

PHASE RETRIEVAL WITH POLARIZATION

BORIS ALEXEEV*, AFONSO S. BANDEIRA†, MATTHEW FICKUS‡, AND DUSTIN G. MIXON§

Abstract. In many areas of imaging science, it is difficult to measure the phase of linear measurements. As such, one often wishes to reconstruct a signal from intensity measurements, that is, perform *phase retrieval*. In this paper, we provide a novel measurement design which is inspired by interferometry and exploits certain properties of expander graphs. We also give an efficient phase retrieval procedure, and use recent results in spectral graph theory to produce a stable performance guarantee which rivals the guarantee for PhaseLift in [14]. We use numerical simulations to illustrate the performance of our phase retrieval procedure, and we compare reconstruction error and runtime with a common alternating-projections-type procedure.

Keywords: phase retrieval, polarization, expander graph, spectral clustering, angular synchronization.

1. Introduction. Given a collection of vectors $\Phi := \{\varphi_\ell\}_{\ell=1}^N \subseteq \mathbb{C}^M$ and a signal $x \in \mathbb{C}^M$, consider measurements of the form

$$z_\ell := |\langle x, \varphi_\ell \rangle|^2 + \nu_\ell, \quad (1.1)$$

where ν_ℓ is noise; we call these noisy *intensity measurements*. Several areas of imaging science, such as X-ray crystallography [25, 37, 38], diffraction imaging [11], astronomy [17] and optics [49], use measurements of this form with the intent of reconstructing the original signal; this inverse problem is called *phase retrieval*. Note that in the measurement process (1.1), we inherently lose some information about x . Indeed, for every $\omega \in \mathbb{C}$ with $|\omega| = 1$, we see that x and ωx produce the same intensity measurements. Thus, the best one can hope to do with the intensity measurements of $x \in \mathbb{C}^M$ is reconstruct the class $[x] \in \mathbb{C}^M/\sim$, where \sim is the equivalence relation of being identical up to a global phase factor.

In practice, phase retrieval falls short of determining the original signal up to global phase. First of all, the intensity measurement process that is used, say $\mathcal{A} : \mathbb{C}^M/\sim \rightarrow \mathbb{R}_{\geq 0}^N$ with $\mathcal{A}(x) = |\Phi^* x|^2$ (entrywise) and viewing $\Phi = [\varphi_1 \cdots \varphi_N]$, often lacks injectivity, making it impossible to reconstruct uniquely. Moreover, the phase retrieval algorithms that are used in practice take alternating projections onto the column space of Φ^* (to bring phase to the measurements) and onto the nonconvex set of vectors y whose entry magnitudes match the intensity measurements $|\Phi^* x|^2$ (to maintain fidelity in the magnitudes) [21, 23, 24]. Unfortunately, the convergence of these algorithms (and various modifications thereof) is particularly sensitive to the choice of initial phases [36].

These deficiencies have prompted two important lines of research in phase retrieval:

- (i) For which measurement designs Φ is $[x] \mapsto |\Phi^* x|^2$ injective?
- (ii) For which injective designs can $[x]$ be reconstructed stably and efficiently?

A first step toward solving (i) is determining how large N must be in order for \mathcal{A} to be injective. It remains an open problem to find the smallest such N [7], but embedding results in differential geometry give that $N \geq (4 + o(1))M$ is necessary [4, 27]. As for sufficiency, Balan, Casazza and Edidin [6] show that for almost every choice of Φ , \mathcal{A} is injective whenever $N \geq 4M - 2$, and recent work has leveraged more structured constructions to slightly decrease this number of measurement vectors [10, 39]. Though the community has investigated various conditions for injectivity, very little is known about how to stably and efficiently

*Department of Mathematics, Princeton University, Princeton, NJ 08544, USA (balaxeev@math.princeton.edu).

†Program in Applied and Computational Mathematics (PACM), Princeton University, Princeton, NJ 08544, USA (ajsb@math.princeton.edu).

‡Department of Mathematics and Statistics, Air Force Institute of Technology, Wright-Patterson Air Force Base, OH 45433, USA (matthew.fickus@gmail.com).

§Department of Mathematics and Statistics, Air Force Institute of Technology, Wright-Patterson Air Force Base, OH 45433, USA (dustin.mixon@afit.edu).

reconstruct in the injective case. In fact, some instances of the phase retrieval problem are known to be NP-complete [45], and so any general reconstruction process is necessarily inefficient, assuming $P \neq NP$.

