Fast Compressive Phase Retrieval

Aditya Viswanathan
Department of Mathematics
Michigan State University
East Lansing, Michigan 48824

Mark Iwen
Department of Mathematics and Department of ECE
Michigan State University
East Lansing, Michigan 48824

***Abstract*—Compressive Phase Retrieval refers to the problem of recovering an unknown sparse signal, upto a global phase constant, given only a small number of phaseless (or magnitude) measurements. This problem occurs in several areas of science – such as optics, astronomy and X-ray crystallography – where the underlying physics of the problem is such that we can only acquire phaseless (or intensity) measurements, and where the underlying signal is sparse (or sparse in an appropriate transform domain). We present here an essentially linear-in-sparsity– time compressive phase retrieval algorithm. We show that it is possible to stably recover k-sparse signals $\mathbf{x} \in \mathbb{C}^n$ from $\mathcal{O}\left(k \log^4 k \cdot \log n\right)$ measurements in only $\mathcal{O}\left(k \log^5 k \cdot \log n\right)$–time. Numerical experiments show that the method is not only fast, but also stable to measurement noise.**

I. Introduction

Let $\mathbf{x} \in \mathbb{C}^n$ be a k-sparse signal, with $k \ll n$. Given the squared magnitude measurements

$$\mathbf{y} = |\mathcal{M}\mathbf{x}|^2 + \mathbf{n}, \tag{1}$$

where $\mathcal{M} \in \mathbb{C}^{m \times n}$ denotes a measurement matrix and $\mathbf{n} \in \mathbb{R}^m$ denotes measurement noise, the compressive phase retrieval problem[1] seeks to recover the unknown signal $\mathbf{x}$ (upto some global phase offset) using only $m \ll n$ phaseless measurements, $\mathbf{y} \in \mathbb{R}^m$. These types of measurements arise in several applications (including optics [1], astronomy [2], quantum mechanics [3] and speech signal processing [4]) on account of the underlying physics. For example, in molecular imaging applications such as X-ray crystallography, we acquire intensity measurements of the diffraction pattern of the underlying specimen [5]. It is possible to show that these measurements correspond to the squared magnitude of the Fourier transform of the underlying specimen or a masked/windowed representation of the specimen. Additionally, it is often the case that the underlying signal $\mathbf{x}$ is sparse or well-approximated by a sparse signal in some transform domain. The recovery of such signals is the subject of this paper. In particular, we are interested in fast (*sub-linear–time*) and efficient compressive phase retrieval methods which use a near-optimal number of phaseless measurements to stably recovery $\mathbf{x}$.

A. Prior Work

Given the numerous practical applications, this problem has attracted the attention of researchers and practitioners across diverse scientific disciplines. Consequently, several computational methods have been proposed for solving this problem. A popular approach involves the extension of general non-sparse phase retrieval methods by incorporating sparsity constraints. For example, imposing additional sparsity constraints in the classical Fienup alternating projection algorithm [6] has been explored in [7], while sparsity enforcing variants of the *PhaseLift* semidefinite relaxation formulation [8] were evaluated in [9]. Problem formulations designed specifically for the compressive phase retrieval problem exist too, such as [10] which is based on a fast local greedy support search procedure, approximate message passing methods such as [11], and combinatorial support identification and estimation procedures such as [12]. In addition, there are also two-stage algorithmic formulations such as [13] and [14]. The results in this paper are motivated, for example, by the construction in [13] which first solves a (non-sparse) phase retrieval problem to recover an intermediate compressed signal, followed by the application of a compressed sensing recovery method such as basis pursuit.

To the best of our knowledge, these and all existing algorithms are super-linear–time in the signal dimension n, which presents a significant computational challenge for very large problems. In this paper, we present theoretical and numerical results showing that it is possible to accurately recover the signal $\mathbf{x}$ in only $\mathcal{O}(k \log^5 k \cdot \log n)$–time using a near-optimal (upto log factors) $\mathcal{O}(k \log^4 k \cdot \log n)$ number of measurements.

