Fast Angular Synchronization for Phase Retrieval via Incomplete Information

Aditya Viswanathan[a] and Mark Iwen[b]

[a]Department of Mathematics, Michigan State University; [b] Department of Mathematics & Department of Electrical and Computer Engineering, Michigan State University.

ABSTRACT

We consider the problem of recovering the phase of an unknown vector, $\mathbf{x} \in \mathbb{C}^d$, given (normalized) phase difference measurements of the form $x_j x_k^* / |x_j x_k^*|$, $j, k \in \{1, \ldots, d\}$, and where x_j^* denotes the complex conjugate of x_j. This problem is sometimes referred to as the *angular synchronization* problem. This paper analyzes a linear-time-in-d eigenvector-based angular synchronization algorithm and studies its theoretical and numerical performance when applied to a particular class of highly incomplete and possibly noisy phase difference measurements. Theoretical results are provided for perfect (noiseless) measurements, while numerical simulations demonstrate the robustness of the method to measurement noise. Finally, we show that this angular synchronization problem and the specific form of incomplete phase difference measurements considered arise in the phase retrieval problem – where we recover an unknown complex vector from phaseless (or magnitude) measurements.

Keywords: angular synchronization, eigen analysis, phase retrieval.

1. INTRODUCTION

We are interested in recovering the phase of a vector, $\mathbf{x} \in \mathbb{C}^d$, given measurements of the form

$$(X')_{j,k} = x_j x_k^*, \quad |j - k \bmod d| < \delta, \tag{1}$$

where $j, k \in [d] := \{1, 2, \ldots, d\}$ and $\delta \in \mathbb{Z}^+$. We note that these measurements describe the diagonal entries of the rank one matrix $\mathbf{x}\mathbf{x}^* \in \mathbb{C}^{d \times d}$; i.e.,

$$(X')_{j,k} = \begin{cases} (\mathbf{x}\mathbf{x}^*)_{j,k} & \text{if } |j - k \bmod d| < \delta \\ 0 & \text{otherwise.} \end{cases} \quad . \tag{2}$$

Normalizing each entry of $X' \in \mathbb{C}^{d \times d}$, we obtain the matrix $A \in \mathbb{C}^{d \times d}$ with entries

$$(A)_{j,k} = \begin{cases} \mathbb{e}^{\mathbb{i}(\phi_j - \phi_k)} & \text{if } |j - k \bmod d| < \delta \\ 0 & \text{otherwise,} \end{cases} \quad , \tag{3}$$

where $x_j = C_j \mathbb{e}^{\mathbb{i}\phi_j}$ for $j \in [d]$. Our objective is to use $A \in \mathbb{C}^{d \times d}$ in order to recover $\phi_1, \ldots, \phi_d \in [0, 2\pi]$ or equivalently, the vector $\tilde{\mathbf{x}} \in \mathbb{C}^d$ with*

$$(\tilde{x})_j := \mathbb{e}^{\mathbb{i}\phi_j} \; = \; x_j / |x_j|. \tag{4}$$

This is an *angular synchronization* problem since we recover d individual phase angles (modulo a global factor of 2π) from the phase angle differences, $\angle(A)_{j,k} = \phi_j - \phi_k$. For $\delta \ll d$, (3) describes an angular synchronization problem with highly incomplete measurements. In addition, it is also likely that these measurements are corrupted by measurement noise, which makes the problem even more challenging. This problem has applications in diverse fields such as network analysis (see[1] for an example), optics,[2] computer vision[3] and phase retrieval (see[4] for more details). We are particularly interested in applications of this problem to phase retrieval, which we describe below.

Aditya Viswanathan's e-mail: `aditya@math.msu.edu` Mark Iwen's e-mail: `markiwen@math.msu.edu`.

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

*We implicitly assume that x has no zero magnitude entries.

1.1 Phase Retrieval

The phase retrieval problem[5–10] involves the reconstruction of an unknown vector $\mathbf{x} \in \mathbb{C}^d$, up to a global phase factor, from the phaseless or magnitude measurements,

$$b_i := |\langle \mathbf{p}_i, \mathbf{x} \rangle|^2 + n_i, \quad i \in \mathbb{Z}^+. \tag{5}$$

Here, $\mathbf{p}_i \in \mathbb{C}^d$ is a measurement vector or *mask* and $n_i \in \mathbb{R}$ is measurement noise. In developing a phase retrieval method, we are interested in designing the measurement mask $\mathbf{p}_i$ such that unknown vector $\mathbf{x}$ can be recovered efficiently using a minimal number of measurements and in a manner robust to measurement noise. Several computational methods for solving this problem have been proposed, including alternating projection algorithms such as,[6] optimization-based approaches such as[7] and frame-theoretic, graph-based methods such as.[8] We restrict our attention, however, to a recently introduced[4] fast (essentially linear time in d) phase retrieval algorithm based on block-circulant measurement constructions. These type of measurements arise, for example, when computing correlations or convolutions with compactly supported masks. Moreover, this algorithm is of particular relevance to the discussion in this paper since it requires the solution of an angular synchronization problem similar to that described in (3). A brief summary of the method follows below while we refer the interested reader to[4] for more details and additional discussion.

