{r}$. This is known as *spectral factorization* (SF). We briefly discuss three SF methods, and refer the reader to [14, Appendix] and [15, Appendix B] for a more complete review.

- Root finding. According to the definition of minimum phase, we can directly calculate the roots of the polynomial defined by the sequence $[\, r_{N-1}^* \; r_{N-2}^* \; ... \; r_1^* \; r_0 \; r_1 \; ... \; r_{N-1}]$. Since $\boldsymbol{r}$ is a valid auto-correlation sequence, the roots come in conjugate reciprocal pairs; therefore, we can pick the ones that are inside the unit circle, expand the expression, and then scale it to have ℓ_2 norm equal to $\sqrt{r_0}$. It is conceptually a simple method, but in practice very sensitive to round-off error when N becomes large.
- SDP based. For a valid auto-correlation $\boldsymbol{r}$, it is shown in [15, Chapter 2.6.1] that the following SDP

$$\begin{aligned} \underset{\boldsymbol{X} \in \mathcal{H}_+^N}{\text{maximize}} \quad & \boldsymbol{X}(1,1) \\ \text{subject to} \quad & r_k = \mathsf{Trace}\left\{ \boldsymbol{T}_k \boldsymbol{X} \right\}, \\ & \forall \, k = 0, 1, ..., N-1, \end{aligned} \tag{6}$$

  has a unique rank one solution, and its rank one component generates the given auto-correlation. Algorithms for SDP are numerically stable, although the complexity could be high if we use a general-purpose SDP solver.
- Kolmogorov's method. It can be shown that if we take the logarithm of a minimum phase signal, then the Hilbert transform of the real part equals the imaginary part. The real part can be calculated from $\boldsymbol{r}$, and the Hilbert transform can be well-approximated by two FFTs if we pick the length to be sufficiently large (typically larger than $20N$ is good enough). This method is both efficient and numerically stable, thus usually the method of choice. Implementation details can be found in [15, Appendix B.4].

Notice that all of these methods give us the minimum phase signal.

We briefly summarize the claims of this subsection here. The PhaseLift problem (4) is a semi-definite relaxation of (2), which does

not generally have (or yield, even with trace norm regularization) a rank one solution. For 1-D Fourier measurement vectors, however, there actually always exist rank one solutions for (4), among other solutions with higher rank [11, 12]. In most cases an SDP algorithm will give us a solution with a higher rank. To obtain a rank one solution, we can construct an auto-correlation sequence $\boldsymbol{r}$ according to (5b), and then perform SF on it. The signal obtained from SF is an optimal solution of problem (2).

3. A NEW MEASUREMENT TECHNIQUE

It is now clear that 1) an arbitrary signal cannot be uniquely determined from the magnitude of its over-sampled 1-D Fourier measurements, because there always exists a minimum phase signal that gives the same measurements; and 2) the least-squares estimate can be efficiently calculated by solving PhaseLift followed by SF, which gives us an optimal solution that is minimum phase. If the true signal $\boldsymbol{s}$ is indeed minimum phase, then we can optimally estimate it in polynomial-time. However, minimum phase is usually an impractical assumption to be imposed to the true signal.

We propose to resolve this ambiguity by deliberately making the signal to be minimum phase before taking the quadratic measurements, so that the augmented signal can be uniquely identified, and then the true signal can be recovered easily.

For an arbitrary complex signal $\boldsymbol{s}$, we propose to add s_0 in front of $\boldsymbol{s}$ before taking measurements, where s_0 satisfies that $|s_0| \geq \|\boldsymbol{s}\|_1$. Denote the augmented signal as $\boldsymbol{s}_{\min}$, then $\boldsymbol{s}_{\min}$ is minimum phase, as explained in the following.

Lemma 1. *[16, Theorem 1.] Let ζ be a zero of the polynomial*

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

where $c_0, ..., c_{N-1} \in \mathcal{C}$ and N is a positive integer, then

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

Equipped with Lemma 1, we show the following result.