[1]It is typical to assume that we are allowed to design both the recovery algorithm $\mathcal{A}_{\mathcal{M}} : \mathbb{R}^m \to \mathbb{C}^n$, as well as the measurement matrix, $\mathcal{M}$.

B. Main Result

Theorem 1. *(**Fast Compressive Phase Retrieval**) There exists a deterministic algorithm $\mathcal{A}_{\mathcal{M}} : \mathbb{R}^m \to \mathbb{C}^n$ for which the following holds: Let $\epsilon \in (0, 1]$, $\mathbf{x} \in \mathbb{C}^n$ with n sufficiently large, and $k \in \{1, 2, \ldots, n\} \subset \mathbb{N}$. Then, one can select a random measurement matrix $\mathcal{M} \in \mathbb{C}^{m\times n}$ such that*

$$\min_{\theta\in[0,2\pi)} \left\|\mathbb{e}^{\mathbb{i}\theta}\mathbf{x} - \mathcal{A}_{\mathcal{M}}\left(|\mathcal{M}\mathbf{x}|^2\right)\right\|_2 \leq \left\|\mathbf{x} - \mathbf{x}_k^{\mathrm{opt}}\right\|_2 + \frac{22\epsilon \left\|\mathbf{x} - \mathbf{x}_{(k/\epsilon)}^{\mathrm{opt}}\right\|_1}{\sqrt{k}}$$

is true with probability at least $1 - \frac{1}{C\cdot\log^2(n)\cdot\log^3(\log n)}$.[2] Here, $\mathbf{x}_k^{\mathrm{opt}}$ denotes the best k-term approximation to $\mathbf{x}$ and m can be chosen to be $\mathcal{O}\big(\frac{k}{\epsilon}\cdot\log^3(\frac{k}{\epsilon})\cdot\log^3\left(\log\frac{k}{\epsilon}\right)\cdot\log n\big)$. Furthermore, the algorithm will run in $\mathcal{O}\left(\frac{k}{\epsilon}\cdot\log^4(\frac{k}{\epsilon})\cdot\log^3\left(\log\frac{k}{\epsilon}\right)\cdot\log n\right)$–time.[3]

This result shows that the runtime and sampling complexities of the proposed algorithm are optimal (upto log factors). Moreover, empirical results show that this method is not only fast, but also stable to measurement noise. The proposed algorithm employs a two-stage construction by incorporating the measurement matrices and recovery methods from: (i) a fast (essentially linear-time) non-sparse phase retrieval algorithm, and (ii) a sub-linear time compressive sensing recovery method.

The rest of the paper is organized as follows: §II describes the fast phase retrieval method, while §III briefly summarizes relevant sub-linear time compressive sensing results and methods. The proposed fast compressive phase retrieval framework is then described in §IV. Numerical results demonstrating the efficiency and robustness are presented in §V, while some concluding comments are offered in §VI.

II. Fast (Non-Sparse) Phase Retrieval

We summarize below a recently introduced fast (*essentially linear-time*) phase retrieval method employing *local correlation*-based measurement matrices. For the sake of brevity, we provide here an illustrative example highlighting the measurement construction as well as salient features of the recovery algorithm, while referring the interested reader to [15] for further details. We start by considering noiseless squared magnitude measurements of the form

$$\mathbf{y} = |\mathcal{P}\mathbf{z}|^2, \tag{2}$$

where $\mathbf{y} \in \mathbb{R}^{12}$, $\mathcal{P} \in \mathbb{C}^{12\times 4}$ denotes a measurement matrix, and $\mathbf{z} \in \mathbb{C}^4$ is the unknown signal we seek to recover. Furthermore, let $\mathcal{P}$ be constructed as follows[4]:

$$\mathcal{P} = \begin{pmatrix}\mathcal{P}_1\\ \mathcal{P}_2\\ \mathcal{P}_3\end{pmatrix}, \quad \mathcal{P}_i \in \mathbb{C}^{4\times 4},\, i \in \{1,2,3\}, \quad \text{where}$$

$$\mathcal{P}_i = \begin{pmatrix} (\mathbf{p}_i)_1^* & (\mathbf{p}_i)_2^* & 0 & 0\\ 0 & (\mathbf{p}_i)_1^* & (\mathbf{p}_i)_2^* & 0\\ 0 & 0 & (\mathbf{p}_i)_1^* & (\mathbf{p}_i)_2^*\\ (\mathbf{p}_i)_2^* & 0 & 0 & (\mathbf{p}_i)_1^* \end{pmatrix}.$$