This leads one to attempt provably stable and efficient reconstruction from measurements of the form (1.1) with particular ensembles Φ . Until recently, this was only known to be possible in cases where $N = \Omega(M^2)$ [5]. By contrast, the state of the art comes from Candès, Strohmer and Voroninski [14], who use semidefinite programming to stably reconstruct from $N = \mathcal{O}(M \log M)$ Gaussian-random measurements. There is other work along this vein [12, 13, 18, 48] which also uses semidefinite programming and provides related guarantees. Typically, semidefinite programs are solved via interior point methods. Since these methods are computationally expensive, in practice, one is inclined to instead use faster numerical methods, but these lack performance guarantees.

While (i) and (ii) above describe what has been a more theoretical approach to phase retrieval, practitioners have meanwhile considered various alternatives to the intensity measurement process (1.1). A common theme among these alternatives is the use of interference to extract more information about the desired signal. For example, *holography* interferes the signal of interest $x \in \mathbb{C}^M$ with a known reference signal $y \in \mathbb{C}^M$, taking measurements of the form $|F(x + \omega y)|^2$, where $\omega \in \mathbb{C}$ has unit modulus and F denotes the Fourier transform [19]; three such measurements (i.e., $3M$ scalar measurements) suffice for injectivity [50]. Alternatively, *spectral phase interferometry for direct electric-field reconstruction (SPIDER)* interferes the signal of interest $x \in \mathbb{C}^M$ with time- and frequency-shifted versions of itself $Sx \in \mathbb{C}^M$, taking measurements of the form $|F(x + Sx)|^2$ [29]; while popular in practice for ultrashort pulse measurement, SPIDER fails to accurately resolve the relative phase of well-separated frequency components [31]. Another interesting approach is *ptychography*, in which overlapping spatial components $P_i x, P_j x \in \mathbb{C}^M$ are interfered with each other, and measurements have the form $|F(P_i x + P_j x)|^2$ [44]. Recently, *vectorial phase retrieval* was proposed, in which two unknown signals $x, y \in \mathbb{C}^M$ are interfered with each other, and the measurements are $|Fx|^2, |Fy|^2, |F(x + y)|^2$ and $|F(x + iy)|^2$ [43]; furthermore, [42] gives that almost every pair of signals is uniquely determined by these $4M$ scalar measurements, in which case both signals can be reconstructed using the polarization identity. While practitioners seem to have identified interference as an instrumental technique for quality phase retrieval, the reconstruction algorithms which are typically used, much like the classical algorithms in [21, 23, 24], are iterative and lack convergence guarantees (e.g., [19, 34], though [42, 43] are noteworthy exceptions).

Returning to measurements of the form (1.1), this paper combines ideas from both state-of-the-art theory and state-of-the-art practice by proposing an exchange of sorts: If you already have $\mathcal{O}(M \log M)$ Gaussian-random measurements vectors (as prescribed in [14]), then we offer a faster reconstruction method with a stable performance guarantee, but at the price of $\mathcal{O}(M \log M)$ additional (non-adaptive) measurements. These new measurement vectors are interferometry-inspired combinations of the originals, and the computational speedups gained in reconstruction come from our use of different spectral methods. While the ideas in this paper can be applied for phase retrieval of 2-D images, we focus on the 1-D case for simplicity. Also, note that the sequel [8] leverages the techniques of this paper to construct masked Fourier measurements, thereby mimicking the illumination methodology of [12]; we suspect that these ideas can be similarly leveraged to tackle a wide variety of practical instances of the phase retrieval problem. To help motivate our measurement design and phase retrieval procedure, we start in the next section by considering the simpler, noiseless case. In this case, the success of our method follows from a neat trick involving the polarization identity along with some well-known results in the theory of expander graphs. In Section 3, we modify the method to obtain provable stability in the noisy case; here, we exploit some recent developments in spectral graph theory. Our results are then corroborated by simulations in Section 4. We give concluding remarks in Section 5, and provide the more technical proofs in the appendix.