Proposition 1. *For an arbitrary complex signal $\boldsymbol{s} = [s_1\ s_2 ... s_{N-1}]^T$, the augmented signal $\boldsymbol{s}_{min} = [\ s_0\ s_1\ s_2\ ...\ s_{N-1}\ ]^T$, where $|s_0| \geq \|\boldsymbol{s}\|_1$, is minimum phase.*

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

$$s_0 + s_1 z^{-1} + ... + s_{N-1} z^{-(N-1)},$$

or equivalently the roots of the polynomial

$$z^{N-1} + \frac{s_1}{s_0} z^{N-2} + ... + \frac{s_{N-1}}{s_0},$$

all lie inside the unit circle. Substituting the coefficients of the above polynomial into the inequality in Lemma 1, the right-hand-side equals 1. Therefore, $\boldsymbol{s}_{\min}$ is minimum phase. □

Conceptually, this approach 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 at the 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, we know that the probability that the magnitude of one element exceeding 3σ is almost negligible; therefore, we can simply construct $\boldsymbol{s}_{\min}$ by setting $s_0 = 3\sigma N$, and then $\boldsymbol{s}_{\min}$ is of minimum phase with very high probability.

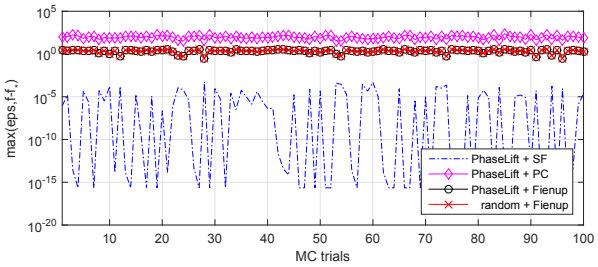

Fig. 1. Optimality gaps $(f - f_\star)$ in each Monte-Carlo trials.

If this new approach is applicable, we also benefit from the fact that M, the number of measurements taken, can be as small as $2N$. If we look at the equivalent reformulation (5), the measurements b_m's are linear with respect to $\boldsymbol{r}$, and 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 semi-definite parametrization (5b). From a unique $\boldsymbol{r}$, a unique minimum phase $\boldsymbol{s}_{\min}$ can be determined from spectral factorization, thus follows the identifiability of phase retrieval with 1-D Fourier measurements.

4. SIMULATIONS

We now present some simulation results to show the effectiveness of our proposed approach. All simulations are done in MATLAB on a Windows desktop. We compare four methods:

- Proposed method, i.e., PhaseLift followed by SF;
- The leading principal component (PC) of the PhaseLift solution;
- Fienup algorithm with 10^5 iterations, initialized by the PC of the PhaseLift solution;
- Fienup algorithm with 10^5 iterations, initialized randomly.

The PhaseLift problem (4) is solved by TFOCS [17], and the SF step is done by the Kolmogorov method with FFT length $32N$.

4.1. Optimization

We first verify that the solution given by spectral factorization indeed attains the lower bound provided by PhaseLift, whereas other methods do not. We want to emphasize that the algorithm always works regardless of whether there is a valid underlying model. To illustrate this, we work with random data. Fixing $N = 128$, we randomly set M as an integer between $2N$ and $8N$, and generate the b_m's from an i.i.d. uniform distribution between $[0, 1]$. The optimality gaps between the fitting error f as in (2) obtained by the aforementioned methods and the theoretical lower bound $f_\star$ given by PhaseLift (4) are shown in Fig. 1, for the 100 Monte-Carlo trials we tested.

As we can see, PhaseLift followed by SF indeed attains the lower bound with negligible gaps, whereas other methods do not. The numerical error of SF using Kolmogorov's method may be as high as 10^{-4}, but still acceptable for practical purposes.

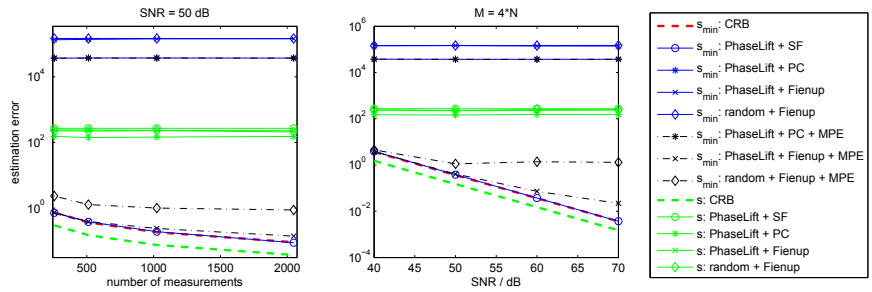

Fig. 2. Averaged estimation error $\|\boldsymbol{s}-\hat{\boldsymbol{s}}\|^2$ while increasing the number of measurements (left) and increasing SNR (right).