This corresponds to (squared magnitude) *correlation* measurements of the unknown signal $\mathbf{z} \in \mathbb{C}^4$ with three *local* masks $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3 \in \mathbb{C}^4$, where $(p_i)_\ell = 0$ for $\ell > 2, i \in \{1, 2, 3\}$. Writing out the correlation sum explicitly and setting $\delta = 2$, we obtain

$$\begin{aligned}(y_i)_\ell &= \left|\sum_{k=1}^{\delta}(\mathbf{p_i})_k^*\cdot z_{\ell+k-1}\right|^2\\ &= \sum_{j,k=1}^{\delta}(\mathbf{p_i})_j\,(\mathbf{p_i})_k^*\,z_{\ell+j-1}^*\,z_{\ell+k-1}\\ &:= \sum_{j,k=1}^{\delta}(\mathbf{p_i})_{j,k}\,z_{\ell+j-1}^*\,z_{\ell+k-1},\end{aligned}$$

where we have used the notation $(\mathbf{p_i})_{j,k} := (\mathbf{p_i})_j(\mathbf{p_i})_k^*$. The resulting *linear* system of equations for the (scaled) phase differences $\{z_i^* z_j\}$ may be written as

$$\widetilde{\mathbf{y}} = \mathcal{P}'\mathbf{b}, \tag{3}$$

where $\widetilde{\mathbf{y}}$ denotes the interleaved vector of measurements

$$[(y_1)_1(y_2)_1(y_3)_1\,(y_1)_2(y_2)_2(y_3)_2\,\ldots\,(y_1)_4(y_2)_4(y_3)_4]^T,$$

$\mathbf{b}$ denotes the vector of scaled phase differences

$$\left[|z_1|^2\;z_1^*z_2\;z_2^*z_1\;|z_2|^2\;z_2^*z_3\;\ldots z_4^*z_3\;|z_4|^2\;z_4^*z_1\;z_1^*z_4\right]^T,$$

and $\mathcal{P}'$ is the block circulant matrix

$$\mathcal{P}' = \begin{bmatrix}\mathcal{P}_1' & \mathcal{P}_2' & 0 & 0\\ 0 & \mathcal{P}_1' & \mathcal{P}_2' & 0\\ 0 & 0 & \mathcal{P}_1' & \mathcal{P}_2'\\ \mathcal{P}_2' & 0 & 0 & \mathcal{P}_1'\end{bmatrix}, \qquad \text{with}$$

$$\mathcal{P}_1' = \begin{bmatrix}(\mathbf{p}_1)_{1,1} & (\mathbf{p}_1)_{1,2} & (\mathbf{p}_1)_{2,1}\\ (\mathbf{p}_2)_{1,1} & (\mathbf{p}_2)_{1,2} & (\mathbf{p}_2)_{2,1}\\ (\mathbf{p}_3)_{1,1} & (\mathbf{p}_3)_{1,2} & (\mathbf{p}_3)_{2,1}\end{bmatrix}, \; \mathcal{P}_2' = \begin{bmatrix}(\mathbf{p}_1)_{2,2} & 0 & 0\\ (\mathbf{p}_2)_{2,2} & 0 & 0\\ (\mathbf{p}_3)_{2,2} & 0 & 0\end{bmatrix}.$$

This block-circulant structure allows for efficient inversion of $\mathcal{P}'$ using FFTs, as well as analysis of its condition

[2]Here $C \in \mathbb{R}^+$ is a fixed absolute constant.