Let $\mathbf{p}_i \in \mathbb{C}^d$ be a mask with δ non-zero entries such that $(\mathbf{p}_i)_\ell = 0$ for $\ell > \delta$; i.e., only its first δ entries are non-zero. Now consider the (squared) correlation measurements, $b_i = |\mathrm{corr}\,(\mathbf{p}_i, \mathbf{x})\,|^2$, $i = 1, 2, \ldots, N$, where N denotes the number of distinct masks used. Expanding the correlation sum, we obtain

$$(b_i)_\ell = \left| \sum_{k=1}^{\delta} (\mathbf{p_i})_k^* \cdot x_{\ell+k-1} \right|^2 = \sum_{j,k=1}^{\delta} (\mathbf{p_i})_j\, (\mathbf{p_i})_k^*\, x_{\ell+j-1}\, x_{\ell+k-1}^*, \qquad \ell \in [d], \quad i = 1, 2, \ldots, N, \tag{6}$$

where the indices are considered modulo-d. Note that we obtain d measurements – $(b_i)_1, (b_i)_2, \ldots, (b_i)_d$ for each mask $\mathbf{p}_i, i = 1, 2, \ldots, N$, and that (6) describes a linear system for the (scaled) phase differences, $(X')_{j,k} = x_j x_k^*$. As an illustrative example, consider the simple case where $d = 4$ and $\delta = 2$. Writing out the linear system (6) for these parameter values and using the notation $(\mathbf{p_i})_{j,k} := (\mathbf{p_i})_j (\mathbf{p_i})_k^*$, we obtain

$$M'\,\mathbf{x}' = \mathbf{b}, \tag{7}$$

where

$$M' = \begin{bmatrix}
(\mathbf{p}_1)_{1,1} & (\mathbf{p_1})_{1,2} & (\mathbf{p}_1)_{2,1} & (\mathbf{p}_1)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
(\mathbf{p}_2)_{1,1} & (\mathbf{p_2})_{1,2} & (\mathbf{p}_2)_{2,1} & (\mathbf{p}_2)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
(\mathbf{p}_3)_{1,1} & (\mathbf{p_3})_{1,2} & (\mathbf{p}_3)_{2,1} & (\mathbf{p}_3)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & (\mathbf{p}_1)_{1,1} & (\mathbf{p_1})_{1,2} & (\mathbf{p_1})_{2,1} & (\mathbf{p}_1)_{2,2} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & (\mathbf{p}_2)_{1,1} & (\mathbf{p_2})_{1,2} & (\mathbf{p_2})_{2,1} & (\mathbf{p}_2)_{2,2} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & (\mathbf{p}_3)_{1,1} & (\mathbf{p_3})_{1,2} & (\mathbf{p_3})_{2,1} & (\mathbf{p}_3)_{2,2} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_1)_{1,1} & (\mathbf{p_1})_{1,2} & (\mathbf{p_1})_{2,1} & (\mathbf{p}_1)_{2,2} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_2)_{1,1} & (\mathbf{p_2})_{1,2} & (\mathbf{p_2})_{2,1} & (\mathbf{p}_2)_{2,2} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_3)_{1,1} & (\mathbf{p_3})_{1,2} & (\mathbf{p_3})_{2,1} & (\mathbf{p}_3)_{2,2} & 0 & 0 \\
(\mathbf{p}_1)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_1)_{1,1} & (\mathbf{p_1})_{1,2} & (\mathbf{p_1})_{2,1} \\
(\mathbf{p}_2)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_2)_{1,1} & (\mathbf{p_2})_{1,2} & (\mathbf{p_2})_{2,1} \\
(\mathbf{p}_3)_{2,2} & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & (\mathbf{p}_3)_{1,1} & (\mathbf{p_3})_{1,2} & (\mathbf{p_3})_{2,1}
\end{bmatrix},$$

$$\mathbf{x}' = \begin{bmatrix} |x_1|^2 & x_1x_2^* & x_2x_1^* & |x_2|^2 & x_2x_3^* & x_3x_2^* & |x_3|^2 & x_3x_4^* & x_4x_3^* & |x_4|^2 & x_4x_1^* & x_1x_4^* \end{bmatrix}^T, \quad \text{and}$$

$$\mathbf{b} = \begin{bmatrix} (b_1)_1 & (b_2)_1 & (b_3)_1 & (b_1)_2 & (b_2)_2 & (b_3)_2 & (b_1)_3 & (b_2)_3 & (b_3)_3 & (b_1)_4 & (b_2)_4 & (b_3)_4 \end{bmatrix}^T.$$