2. The noiseless case. In this section, we provide a new measurement design and reconstruction algorithm for phase retrieval. Here, we specifically address the noiseless case, in which ν_ℓ in (1.1) is zero for every $\ell = 1, \dots, N$; this case will give some intuition for a more stable version of our techniques, which we introduce in the next section. In the noiseless case, we will use on the order of the fewest measurements

possible, namely $N = \mathcal{O}(M)$, where M is the dimension of the signal.

Before stating our measurement design and phase retrieval procedure, we motivate both with some discussion. Take a finite set V , and suppose we take intensity measurements of $x \in \mathbb{C}^M$ with a set $\Phi_V := \{\varphi_i\}_{i \in V}$ that spans \mathbb{C}^M . Again, we wish to recover x up to a global phase factor. Having $|\langle x, \varphi_i \rangle|$ for every $i \in V$, we claim it suffices to determine the relative phase between $\langle x, \varphi_i \rangle$ and $\langle x, \varphi_j \rangle$ for all pairs $i \neq j$. Indeed, if we had this information, we could arbitrarily assign some nonzero coefficient $c_i = |\langle x, \varphi_i \rangle|$ to have positive phase. If $\langle x, \varphi_j \rangle$ is also nonzero, then it has well-defined relative phase

$$\rho_{ij} := \left(\frac{\langle x, \varphi_i \rangle}{|\langle x, \varphi_i \rangle|} \right)^{-1} \frac{\langle x, \varphi_j \rangle}{|\langle x, \varphi_j \rangle|}, \quad (2.1)$$

which determines the coefficient by multiplication: $c_j = \rho_{ij} |\langle x, \varphi_j \rangle|$. Otherwise when $\langle x, \varphi_j \rangle = 0$, we naturally take $c_j = 0$, and for notational convenience, we arbitrarily take $\rho_{ij} = 1$. From here, the original signal's equivalence class $[x] \in \mathbb{C}^M / \sim$ can be identified by applying the canonical dual frame $\{\tilde{\varphi}_j\}_{j \in V}$, namely the Moore-Penrose pseudoinverse, of Φ_V :

$$\sum_{j \in V} c_j \tilde{\varphi}_j = \sum_{j \in V} \rho_{ij} |\langle x, \varphi_j \rangle| \tilde{\varphi}_j = \left(\frac{\langle x, \varphi_i \rangle}{|\langle x, \varphi_i \rangle|} \right)^{-1} \sum_{j \in V} \langle x, \varphi_j \rangle \tilde{\varphi}_j = \left(\frac{\langle x, \varphi_i \rangle}{|\langle x, \varphi_i \rangle|} \right)^{-1} x \in [x]. \quad (2.2)$$

Having established the utility of the relative phase between coefficients, we now seek some method of extracting this information. To this end, we turn to a special version of the polarization identity:

LEMMA 2.1 (Mercedes-Benz Polarization Identity). *Take $\zeta := e^{2\pi i/3}$. Then for any $a, b \in \mathbb{C}$,*

$$\bar{a}b = \frac{1}{3} \sum_{k=0}^2 \zeta^k |a + \zeta^{-k}b|^2. \quad (2.3)$$

Proof. We start by expanding the right-hand side of (2.3):

$$\text{RHS} := \frac{1}{3} \sum_{k=0}^2 \zeta^k |a + \zeta^{-k}b|^2 = \frac{1}{3} \sum_{k=0}^2 \zeta^k (|a|^2 + 2\text{Re}(\zeta^{-k}\bar{a}b) + |b|^2) = \frac{2}{3} \sum_{k=0}^2 \zeta^k \text{Re}(\zeta^{-k}\bar{a}b).$$

Multiplying, we find

$$\text{Re}(\zeta^{-k}\bar{a}b) = \text{Re}(\zeta^{-k})\text{Re}(\bar{a}b) - \text{Im}(\zeta^{-k})\text{Im}(\bar{a}b) = \text{Re}(\zeta^k)\text{Re}(\bar{a}b) + \text{Im}(\zeta^k)\text{Im}(\bar{a}b).$$