[3]For the sake of simplicity, we assume $k = \Omega(\log n)$ when stating the measurement and runtime bounds above.

[4]The notation $(\mathbf{p}_i)_\ell$ denotes the ℓ-th entry of the i-th mask and $(\mathbf{p}_i)_\ell^*$ denotes the complex conjugate of $(\mathbf{p}_i)_\ell$.

number. For example, it is shown in [15] that choosing the mask entries to be

$$(\mathbf{p}_i)_k = \begin{cases} \frac{\mathbb{e}^{-k/a}}{\sqrt[4]{2\delta-1}} \cdot \mathbb{e}^{\frac{2\pi \mathbb{i}\cdot(k-1)\cdot(i-1)}{2\delta-1}} & \text{if } k \leq \delta \\ 0 & \text{if } k > \delta \end{cases} \quad (4)$$

guarantees that the condition number of $\mathcal{P}'$ grows no worse than $\mathcal{O}(\delta^2)$. Since δ is typically small, this ensures well conditioned measurements, irrespective of the problem dimension.

Note that by solving the linear system (3), the magnitude of $\mathbf{z}$ is automatically recovered. In particular,

$$|z_1|^2 = b_1, \quad |z_2|^2 = b_4, \quad |z_3|^2 = b_7, \quad |z_4|^2 = b_{10}.$$

Additionally, by normalizing (to unit magnitude) the entries of $\mathbf{b}$, one also recovers phase difference estimates $\phi_{i,j} := \arg(z_j) - \arg(z_i)$, for $i, j \in \{1, 2, 3, 4\}$ and $|i - j \mod 4| = 1$. For example, $\phi_{1,2} = \arg(z_2) - \arg(z_1) = b_2/|b_2|$. It is now possible to recover $\arg(\mathbf{z})$ by using a greedy procedure. This is commonly referred to as the *angular synchronization* problem [16].

Assume, without loss of generality, that $|z_1| \geq |z_i|$, $i \in \{2, 3, 4\}$. Start by setting $\arg(z_1) = 0$.[5] It is now possible to set the phase of z_2 and z_4 using the estimated phase differences $\phi_{1,2}$ and $\phi_{1,4}$ respectively; i.e.,

$$\arg(z_2) = \arg(z_1) + \phi_{1,2} = \phi_{1,2},$$
$$\arg(z_4) = \arg(z_1) + \phi_{1,4} = \phi_{1,4}.$$

Similarly, one can set $\arg(z_3) = \arg(z_2) + \phi_{2,3}$, thereby recovering all of the entries' unknown phases. Note that the computational cost of this procedure is essentially linear in the dimension of $\mathbf{z}$.

The above discussion assumes that $\mathbf{z}$ is "flat"; i.e., it does not contain a long string of zeros or entries of very small magnitude. In this case, the network of phase differences is broken, thereby causing the angular synchronization method to fail. Nevertheless, for recovering such vectors (and, in general, arbitrary vectors $\mathbf{z}$), the measurement matrix $\mathcal{P}$ may be modified by multiplication with a fast (FFT-time) Johnson-Lindentrauss transform matrix [15] to "flatten" $\mathbf{z}$ and ensure it does not contain a long string of zeros or small entries. While we refer the reader to [15] for more details, we conclude with a noiseless recovery guarantee.

Theorem 2. *(***Fast Phase Retrieval***) Let $\mathbf{z} \in \mathbb{C}^d$ with d sufficiently large. Then, one can select a random measurement matrix $\mathcal{P} \in \mathbb{C}^{m \times d}$ such that the following holds with probability at least $1 - \frac{1}{c \cdot \log^2(d) \cdot \log^3(\log d)}$: The above fast phase retrieval method will recover an $\tilde{\mathbf{z}} \in \mathbb{C}^d$ with*

$$\min_{\theta \in [0, 2\pi)} \left\| \mathbf{z} - \mathbb{e}^{\mathbb{i}\theta} \tilde{\mathbf{z}} \right\|_2 = 0$$

when given the noiseless magnitude measurements $|\mathcal{P}\mathbf{x}|^2 \in \mathbb{R}^m$. Here m can be chosen to be $\mathcal{O}(d \cdot \log^2(d) \cdot \log^3(\log d))$. Furthermore, the algorithm will run in $\mathcal{O}(d \cdot \log^3(d) \cdot \log^3(\log d))$-time in that case.