With $\delta = 2$, there are $(2\delta - 1)d = 12$ unknown phase differences to be estimated. Equation (7) solves for these using measurements generated using $N = 3$ distinct masks, $\mathbf{p}_i, i = 1, 2, 3$. We draw attention to the block-circulant structure of the matrix M', which allows for efficient inversion using FFTs. This is especially desirable for large dimensional problems. Further, it can be shown that the matrix M' is well conditioned – we refer the reader to[4] for details and explicit (including deterministic) constructions of the masks $\mathbf{p}_i$.

The above framework allows us to recover phase differences even though we start with phaseless measurements. It is a specific realization of a more general framework called *lifting* – instead of solving for the unknown vector $\mathbf{x} \in \mathbb{C}^d$ directly, we first solve for the quadratic form $\mathbf{x}\mathbf{x}^* \in \mathbb{C}^{d\times d}$ using masked measurements. Lifting methods are analytically and computationally appealing since they transform non-linear (and indeed non-convex) problems such as the phase retrieval problem (5) to linear forms such as (6). However, for the block-circulant measurements described by (6) and of interest to us, we only recover the diagonal entries of the matrix $\mathbf{x}\mathbf{x}^* \in \mathbb{C}^{d\times d}$ (recall (2)).

Further, since δ is typically much smaller than d, recovering $\mathbf{x}$ from X' is no longer a trivial problem. Although the magnitude of $\mathbf{x}$ can be easily approximated using $|x_j| = \sqrt{(X')_{j,j}}$, $j \in [d]$, recovering the phase of $\mathbf{x}$ is less straightforward. The solution is to solve the angular synchronization problem determined by (3). This may be solved using the eigenvector-based method proposed in this paper to recover the phase of $\mathbf{x}$, $(\tilde{x})_j := \mathbb{e}^{\mathbb{i}\phi_j} = x_j/|x_j|$, $j \in [d]$. Therefore, efficient, accurate and robust angular synchronization methods are essential for solving phase retrieval problems arising from block-circulant measurements constructions such as those introduced above.

1.2 Related Work

Angular synchronization has mainly been studied in the context of its applications (see[1–3] for some examples). A more general formulation and mathematical analysis of the problem can be found in.[11] Further connections to graph theory, maximum likelihood estimation and convex relaxations are explored in.[12,13] In most of these works, recovery algorithms and guarantees are developed under the assumption that measurements are missing (or corrupted) at random. Their applicability to the structured and highly incomplete[†] measurement constructions of (2) is not immediately obvious.

The rest of the paper is organized as follows: §2 provides a theoretical analysis of eigenvector-based angular synchronization, while §3 provides numerical simulations showing the accuracy of the proposed method and its robustness to measurement noise. Finally, we provide some concluding remarks and future research directions in §4.

2. EIGENVECTOR-BASED ANGULAR SYNCHRONIZATION

Recall that our objective in solving the angular synchronization problem is to recover the phases of the entries of $\mathbf{x}$ given A; i.e., we want to recover $(\tilde{x})_j := \mathbb{e}^{\mathbb{i}\phi_j} = x_j/|x_j|$, $j \in [d]$ using the normalized and banded entries of the matrix $A \in \mathbb{C}^{d\times d}$ defined in (3). We show below that $\tilde{\mathbf{x}}$ can be accurately approximated using the leading eigenvector of A. This is computationally attractive since the leading eigenvector of A can be computed efficiently using the power method[14] in essentially linear time. For the discussion below, we let $D = (2\delta - 1)d$ and we always order the d real eigenvalues of a $d \times d$ Hermitian matrix as $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_d$, with possible repetitions.

The following lemmas demonstrate that eigenvector-based angular synchronization should indeed be possible. In particular, this first lemma shows that $\tilde{\mathbf{x}}$ is always an eigenvector for the largest eigenvalue of A.

Lemma 1. The largest eigenvalue of A is $\lambda_1 = 2\delta - 1$, and $\tilde{\mathbf{x}}$ is an eigenvector of A with this maximal eigenvalue.

Proof: Let $\lambda_{\max}$ denote the largest magnitude eigenvalue of A. Since A is Hermitian, the spectral radius of $A = \|A\|_2 = |\lambda_{\max}|$. Furthermore, one can see that

$$|\lambda_{\max}| \leq \|A\|_\infty = \max_{j\in[d]} \sum_{k=1}^{d} |A_{j,k}| = 2\delta - 1,$$

using standard results concerning matrix norms.[15] Finally, one can also see that $A\tilde{\mathbf{x}} = (2\delta - 1)\tilde{\mathbf{x}}$. □

The next lemma shows that the smallest eigenvalue of A, which is generally negative, is strictly smaller in magnitude than λ_1. Hence, λ_1 is indeed the unique *largest magnitude* eigenvalue of A *if* $\lambda_1 > \lambda_2$.