We substitute this into our expression for RHS:

$$\begin{aligned} \text{Re}(\text{RHS}) &= \frac{2}{3} \left[\text{Re}(\bar{a}b) \sum_{k=0}^2 (\text{Re}(\zeta^k))^2 + \text{Im}(\bar{a}b) \sum_{k=0}^2 \text{Re}(\zeta^k)\text{Im}(\zeta^k) \right], \\ \text{Im}(\text{RHS}) &= \frac{2}{3} \left[\text{Re}(\bar{a}b) \sum_{k=0}^2 \text{Re}(\zeta^k)\text{Im}(\zeta^k) + \text{Im}(\bar{a}b) \sum_{k=0}^2 (\text{Im}(\zeta^k))^2 \right]. \end{aligned}$$

Finally, we apply the following easy-to-verify identities:

$$\sum_{k=0}^2 (\text{Re}(\zeta^k))^2 = \sum_{k=0}^2 (\text{Im}(\zeta^k))^2 = \frac{3}{2}, \quad \sum_{k=0}^2 \text{Re}(\zeta^k)\text{Im}(\zeta^k) = 0,$$

which yield $\text{RHS} = \bar{a}b$. \square

The above polarization identity can also be proved by viewing $\{\zeta^k\}_{k=0}^2$ as a Mercedes-Benz frame in \mathbb{R}^2 and $\frac{2}{3} \sum_{k=0}^2 \zeta^k \text{Re}(\zeta^{-k}u)$ as the corresponding reconstruction formula for $u \in \mathbb{C} = \mathbb{R}^2$. We can now use this polarization identity to determine relative phase (2.1):

$$\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle = \frac{1}{3} \sum_{k=0}^2 \zeta^k |\langle x, \varphi_i \rangle + \zeta^{-k} \langle x, \varphi_j \rangle|^2 = \frac{1}{3} \sum_{k=0}^2 \zeta^k |\langle x, \varphi_i + \zeta^k \varphi_j \rangle|^2. \quad (2.4)$$

Thus, if in addition to Φ_V we measure with $\{\varphi_i + \zeta^k \varphi_j\}_{k=0}^2$, we can use (2.4) to determine $\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle$ and then normalize to get the relative phase:

$$\rho_{ij} := \left(\frac{\langle x, \varphi_i \rangle}{|\langle x, \varphi_i \rangle|} \right)^{-1} \frac{\langle x, \varphi_j \rangle}{|\langle x, \varphi_j \rangle|} = \frac{\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle}{|\langle x, \varphi_i \rangle \langle x, \varphi_j \rangle|}, \quad (2.5)$$

provided both $\langle x, \varphi_i \rangle$ and $\langle x, \varphi_j \rangle$ are nonzero. To summarize our discussion of reconstructing a single signal, if we measure with Φ_V and $\{\varphi_i + \zeta^k \varphi_j\}_{k=0}^2$ for every pair $i, j \in V$, then we can recover $[x]$. However, such a method uses $|V| + 3 \binom{|V|}{2}$ measurements, and since Φ_V must span \mathbb{C}^M , we necessarily have $|V| \geq M$ and thus a total of $\Omega(M^2)$ measurements. Note that a nearly identical formulation of these $\Omega(M^2)$ measurements appears in Theorem 5.2 of [5]. The proof of this result shows how one can adaptively appeal to only $\mathcal{O}(M)$ of the measurements to perform phase retrieval, suggesting that most of these measurements are actually unnecessary. However, since the $\mathcal{O}(M)$ measurements that end up being used are highly dependent on the signal being measured, one cannot blindly restrict to a particular subcollection of $\mathcal{O}(M)$ measurement vectors a priori without forfeiting injectivity.

In pursuit of $\mathcal{O}(M)$ measurements, take some simple graph $G = (V, E)$, arbitrarily assign a direction to each edge, and only take measurements with Φ_V and $\Phi_E := \bigcup_{(i,j) \in E} \{\varphi_i + \zeta^k \varphi_j\}_{k=0}^2$. To recover $[x]$, we again arbitrarily assign some nonzero vertex measurement to have positive phase, and then we propagate relative phase information along the edges by multiplication to determine the phase of the other vertex measurements relative to the original vertex measurement:

$$\rho_{ik} = \rho_{ij} \rho_{jk}. \quad (2.6)$$

However, if x is orthogonal to a given vertex vector, then that measurement is zero, and so relative phase information cannot propagate through the corresponding vertex; indeed, such orthogonality has the effect of removing the vertex from the graph, and for some graphs, this will prevent recovery. For example, if G is a star, then x could be orthogonal to the vector corresponding to the internal vertex, whose removal would render the remaining graph edgeless. That said, we should select Φ_V and G so as to minimize the impact of orthogonality with vertex vectors.