III. Sub-Linear Time Compressive Sensing

In the last few years, several low-complexity recovery algorithms for recovering a k-sparse vector $\mathbf{x}$ from $d \ll n$ compressed measurements $\mathcal{C}\mathbf{x} \in \mathbb{C}^d$ have been proposed (see, for example [17] and [18]). For the results in this paper, we use the measurement constructions and recovery methods from [19]. Here, the compressed sensing matrix $\mathcal{C} \in \mathbb{C}^{d \times n}$ is a random sparse binary matrix obtained by randomly sub-sampling the rows of another suitably well-chosen incoherent matrix (for example, the adjacency matrix of certain unbalanced expander graphs). It can be shown that the resulting matrices satisfy certain strong combinatorial properties which permit the use of low complexity (sub-linear in the problem size) compressed sensing recovery algorithms.

The recovery algorithm then proceeds in two phases:

1) Identify the k largest magnitude entries of $\mathbf{x}$ using standard bit-testing techniques.
2) Estimate these k largest entries using median estimates and techniques from computer science streaming literature.

The sampling and runtime complexities of this method are both $\mathcal{O}(k \cdot \log k \cdot \log n)$. While we refer the interested reader to [19] for further details, we list below the main result from [19] of relevance to our discussion.

Theorem 3. *(***Sub-Linear Time Compressive Recovery***) Let $\epsilon \in (0, 1]$, $\sigma \in [2/3, 1)$, $\mathbf{x} \in \mathbb{C}^n$, and $k \in \{1, 2, \ldots, n\}$. With probability at least σ the deterministic compressive sensing algorithm from [19] will output a vector $\widetilde{\mathbf{x}} \in \mathbb{C}^n$ satisfying*

$$\|\mathbf{x} - \widetilde{\mathbf{x}}\|_2 \;\leq\; \left\|\mathbf{x} - \mathbf{x}_k^{\rm opt}\right\|_2 + \frac{22\epsilon \left\|\mathbf{x} - \mathbf{x}_{(k/\epsilon)}^{\rm opt}\right\|_1}{\sqrt{k}} \quad (5)$$

when executed with random linear input measurements $\mathcal{C}\mathbf{x} \in \mathbb{C}^d$. Here $d = \mathcal{O}\left(\frac{k}{\epsilon} \cdot \log\left(\frac{k/\epsilon}{1-\sigma}\right) \log n\right)$ suffices. The required runtime of the algorithm is $\mathcal{O}\left(\frac{k}{\epsilon} \cdot \log\left(\frac{k/\epsilon}{1-\sigma}\right) \log\left(\frac{n}{1-\sigma}\right)\right)$ in this case.[6]

[5]Recall that we can only recover $\mathbf{z}$ up to an unknown global phase factor which, in this case, will be the true phase of z_1.

[6]For the sake of simplicity, we assume $k = \Omega(\log n)$ when stating the measurement and runtime bounds above.

IV. Fast Compressive Phase Retrieval

We now present a simple two-stage formulation for our fast compressive phase retrieval algorithm. Let $\mathcal{P} \in \mathbb{C}^{m\times d}$ denote a phase retrieval matrix associated with the phase retrieval method Δ_P, and let $\mathcal{C} \in \mathbb{C}^{d\times n}$ denote a compressive sensing matrix associated with the sub–linear time compressive sensing algorithm Δ_C. Let the measurement matrix $\mathcal{M}$ for the compressive phase retrieval problem (1) be constructed as $\mathcal{M} = \mathcal{PC}$. Then, we can show that $\Delta_C \circ \Delta_P : \mathbb{R}^m \to \mathbb{C}^n$ recovers the unknown signal $\mathbf{x}$ upto a global phase factor stably and accurately. For results presented in this paper, we choose Δ_P to be the fast (FFT–time) phase retrieval method described in §II. For Δ_C, we use the sub–linear time recovery methods proposed in [19], and summarized in §III.