[†]especially those arising from problems in phase retrieval using block-circulant measurements

Lemma 2. The smallest eigenvalue of A has magnitude $|\lambda_d| \leq 2\delta - 3 = \lambda_1 - 2$ for all $\delta \geq 3$.

Proof: Suppose $\lambda_d \geq 0$. In this case, we have that

$$\lambda_1(d - \lambda_1) = \|A\|_F^2 - \lambda_1^2 = \sum_{j=2}^{d} \lambda_j^2 \geq (d-1)\lambda_d^2.$$

Thus, $\lambda_d \leq \sqrt{\lambda_1} \leq \lambda_1 - 2 = 2\delta - 3$ since $\lambda_1 \geq 4$. Now suppose that $\lambda_d < 0$. Let $\mathbf{y}$ be an eigenvector of A with eigenvalue λ_d, and let $j \in [d]$ be such that $|y_j| = \|\mathbf{y}\|_\infty$. Considering the equation $(A - \lambda_d I)\mathbf{y} = \mathbf{0}$, we note that the inner product of the j^{th} row of $A - \lambda_d I$ with $\mathbf{y}$ must satisfy

$$-(1 + |\lambda_d|)y_j = \sum_{0 < |k-j \mod d| < \delta} y_k \mathbb{e}^{\mathbb{i}(\phi_j - \phi_k)}$$

so that

$$\|\mathbf{y}\|_\infty = |y_j| = \frac{1}{1 + |\lambda_d|} \left| \sum_{0 < |k-j \mod d| < \delta} y_k \mathbb{e}^{\mathbb{i}(\phi_j - \phi_k)} \right| \leq \frac{2\delta - 2}{1 + |\lambda_d|} \cdot \|y\|_\infty.$$

This last inequality can not hold unless $1 + |\lambda_d| \leq 2\delta - 2$. The desired result follows. $\square$

This third lemma proves that $\tilde{\mathbf{x}}$ is the *only* eigenvector of A with eigenvalue λ_1. This establishes that $\lambda_1 > \lambda_2$ does indeed hold.

Lemma 3. The largest two eigenvalues of A are not equal, $\lambda_1 \neq \lambda_2$, for all $\delta \geq 2$.

Proof: Let $\mathbf{y}$ be an eigenvector of A with eigenvalue λ_1, and let $j \in [d]$ be such that $|y_j| = \|\mathbf{y}\|_\infty$ as above. Considering the equation $(A - \lambda_1 I)\mathbf{y} = \mathbf{0}$, we see that the inner product of the j^{th} row of $A - \lambda_1 I$ with $\mathbf{y}$ must satisfy

$$(2\delta - 2)y_j = \sum_{0 < |k-j \mod d| < \delta} y_k \mathbb{e}^{\mathbb{i}(\phi_j - \phi_k)} \tag{8}$$

so that

$$\|\mathbf{y}\|_\infty = |y_j| = \frac{1}{2\delta - 2} \left| \sum_{0 < |k-j \mod d| < \delta} y_k \mathbb{e}^{\mathbb{i}(\phi_j - \phi_k)} \right| \leq \frac{1}{2\delta - 2} \sum_{0 < |k-j \mod d| < \delta} |y_k|.$$

Note that this last inequality can not hold unless $|y_k| = \|\mathbf{y}\|_\infty$ also holds for all $2\delta - 2$ values of k with $0 < |k - j \mod d| < \delta$. Repeating this argument we come to the conclusion that $|y_k| = \|\mathbf{y}\|_\infty$ must indeed hold for all $k \in [d]$. Revisiting (8) and dividing through by y_j we learn that

$$(2\delta - 2) = \sum_{0 < |k-j \mod d| < \delta} \mathbb{e}^{\mathbb{i}(\theta_k - \theta_j)} \mathbb{e}^{\mathbb{i}(\phi_j - \phi_k)}$$

holds, where $y_j = \|\mathbf{y}\|_\infty \mathbb{e}^{\mathbb{i}\theta_j}$, for all $j \in [d]$. As a result, it must be the case that $\theta_k - \phi_k = \theta_j - \phi_j$ holds for all $j \in [d], 0 < |k - j \mod d| < \delta$. It follows that $\mathbf{y}$ must be scalar multiple of $\tilde{\mathbf{x}}$. $\square$

Finally, one can also see that the largest eigenvalue will contain a substantial fraction of the spectral energy of A if δ is taken to be sufficiently large.