4.2. Estimation

Clearly the ultimate goal is not merely to solve problem (2) to optimality, but to recover the signal $\boldsymbol{s}$. In this subsection we verify that our proposed new measuring technique is able to recover $\boldsymbol{s}$ up to global phase ambiguity. For a fixed signal $\boldsymbol{s} \in \mathcal{C}^{N-1}$ with elements randomly generated from i.i.d. $\mathcal{CN}(0,1)$ and $N = 128$, $\boldsymbol{s}_{\min}$ is constructed by inserting $s_0 = 3N$ in front of $\boldsymbol{s}$ so that it is minimum phase with very high probability. We compare the aforementioned algorithms in two measurement setups: direct Fourier measurements on $\boldsymbol{s}$, and using our proposed technique, i.e., measuring $\boldsymbol{s}_{\min}$, and deleting the first element of the solution to obtain $\hat{\boldsymbol{s}}$. White Gaussian noise with variance ς^2 is added to the squared magnitude measurements, fixing the signal-to-noise ratio (SNR) to be the same in both cases:

$$\text{SNR} = 10\log_{10}\frac{\sum_{m=1}^{M}|\boldsymbol{f}_m^H(\boldsymbol{s} \text{ or } \boldsymbol{s}_{\min})|^4}{M\varsigma^2}.$$

Notice that this implies much higher noise power in our proposed setup, because the addition of the impulse artificially inflates the SNR.

After resolving the global phase ambiguity, the estimation error $\|\boldsymbol{s}-\hat{\boldsymbol{s}}\|^2$ from each method is compared with the Cramér-Rao bound (CRB) derived in [18], which lower bounds the averaged squared error that can be attained. The CRB depends on the true value of $\boldsymbol{s}$, and that is why we used the same $\boldsymbol{s}$ throughout the simulations we show here—we want the CRB curve to be consistent with how we change one parameter setting of the simulation. The performance is actually very robust to inputs, as long as we make the signal minimum phase by inserting $s_0 = 3\sigma N$. If $\boldsymbol{s}_{\min}$ is measured, we simply ignore the first element when summing up the CRB for each element, leading to the CRB of the desired signal $\boldsymbol{s}$.

Minimum phase equivalent (MPE) signal: Our method is using the assumption that $\boldsymbol{s}_{\min}$ is minimum phase. It would be unfair for the other methods not to use it as well when trying to recover $\boldsymbol{s}_{\min}$. We therefore compute the minimum phase equivalent (MPE) signal that gives the same fitting error, using the root finding procedure that we discussed in Sec. 2.1, i.e., taking the conjugate reciprocal of the roots that are outside the unit circle. Numerically, the roots of a polynomial are found by calculating the eigenvalues of a companion matrix. The eigenvectors of a companion matrix are Vandermonde vectors generated by the corresponding eigenvalue. So after we take the conjugate reciprocals of some of the roots, assuming all the roots are distinct and non-zero, the companion matrix corresponding to the minimum phase signal can be conversely formed. The equivalent minimum phase signal can then easily be read out from the last row of the new companion matrix[1]. An interesting side observation here is that, although one would normally prefer not lifting the variables, working with matrices of size $N \times N$ seems unavoidable if we want to use the minimum phase prior for identifiability.