First, we can take Φ_V to be *full spark*, that is, Φ_V has the property that every subcollection of M vectors spans. Full spark frames appear in a wide variety of applications. Explicit deterministic constructions of them are given in [1, 41]. For example, we can select the first M rows of the $|V| \times |V|$ discrete Fourier transform matrix, and take Φ_V to be the columns of the resulting $M \times |V|$ matrix; in this case, the fact that Φ_V is full spark follows from the Vandermonde determinant formula. In our application, Φ_V being full spark will be useful for two reasons. First, this implies that $x \neq 0$ is orthogonal to at most $M - 1$ members of Φ_V , thereby limiting the extent of x 's damage to our graph. Additionally, Φ_V being full spark frees us from requiring the graph to be connected after the removal of vertices; indeed, any remaining component of size M or more will correspond to a subcollection of Φ_V that spans, meaning it has a dual frame to reconstruct with. It remains to find a graph of $\mathcal{O}(M)$ vertices and edges that maintains a size- M component after the removal of any $M - 1$ vertices.

To this end, we consider a well-studied family of sparse graphs known as *expander graphs*. We choose these graphs for their notably strong connectivity properties. There is a combinatorial definition of expander graphs, but we will focus on the spectral definition. Given a d -regular graph G of n vertices, consider its adjacency matrix A , and define the *Laplacian* to be $L := I - \frac{1}{d}A$; if G were not regular, we would consider

the diagonal matrix D of vertex degrees and define the Laplacian to be $L := I - D^{-1/2}AD^{-1/2}$. This is often called the *normalized* Laplacian in the literature, but we make no distinction here. We are particularly interested in the eigenvalues of the Laplacian: $0 = \lambda_1 \leq \dots \leq \lambda_n$. The second eigenvalue λ_2 of the Laplacian is called the *spectral gap* of the graph, and as we shall see, this value is particularly useful in evaluating the graph's connectivity. We say G has *expansion* λ if $\{\lambda_2, \dots, \lambda_n\} \subseteq [1 - \lambda, 1 + \lambda]$; note that since $1 - \lambda \leq \lambda_2$, small expansion implies large spectral gap. Furthermore, a family of d -regular graphs $\{G_i\}_{i=1}^\infty$ is a *spectral expander family* if there exists $c < 1$ such that every G_i has expansion $\lambda(G_i) \leq c$. Since d is constant over an expander family, expanders with many vertices have particularly few edges. There are many results which describe the connectivity of expanders, but the following is particularly relevant to our application:

LEMMA 2.2 (Spectral gap grants connectivity [26]). *Consider a d -regular graph G of n vertices with spectral gap λ_2 . For all $\varepsilon \leq \frac{\lambda_2}{6}$, removing any εdn edges from G results in a connected component of size $\geq (1 - \frac{2\varepsilon}{\lambda_2})n$.*

Note that removing εn vertices from a d -regular graph necessarily removes $\leq \varepsilon dn$ edges, and so this lemma directly applies. For our application, we want to guarantee that the removal of any $M - 1$ vertices maintains a size- M component. To do this, we will ensure both (i) $M - 1 \leq \varepsilon n$ and (ii) $M - 1 < (1 - \frac{2\varepsilon}{\lambda_2})n$, and then invoke the above lemma. Note that since $n \geq M \geq 2$,

$$\varepsilon \leq \frac{\lambda_2}{6} \leq \frac{\text{Tr}[L]}{6(n-1)} = \frac{n}{6(n-1)} \leq \frac{1}{3} < \frac{2}{3} \leq 1 - \frac{2\varepsilon}{\lambda_2},$$

where the last inequality is a rearrangement of $\varepsilon \leq \frac{\lambda_2}{6}$. Thus $\varepsilon n < (1 - \frac{2\varepsilon}{\lambda_2})n$, meaning (i) implies (ii), and so it suffices to have $M \leq \varepsilon n + 1$. Overall, we use the following criteria to pick our expander graph: Given the signal dimension M , use a d -regular graph $G = (V, E)$ of n vertices with spectral gap λ_2 such that $M \leq (\frac{\lambda_2}{6})n + 1$. Then by the previous discussion, the total number of measurements is $N = |V| + 3|E| = (\frac{3}{2}d + 1)n$. If we think of the degree d as being fixed, then the number of vertices n in the graph is proportional to the total number of measurements N (this is the key distinction from the previous complete-graph case).