Lemma 4. Let $c \in \mathbb{R}^+$. We will have $c\sqrt{\sum_{j=2}^{d} \lambda_j^2} \leq \lambda_1$ whenever $\delta \geq \frac{1+(d+1)c^2}{2(1+c^2)}$.

Proof: Whenever $\delta \geq \frac{1+(d+1)c^2}{2(1+c^2)}$ we will have $dc^2 \leq (2\delta - 1)(1 + c^2) = \lambda_1(1 + c^2)$, using Lemma 1. Continuing, we learn that

$$\lambda_1^2 \geq \lambda_1(d - \lambda_1)c^2 = (\|A\|_F^2 - \lambda_1^2)c^2 = \left(\sum_{j=2}^{d} \lambda_j^2 \right) c^2.$$

Taking square roots now yields the desired result. $\square$

We now have achieved the desired final result of our analysis: In the noiseless case a simple power method may be used on A from (3) in order to recover $\tilde{\mathbf{x}}$ from (4). Furthermore, this simple scheme will be efficient whenever λ_1 dominates both λ_2 and λ_d sufficiently well (which is a function of the choice of δ for fixed d).

3. NUMERICAL RESULTS

We now present numerical results studying the spectral properties of the matrix A and demonstrating the robustness of eigenvector-based angular synchronization to measurement noise. In all simulations below, i.i.d. standard complex Gaussian signals are used as the test signals $\mathbf{x}$. All simulations were performed in Matlab; in particular, Matlab's `eigs` command was used to estimate the desired eigenvalues and eigenvectors. Each data point in the figures below was obtained as the average of 100 trials.

We begin by studying the spectral properties of A – in particular, the spectral gap, $\lambda_1 - \lambda_2$, and the ratio $\frac{\lambda_1}{\max(|\lambda_2|, |\lambda_d|)}$. As discussed in §2, these quantities are important in ensuring that the recovered solution is unique, accurate, and can be efficiently computed using an iterative scheme such as the power method. Figure 1a illustrates the dependence of the spectral gap $\lambda_1 - \lambda_2$ on the band parameter δ for a fixed problem dimension d. Plots for $d = 64, 256$ and 1024 are provided which show that the spectral gap is always positive, increases with δ (the plots indicate that the spectral gap appears to grow cubically with δ), and that $\lambda_1 - \lambda_2 \approx d$ when $\delta \approx d/2$ (i.e., as expected, the matrix A becomes rank-1 when we measure all possible phase differences). We also plot the dependence of the ratio $\frac{\lambda_1}{\max(|\lambda_2|, |\lambda_d|)}$ on δ in Figure 1b. Note that this ratio is always greater than 1 as predicted in §2, indicating that the iterative method used to compute the leading eigenvector (such as the power method) will converge.

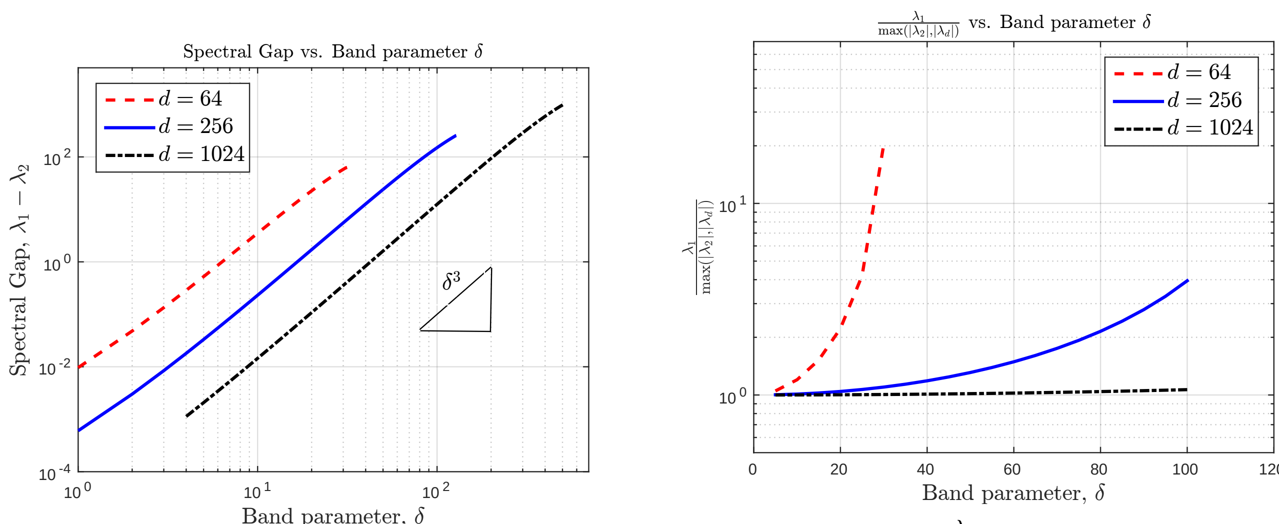

(a) Spectral gap, $\lambda_1 - \lambda_2$ vs. δ.