The recovery algorithm proceeds in the following two stages:

1) First, apply the fast phase retrieval method, $\Delta_P : \mathbb{R}^m \to \mathbb{C}^d$, to the phaseless measurements $\mathbf{y}$ and recover an intermediate compressed signal $\mathbf{z} \in \mathbb{C}^d$. From §II and Theorem 2, we known that $m = \mathcal{O}(d \operatorname{polylog} d)$ phaseless measurements suffice for accurate recovery of the intermediate compressed signal $\mathbf{z}$. Further, we also know that it is possible to design Δ_P to run in $\mathcal{O}(d \operatorname{polylog} d)$–time. Additionally, from §III and Theorem 3, we known that the dimension of the intermediate compressed signal is $d = \mathcal{O}\left(k \log k \cdot \log n\right)$.
2) Next, we use a sub–linear time compressive sensing algorithm, $\Delta_C : \mathbb{C}^d \to \mathbb{C}^n$, to recover the unknown signal $\mathbf{x}$ (upto a global phase factor). From §III and Theorem 3, we know that $\mathbf{x}$ can be recovered in $\mathcal{O}\left(k \log k \cdot \log n\right)$–time.

We also note that the proof of the main result, Theorem 1, follows directly from Theorem 2 and Theorem 3.

V. Representative Numerical Simulations

We now present some representative numerical results demonstrating the efficiency of the proposed method. In each case, the test signals were generated to have i.i.d complex Gaussian non-zero entries, with non-zero index locations chosen by k-permutations. Simulations were performed on a laptop computer with an Intel® Core™i3-3120M processor, 6GB RAM and Matlab R2015b. Open source Matlab code[7] used to generate these numerical results can be found at [22].

[7]This code uses a fast (sub-linear time) compressive recovery method detailed in [20] and implemented in [21].

Fig. 1 plots the computational time (in seconds, averaged over 100 trials) taken to solve the compressive phase retrieval problem as a function of the signal sparsity k for different problem sizes. The recovery is deemed to be successful if $\min_{\theta\in[0,2\pi)} \|e^{\mathrm{i}\theta}\mathbf{x} - \widehat{\mathbf{x}}\|_2 \leq 10^{-10}$, where $\widehat{\mathbf{x}}$ denotes the recovered signal. Observe that the

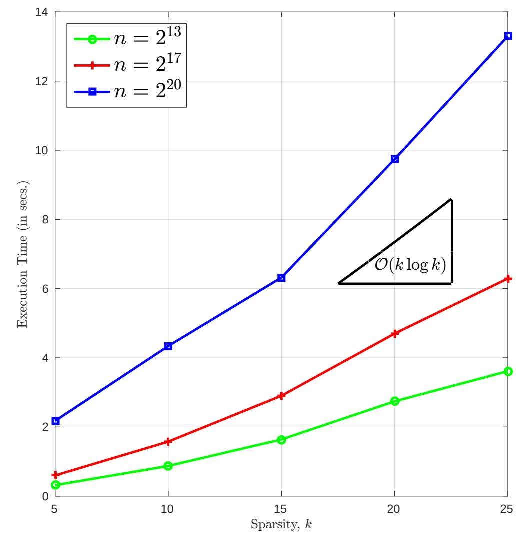

Fig. 1. Computational Efficiency of the Proposed Compressive Phase Retrieval Algorithm

overall execution time is sub-linear in the problem size n and (poly) log-linear in the sparsity k. For example, a 20-sparse million-length complex vector can be recovered in about 10 seconds. For completeness, the (average, over 100 trials) number of measurements acquired and the intermediate signal dimension for the simulation run with signal size $n = 2^{20}$ is tabulated in Table I. We remark that the sub-linear compressive recovery software implementation [21] used to generate these results is not optimized; we expect improved performance (and fewer total measurements) when using more refined and optimized software implementations.