Recall that we seek $N = \mathcal{O}(M)$ measurements. To minimize the redundancy $\frac{N}{M}$ for a fixed degree d , we would like a maximal spectral gap λ_2 , and it suffices to seek minimal spectral expansion λ . Spectral graph families known as *Ramanujan graphs* are asymptotically optimal in this sense; taking \mathcal{G}_n^d to be the set of connected d -regular graphs with $\geq n$ vertices, Alon and Boppana (see [2]) showed that for any fixed d ,

$$\lim_{n \rightarrow \infty} \inf_{G \in \mathcal{G}_n^d} \lambda(G) \geq \frac{2\sqrt{d-1}}{d},$$

while Ramanujan graphs are defined to have spectral expansion $\leq \frac{2\sqrt{d-1}}{d}$. To date, Ramanujan graphs have only been constructed for certain values of d . One important construction was given by Lubotzky, Phillips, and Sarnak [33], which produces a Ramanujan family whenever $d - 1 \equiv 1 \pmod{4}$ is prime. Among these graphs, we get the smallest redundancy $\frac{N}{M}$ when $M = \lfloor (1 - \frac{2\sqrt{d-1}}{d})\frac{n}{6} + 1 \rfloor$ and $d = 6$:

$$\frac{N}{M} \leq \frac{(\frac{3}{2}d + 1)n}{(1 - \frac{2\sqrt{d-1}}{d})\frac{n}{6}} = 45(3 + \sqrt{5}) \approx 235.62.$$

Thus, in such cases, our techniques allow for phase retrieval with only $N \leq 236M$ measurements. However, the number of vertices in each Ramanujan graph from [33] is of the form $q(q^2 - 1)$ or $\frac{q(q^2 - 1)}{2}$, where $q \equiv 1 \pmod{4}$ is prime, and so any bound on redundancy $\frac{N}{M}$ using these graphs will only be valid for particular values of M .

In order to get $N = \mathcal{O}(M)$ in general, we use the fact that random graphs are nearly Ramanujan with high probability. In particular, for every $\varepsilon > 0$ and even d , a random d -regular graph has spectral expansion $\lambda \leq \frac{2\sqrt{d-1} + \varepsilon}{d}$ with high probability as $n \rightarrow \infty$ [22]. Thus, picking ε and d to satisfy $\frac{2\sqrt{d-1} + \varepsilon}{d} < 1$, we may take $M = \lfloor (1 - \frac{2\sqrt{d-1} + \varepsilon}{d})\frac{n}{6} + 1 \rfloor$ to get

$$\frac{N}{M} \leq \frac{(\frac{3}{2}d + 1)n}{(1 - \frac{2\sqrt{d-1} + \varepsilon}{d})\frac{n}{6}},$$

and this choice will satisfy $M \leq (\frac{\lambda_2}{6})n + 1$ with high probability. To see how small this redundancy is, note that taking $\varepsilon = 0.1$ and $d = 8$ gives $N \leq 240M$. While the desired expansion properties of a random graph are only present with high probability, estimating the spectral gap is inexpensive, and so it is computationally feasible to verify whether a randomly drawn graph is good enough. Moreover, n can be any sufficiently large integer, and so the above bound is valid for all sufficiently large M , i.e., our procedure can perform phase retrieval with $N = \mathcal{O}(M)$ measurements in general.

Combining this with the above discussion, we have the following measurement design and phase retrieval procedure:

Measurement Design A (noiseless case)

- Fix $d > 2$ even and $\varepsilon \in (0, d - 2\sqrt{d-1})$.
- Given M , pick some d -regular graph $G = (V, E)$ with spectral gap $\lambda_2 \geq \lambda' := 1 - \frac{2\sqrt{d-1} + \varepsilon}{d}$ and $|V| = \lceil \frac{6}{\lambda'}(M-1) \rceil$, and arbitrarily direct the edges.
- Design the measurements $\Phi := \Phi_V \cup \Phi_E$ by taking $\Phi_V := \{\varphi_i\}_{i \in V} \subseteq \mathbb{C}^M$ to be full spark and $\Phi_E := \bigcup_{(i,j) \in E} \{\varphi_i + \zeta^k \varphi_j\}_{k=0}^2$.