(b) $\frac{\lambda_1}{\max(|\lambda_2|, |\lambda_d|)}$ vs. δ.

Figure 1: Spectral Properties of A – Dependence of the spectral gap, $\lambda_1 - \lambda_2$ and the ratio $\frac{\lambda_1}{\max(|\lambda_2|, |\lambda_d|)}$ on the band parameter δ.

Similarly, Figure 2 illustrates the dependence of the spectral gap and the ratio $\frac{\lambda_1}{\max(|\lambda_2|, |\lambda_d|)}$ on the problem dimension d. For a fixed band parameter δ, Figure 2a shows that the spectral gap appears to vary like $\frac{1}{d^2}$. For completeness, we also plot the variation in the ratio $\frac{\lambda_1}{\max(|\lambda_2|, |\lambda_d|)}$ as a function of d in Figure 2b.

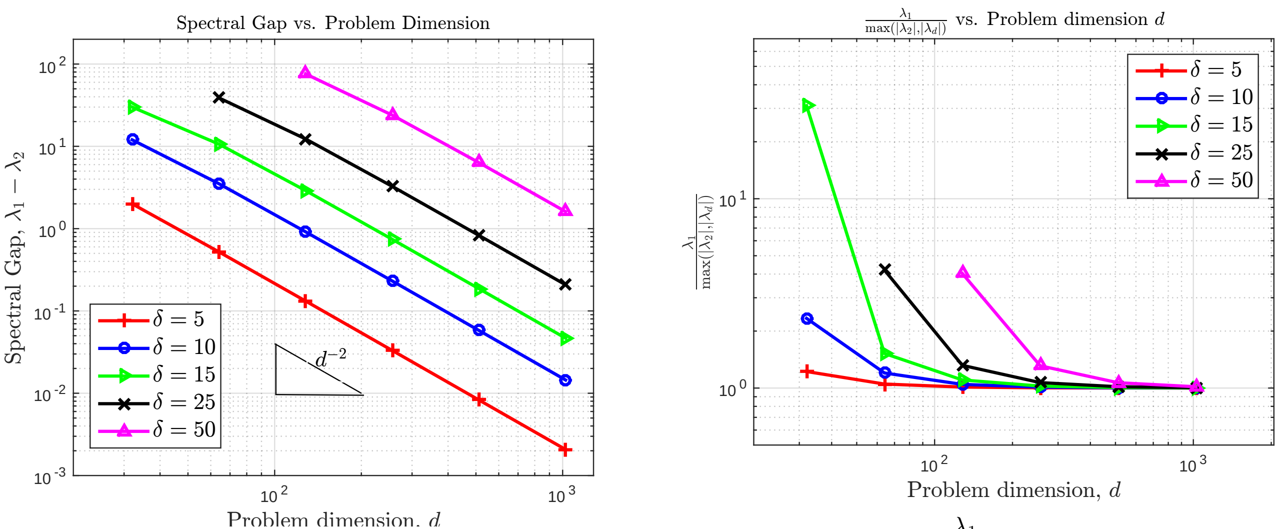

(a) Spectral gap, $\lambda_1 - \lambda_2$ vs. d.

(b) $\dfrac{\lambda_1}{\max(|\lambda_2|, |\lambda_d|)}$ vs. d.

Figure 2: Spectral Properties of A – Dependence of the spectral gap, $\lambda_1 - \lambda_2$ and the ratio $\dfrac{\lambda_1}{\max(|\lambda_2|, |\lambda_d|)}$ on the problem dimension d.

Next, we present numerical simulations demonstrating the robustness of the proposed method to measurement noise. Although we defer a rigorous theoretical analysis to future work, Figures 3 and 4 show that eigenvector-based angular synchronization is indeed robust to measurement errors numerically. For the results below, we assume that following noisy measurement model:

$$(Y')_{j,k} = (X')_{j,k} + \eta_{j,k} = \begin{cases} (\mathbf{x}\mathbf{x}^*)_{j,k} + \eta_{j,k} & \text{if } |j - k \bmod d| < \delta \\ 0 & \text{otherwise,} \end{cases} . \tag{9}$$

where $\eta_{j,k}$ denotes additive measurement noise and is drawn from an i.i.d complex Gaussian distribution. Moreover, we apply the eigenvector method to $(Y' + (Y')^*)/2$ to ensure that Hermitian symmetry is maintained. We report added noise in terms of signal to noise ratio (SNR, in dB), with

$$\text{Added Noise SNR (dB)} = 10 \log_{10} \left(\frac{\|X'\|_F^2}{\|Y' - X'\|_F^2} \right).$$