Finally, we present empirical evidence of the robustness of the proposed method to additive measurement noise. Fig. 2 plots the reconstruction error (in dB, averaged over 100 trials) as a function of the added noise level (in dB) for recovering a $n = 2^{20}$-length signal with sparsities $k = 5$ and $k = 25$. Additive i.i.d. Gaussian noise at the prescribed SNRs were added to the test signals prior to reconstruction. We observe that the recovery algorithm displays graceful degradation with the noise level, although it is possible to improve this

k	5	10	15	20	25
m (%)	$7,371$ (0.70)	$20,475$ (1.95)	$38,304$ (3.65)	$53,550$ (5.11)	$75,411$ (7.19)
d	$2,457$	$6,825$	$12,768$	$17,850$	$25,137$

TABLE I
NUMBER OF MEASUREMENTS REQUIRED FOR COMPRESSIVE PHASE RETRIEVAL (SIGNAL LENGTH, $n = 2^{20} = 1,048,576$). HERE, k IS THE SIGNAL SPARSITY, m IS THE NUMBER OF MEASUREMENTS ACQUIRED, $\%$ DENOTES THE NUMBER OF MEASUREMENTS AS A PERCENTAGE OF THE SIGNAL SIZE n, AND d IS THE INTERMEDIATE COMPRESSED SIGNAL DIMENSION.

performance through use of more robust phase retrieval and compressive recovery methods.

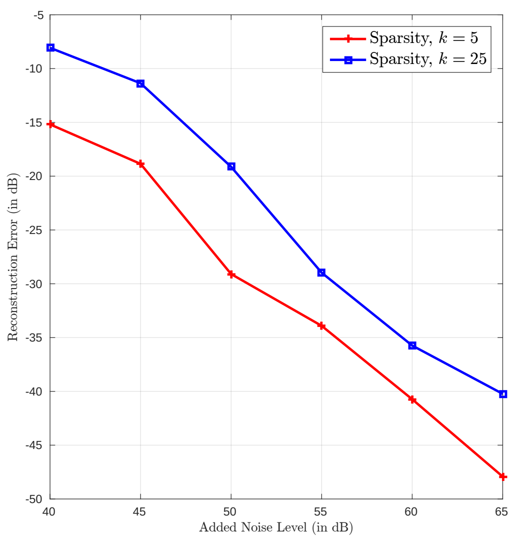

Fig. 2. Robustness of the Proposed Compressive Phase Retrieval Algorithm to Additive Measurement Noise

VI. CONCLUDING REMARKS

We have presented an essentially linear-in-sparsity–time phase retrieval algorithm which is capable of recovering k-sparse signals $\mathbf{x} \in \mathbb{C}^n$ from $\mathcal{O}(k \log^4 k \cdot \log n)$ measurements in only $\mathcal{O}(k \log^5 k \cdot \log n)$–time. Representative numerical results demonstrate the computational efficiency of the method. Future research directions include obtaining robust recovery guarantees for the algorithm in the presence of measurement noise as well as a comprehensive comparison against other compressive phase retrieval methods.

ACKNOWLEDGMENT

M.A. Iwen was supported in part by NSF DMS-1416752 and NSA H98230-13-1-0275.