The performance is shown in Fig. 2, where each point is an average of 10 Monte-Carlo trials. On the left, we fix SNR$=$ 50dB, and increase the number of measurements M from $2N$ to $16N$. On the right, we fix $M = 4N$, and increase the SNR from 40dB to 70dB. The SNR may seem high here, but notice that 1-D Fourier phase retrieval is one of the most difficult phase retrieval problems, as is also shown by the CRB that if we further decrease SNR, simply guessing $\hat{\boldsymbol{s}} = \boldsymbol{0}$ would attain the CRB. As is shown in Fig. 2, if we follow the proposed procedure, the estimation error is able to attain the CRB, even for as few as $M = 2N$ measurements. None of the other methods are able to attain the theoretical bound in all cases, especially if their solution is not refined by Fienup, or if the minimum phase prior is not used. On the other hand, if we directly measure $\boldsymbol{s}$, none of the algorithms are able to recover the signal correctly, even though in this case the corresponding CRB is smaller (since it has nothing to do with identifiability). Notice that the MPE solution of Fienup initialized with PhaseLift comes second best; however, if we need to solve PhaseLift in the first place, SF with Kolmogorov's method only requires two FFTs with moderately increased length, whereas Fienup requires a large number of iterations, each with FFT complexity, and an additional MPE step.

5. CONCLUSION

In this paper we studied the phase retrieval problem with 1-D oversampled Fourier measurements. As we explained:

1. The least-squares formulation of this problem (2) can be optimally solved, because there always exists a rank one solution for the relaxation method PhaseLift, which can be found by spectral factorization;
2. Without any prior information on the true signal, the solution is not unique. On the other hand, the solution given by spectral factorization is minimum phase.
3. We therefore proposed a simple measurement procedure that first makes any signal minimum phase, then measures the magnitude of its Fourier transform. To make any signal minimum phase, all we need to do is to add an impulse with large enough amplitude to the front of the signal, which is conceptually easy.

As a result, the model is now identifiable, and we have an algorithm that provably solves it to optimality. Simulations on synthetic data confirmed our claims.

[1] Alternatively, one can directly calculate the coefficients of the polynomial from the roots, e.g., using the `poly` function in MATLAB. However, we found that the roundoff error becomes significant as N approaches 64. Using the method we discussed above gives acceptable round-off error, although it does not work if the Vandermonde matrix is singular.

6. REFERENCES

[1] 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*, vol. 32, no. 3, pp. 87–109, 2015.

[2] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval from coded diffraction patterns," *Applied and Computational Harmonic Analysis*, vol. 39, no. 2, pp. 277–299, 2015.

[3] 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)*, 2015, pp. 3277–3281.

[4] L. S. Taylor, "The phase retrieval problem," *IEEE Transactions on Antennas and Propagation*, vol. 29, no. 2, pp. 386–391, 1981.

[5] 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*, vol. 47, no. 11, pp. 2978–2989, 1999.

[6] A. V. Oppenheim and R. W. Schafer, *Discrete-time Signal Processing*, Prentice Hall, 3rd edition, 2010.

[7] Z.-Q. Luo, W.-K. Ma, M.-C. So, Y. Ye, and S. Zhang, "Semidefinite relaxation of quadratic optimization problems," *IEEE Signal Processing Magazine*, vol. 27, no. 3, pp. 20–34, 2010.

[8] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM Journal on Imaging Sciences*, vol. 6, no. 1, pp. 199–225, 2013.

[9] E. J. Candès, T. Strohmer, and V. Voroninski, "PhaseLift: exact and stable signal recovery from magnitude measurements via convex programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1241–1274, 2013.

[10] J. R. Fienup, "Reconstruction of an object from the modulus of its Fourier transform," *Optics Letters*, vol. 3, no. 1, pp. 27–29, 1978.

[11] E. Karipidis, N. D. Sidiropoulos, and Z.-Q. Luo, "Far-field multicast beamforming for uniform linear antenna arrays," *IEEE Transactions on Signal Processing*, vol. 55, no. 10, pp. 4916–4927, 2007.

[12] A. Konar and N. D. Sidiropoulos, "Hidden convexity in QCQP with Toeplitz-Hermitian quadratics," *IEEE Signal Processing Letters*, vol. 22, no. 10, pp. 1623–1627, 2015.

[13] B. Alkire and L. Vandenberghe, "Convex optimization problems involving finite autocorrelation sequences," *Mathematical Programming*, vol. 93, no. 3, pp. 331–359, 2002.

[14] 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*, pp. 215–245. Springer, 1999.

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

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

[17] S. R. Becker, E. J. Candès, and M. C. Grant, "Templates for convex cone problems with applications to sparse signal recovery," *Mathematical Programming Computation*, vol. 3, no. 3, pp. 165–218, 2011.

[18] 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*, 2015, submitted.