Similarly, we report reconstruction errors in dB, with

$$\text{Reconstruction Error (dB)} = 10 \log_{10} \left(\frac{\min\limits_{\theta \in [0, 2\pi)} \|\tilde{\mathbf{y}} - e^{\mathrm{i}\theta} \tilde{\mathbf{x}}\|_2^2}{\|\tilde{\mathbf{x}}\|_2^2} \right),$$

where $\tilde{\mathbf{x}}$ is the true phase of the unknown signal $\mathbf{x}$ (refer to (4)), $\tilde{\mathbf{y}}$ is the eigenvector-based approximation to $\tilde{\mathbf{x}}$ and θ is a global phase offset. Figure 3 plots the reconstruction error (in dB) as a function of the added noise level for a fixed problem dimension d. Figure 3a plots the reconstruction performance for $d = 256$, while Figure 3b plots the reconstruction errors for $d = 1024$. In each case, plots for a few different values of the band parameter δ are included to illustrate its effect on reconstruction performance. As expected, the reconstruction error is roughly of the same order as added noise, with performance improving with increase in δ. We note that for large problem dimensions and with high levels of added noise, choosing a small value of δ can result in poor performance (see, for example, Figure 3b). This is not surprising, given that the spectral gap is likely to be small in this case. Similar conclusions can be drawn from Figure 4, where the reconstruction error (in dB) is plotted as a function of the added noise level for a fixed value of the band parameter δ.

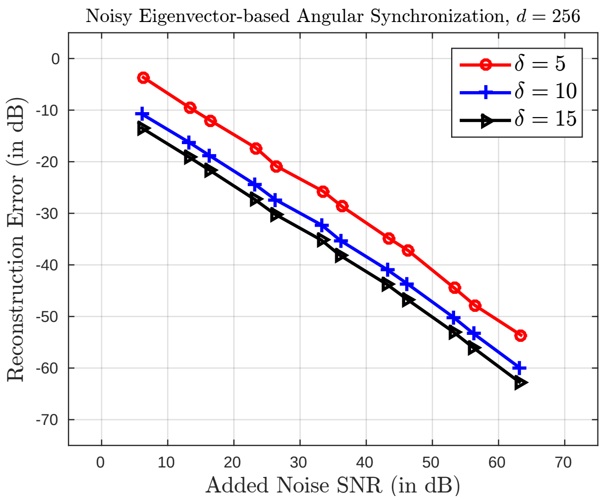

(a) Problem dimension, $d = 256$.

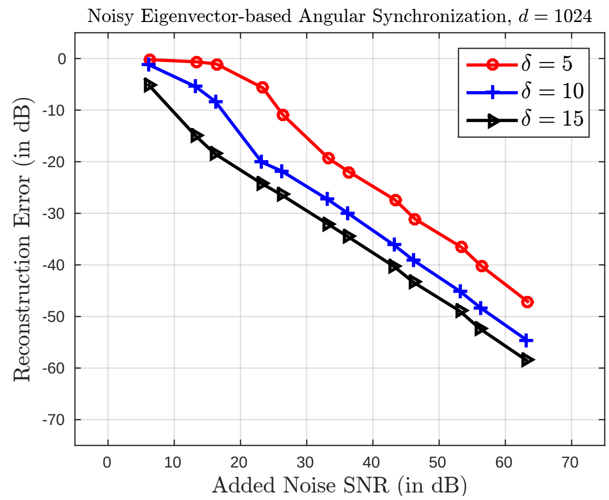

(b) Problem dimension, $d = 1024$.

Figure 3: Performance of eigenvector-based angular synchronization in the presence of measurement noise. This figure plots the reconstruction error as a function of added noise level for a fixed problem dimension, d.

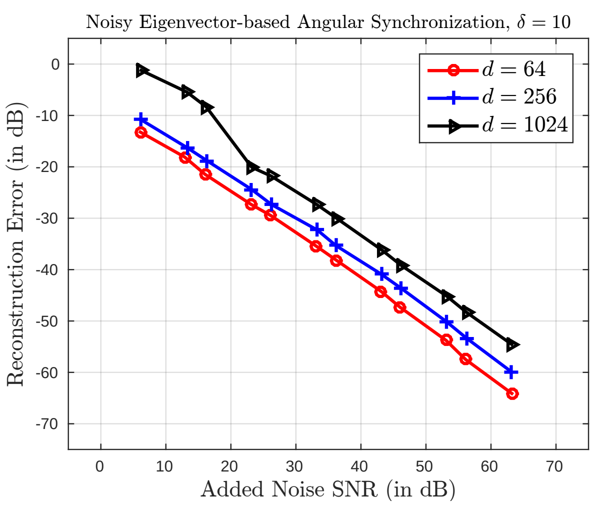

(a) Band parameter, $\delta = 10$.