REFERENCES

[1] A. Walther, “The question of phase retrieval in optics,” *J. Modern Optics*, vol. 10, no. 1, pp. 41–49, 1963.
[2] C. Fienup and J. Dainty, “Phase retrieval and image reconstruction for astronomy,” *Image Recovery: Theory and Applicat.*, pp. 231–275, 1987.
[3] J. Corbett, “The Pauli problem, state reconstruction and quantum-real numbers,” *Rep. Math. Physics*, vol. 57, no. 1, pp. 53–68, 2006.
[4] R. Balan, P. Casazza, and D. Edidin, “On signal reconstruction without phase,” *Appl. Comput. Harmonic Anal.*, vol. 20, no. 3, pp. 345–356, 2006.
[5] R. W. Harrison, “Phase problem in crystallography,” *J. Optical Soc. Amer.*, vol. 10, no. 5, pp. 1046–1055, 1993.
[6] J. R. Fienup, “Reconstruction of an object from the modulus of its Fourier transform,” *Optics Lett.*, vol. 3, pp. 27–29, 1978.
[7] S. Mukherjee and C. S. Seelamantula, “Fienup algorithm with sparsity constraints: Application to frequency-domain optical-coherence tomography,” *IEEE Trans. Signal Process.*, vol. 62, no. 18, pp. 4659–4672, 2014.
[8] E. J. Candes, T. Strohmer, and V. Voroninski, “Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming,” *Commun. Pure Appl. Math.*, vol. 66, no. 8, pp. 1241–1274, 2013.
[9] H. Ohlsson, A. Yang, R. Dong, and S. Sastry, “CPRL: An Extension of Compressive Sensing to the Phase Retrieval Problem,” in *Proc. 26th Conf. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1376–1384.
[10] Y. Shechtman, A. Beck, and Y. C. Eldar, “GESPAR: Efficient Phase Retrieval of Sparse Signals,” *IEEE Trans. Signal Process.*, vol. 62, no. 4, pp. 928–938, 2014.
[11] P. Schniter and S. Rangan, “Compressive Phase Retrieval via Generalized Approximate Message Passing,” in *Proc. Allerton Conf. Commun. Control and Comput.*, 2012.
[12] K. Jaganathan, S. Oymak, and B. Hassibi, “Sparse phase retrieval: Convex algorithms and limitations,” in *Inf. Theory Proc. (ISIT), 2013 IEEE Int. Symp. on*, 2013, pp. 1022–1026.
[13] M. Iwen, A. Viswanathan, and Y. Wang, “Robust Sparse Phase Retrieval Made Easy,” *Appl. Comput. Harmonic Anal.*, 2015, in Press.
[14] S. Bahmani and J. Romberg, “Efficient compressive phase retrieval with constrained sensing vectors,” 2015, preprint, arXiv:1507.08254.
[15] M. Iwen, A. Viswanathan, and Y. Wang, “Fast Phase Retrieval For High-Dimensions,” 2015, preprint, arXiv:1501.02377.
[16] A. Singer, “Angular synchronization by eigenvectors and semidefinite programming,” *Appl. Comput. Harmonic Anal.*, vol. 30, pp. 20–36, 2011.
[17] G. Cormode and S. Muthukrishnan, “Combinatorial algorithms for compressed sensing,” in *Structural Inf. and Commun. Complexity*. Springer, 2006, pp. 280–294.
[18] A. C. Gilbert, Y. Li, E. Porat, and M. J. Strauss, “Approximate sparse recovery: optimizing time and measurements,” in *Proc. 42^{nd} ACM Symp. Theory Comput.*, 2010, pp. 475–484.
[19] M. Iwen, “Compressed sensing with sparse binary matrices: Instance optimal error guarantees in near-optimal time,” *J. Complexity*, vol. 30, no. 1, pp. 1–15, 2014.
[20] A. C. Gilbert, M. J. Strauss, J. A. Tropp, and R. Vershynin, “Sublinear approximation of compressible signals,” in *Proc. SPIE Intell. Integrated Microsystems (IIM)*, 2006, pp. 623 206.01–09.
[21] J. Tropp, “Chaining Pursuit: a reconstruction streaming algorithm,” http://www-personal.umich.edu/~romanv/software/software.html, 2005.
[22] M. Iwen, Y. Wang, and A. Viswanathan, “SparsePR: Matlab software for sparse phase retrieval, version 1.0,” https://bitbucket.org/charms/sparsepr, 2014.