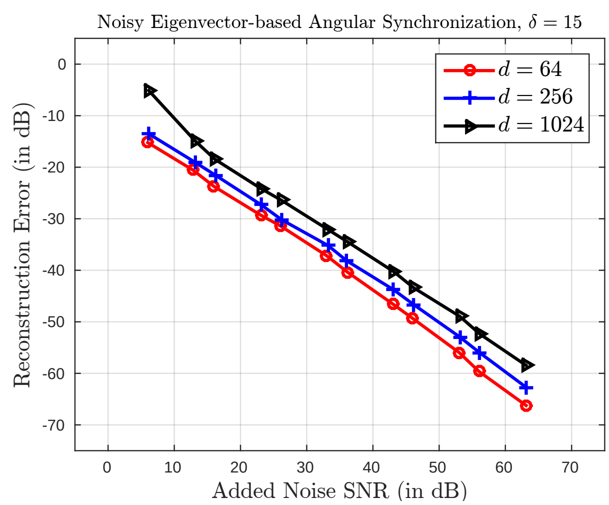

(b) Band parameter, $\delta = 15$.

Figure 4: Performance of eigenvector-based angular synchronization in the presence of measurement noise. This figure plots the reconstruction error as a function of added noise level for a fixed value of the band parameter, δ.

4. CONCLUDING REMARKS

In this paper, we have considered angular synchronization from highly incomplete information – a problem which arises from applications in phase retrieval using block circulant measurement constructions. We use an eigenvector-based method to solve this angular synchronization problem and analyze its theoretical performance. Numerical results show that this method is accurate, efficient, and robust to measurement errors. Some future research directions include a theoretical analysis of the method in the presence of measurement errors, connections and comparisons of this method to other angular synchronization algorithms such as alternating projection methods, as well as applying the theoretical and numerical findings presented here to the phase retrieval problem.

REFERENCES

[1] Giridhar, A. and Kumar, P. R., "Distributed clock synchronization over wireless networks: Algorithms and analysis," in [*Proc. 45th IEEE Conf. Decis. Control*], 4915–4920, IEEE (2006).

[2] Rubinstein, J. and Wolansky, G., "Reconstruction of optical surfaces from ray data," *Opt. Rev.* **8**(4), 281–283 (2001).

[3] Frankot, R. T. and Chellappa, R., "A method for enforcing integrability in shape from shading algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.* **10**(4), 439–451 (1988).

[4] Iwen, M. A., Viswanathan, A., and Wang, Y., "Fast phase retrieval for high-dimensions," *preprint, arXiv:1501.02377* (2015).

[5] Balan, R., Casazza, P., and Edidin, D., "On signal reconstruction without phase," *Appl. Comput. Harmon. Anal.* **20**(3), 345–356 (2006).

[6] Fienup, J. R., "Reconstruction of an object from the modulus of its Fourier transform," *Opt. Lett.* **3**, 27–29 (1978).

[7] Candes, E. J., Strohmer, T., and Voroninski, V., "PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming," *Commun. Pure Appl. Math.* **66**(8), 1241–1274 (2013).

[8] Alexeev, B., Bandeira, A. S., Fickus, M., and Mixon, D. G., "Phase retrieval with polarization," *SIAM J. Imag. Sci.* **7**(1), 35–66 (2014).

[9] Eldar, Y. C. and Mendelson, S., "Phase retrieval: Stability and recovery guarantees," *Appl. Comput. Harmon. Anal.* **36**(3), 473–494 (2014).

[10] Fickus, M., Mixon, D., Nelson, A., and Wang, Y., "Phase retrieval from very few measurements," *Linear Algebra Appl.* **449**, 475–499 (2014).

[11] Singer, A., "Angular synchronization by eigenvectors and semidefinite programming," *Appl. Comput. Harmon. Anal.* **30**, 20–36 (2011).

[12] Bandeira, A. S., Singer, A., and Spielman, D. A., "A Cheeger inequality for the graph connection laplacian," *SIAM J. Matrix Anal. Appl.* **34**(4), 1611–1630 (2013).

[13] Bandeira, A. S., Boumal, N., and Singer, A., "Tightness of the maximum likelihood semidefinite relaxation for angular synchronization," (2014). preprint, arXiv:1411.3272.

[14] Golub, G. H. and Van Loan, C. F., [*Matrix Computations*], The Johns Hopkins University Press, 3 ed. (1996).

[15] Horn, R. A. and Johnson, C. R., [*Matrix Analysis*], Cambridge University Press, 2 ed. (2012).