Phase Retrieval Procedure A (noiseless case)

- Given $\{|\langle x, \varphi \rangle|^2\}_{\varphi \in \Phi}$, delete the vertices $i \in V$ with $|\langle x, \varphi_i \rangle|^2 = 0$.
- In the remaining induced subgraph, find a connected component of $\geq M$ vertices V' .
- Pick a vertex in V' to have positive phase and propagate/multiply relative phases (2.6), which are calculated by normalizing (2.4), see (2.5).
- Having $\{\langle x, \varphi_i \rangle\}_{i \in V'}$ up to a global phase factor, find the least-squares estimate of $[x]$ by applying the Moore-Penrose pseudoinverse of $\{\varphi_i\}_{i \in V'}$, see (2.2).

Note that this phase retrieval procedure is particularly fast. Indeed, if we use $E \subseteq V^2$ to store G , then we can delete vertices $i \in V$ with $|\langle x, \varphi_i \rangle|^2 = 0$ by deleting the edges for which (2.4) is zero, which takes $\mathcal{O}(|E|)$ time. Next, if the members of E are ordered lexicographically, the remaining subgraph can be easily partitioned into connected components in $\mathcal{O}(|E|)$ time by collecting edges with common vertices, and then propagating relative phase in the largest component is performed in $\mathcal{O}(|E|)$ time using a depth- or breadth-first search. Overall, we only use $\mathcal{O}(M)$ time before the final least-squares step of the phase retrieval procedure, which happens to be the bottleneck, depending on the subcollection $\Phi_{V'}$. In general, we can find the least-squares estimate in $\mathcal{O}(M^3)$ time using Gaussian elimination, but if $\Phi_{V'}$ has special structure (e.g., it is a submatrix of the discrete Fourier transform matrix), then one might exploit that structure to gain speedups (e.g., use the fast Fourier transform in conjunction with an iterative method). Regardless, our procedure reduces the nonlinear phase retrieval problem to the much simpler problem of solving an overdetermined *linear* system.

While this measurement design and phase retrieval procedure is particularly efficient, it certainly lacks stability. Perhaps most notably, we have not imposed anything on Φ_V that guarantees stability with inverting $\Phi_{V'}$; indeed, we have merely enforced linear independence between vectors, while stability will require well-conditioning. Another noteworthy source of instability is our method of phase propagation, which naturally accumulates error; it would be better if the relative phases were combined using a more democratic process that encourages noise cancellation. In the next section, we will address these concerns (and others) and modify our procedure accordingly; the revised procedure will be stable, but at the price of a log factor in the number of measurements: $N = \mathcal{O}(M \log M)$. As we mention in the concluding remarks, we do not think this log factor is necessary, but we leave this pursuit for future work.

3. The noisy case. In this section, we consider a noise-robust version of the measurement design and phase retrieval procedure of the previous section. In the end, the measurement design will be nearly identical: vertex measurements will be independent complex Gaussian vectors (thereby being full spark with probability 1), and the edge measurements will be the same sort of linear combinations of vertex measurements. Our use

of randomness in this version will enable the vertex measurements to simultaneously satisfy two important conditions with high probability: *projective uniformity with noise* and *numerical erasure robustness*. Before defining these conditions, we motivate them by considering a noisy version of our phase retrieval procedure.

Recall that our noiseless procedure starts by removing the vertices $i \in V$ for which $|\langle x, \varphi_i \rangle|^2 = 0$. Indeed, since we plan to propagate relative phase information along edges, these 0-vertices are of no use, as relative phase with these vertices is not well defined. Since we calculate relative phase by normalizing (2.4), we see that relative phase is sensitive to perturbations when (2.4) is small, meaning either $\langle x, \varphi_i \rangle$ or $\langle x, \varphi_j \rangle$ is small. As such, while 0-vertices provide no relative phase information in the noiseless case, small vertices provide *unreliable* information in the noisy case, and so we wish to remove them accordingly (alternatively, one might use weights according to one's confidence in the information, but we decided to use hard thresholds to simplify the analysis). However, we also want to ensure that there are only a few small vertices. In the noiseless case, we limit the number of 0-vertices by using a full spark frame; in the noisy case, we make use of a new concept we call *projective uniformity*:

DEFINITION 3.1. *The α -projective uniformity of $\Phi = \{\varphi_i\}_{i=1}^n \subseteq \mathbb{C}^M$ is given by*

$$\text{PU}(\Phi; \alpha) = \min_{\substack{x \in \mathbb{C}^M \\ \|x\|=1}} \max_{\substack{\mathcal{I} \subseteq \{1, \dots, n\} \\ |\mathcal{I}| \geq \alpha n}} \min_{i \in \mathcal{I}} |\langle x, \varphi_i \rangle|^2.$$

In words, projective uniformity gives the following guarantee: For every unit-norm signal x , there exists a collection of vertices $\mathcal{I} \subseteq V$ of size at least $\alpha|V|$ such that $|\langle x, \varphi_i \rangle|^2 \geq \text{PU}(\Phi_V; \alpha)$ for every $i \in \mathcal{I}$. As such, projective uniformity effectively limits the total number of small vertices possible, at least before the measurements are corrupted by noise. However, the phase retrieval algorithm will only have access to noisy versions of the measurements, and so we must account for this subtlety in our procedure. In an effort to isolate the reliable pieces of relative phase information, Algorithm 1 removes the vertices corresponding to small noisy edge combinations (2.4).

Algorithm 1: Pruning for reliability

Input: Graph $G = (V, E)$, function $f: E \rightarrow \mathbb{R}$ such that $f(i, j) = |\langle x, \varphi_i \rangle \langle x, \varphi_j \rangle + \varepsilon_{ij}|$, pruning parameter α
Output: Subgraph H with a larger smallest edge weight
Initialize $H \leftarrow G$
for $i = 1$ **to** $\lfloor (1 - \alpha)|V| \rfloor$ **do**
 Find the minimizer $(i, j) \in E$ of f
 $H \leftarrow H \setminus \{i, j\}$
end

We now explain why only reliable pieces of relative phase information will remain after running the above algorithm, provided Φ_V has sufficient projective uniformity. The main idea is captured in the following:

LEMMA 3.2. *Define $\|\theta\|_{\mathbb{T}} := \min_{k \in \mathbb{Z}} |\theta - 2\pi k|$ for all angles $\theta \in \mathbb{R}/2\pi\mathbb{Z}$. Then for any $z, \varepsilon \in \mathbb{C}$,*

$$\|\arg(z + \varepsilon) - \arg(z)\|_{\mathbb{T}} \leq \pi \frac{|\varepsilon|}{|z|}.$$

Proof. If $|\varepsilon| \geq |z|$, the result is trivial. Suppose $|\varepsilon| < |z|$, and consider the triangle whose vertices are 0, z , and $z + \varepsilon$. Denoting the angle at 0 by $\theta = \|\arg(z + \varepsilon) - \arg(z)\|_{\mathbb{T}}$ and the angle at $z + \varepsilon$ by α , the law of sines gives

$$\frac{\sin \theta}{|\varepsilon|} = \frac{\sin \alpha}{|z|} \leq \frac{1}{|z|}. \tag{3.1}$$

Next, the law of cosines gives

$$2|z||z + \varepsilon| \cos \theta = |z|^2 + |z + \varepsilon|^2 - |\varepsilon|^2 > |z + \varepsilon|^2 > 0,$$

and so $\cos \theta > 0$, i.e., $\theta \in [0, \frac{\pi}{2})$. Finally, by the concavity of $\sin(\cdot)$ and then (3.1), we conclude that

$$\frac{2}{\pi} \theta \leq \sin \theta \leq \frac{|\varepsilon|}{|z|},$$

which implies the result. \square

By taking $z = \overline{\langle x, \varphi_i \rangle \langle x, \varphi_j \rangle} + \varepsilon_{ij}$ and $\varepsilon = -\varepsilon_{ij}$, we can use this lemma to bound the relative phase error we incur when normalizing z . In fact, consider the minimum of f when Algorithm 1 is complete. Since the algorithm deletes vertices from G according to the input signal x , this minimum will vary with x ; let PUN denote the smallest possible minimum value. Then the relative phase error incurred with $(i, j) \in E$ is no more than $\pi|\varepsilon_{ij}|/\text{PUN}$, regardless of the signal measured. Indeed, our use of *projective uniformity with noise* (i.e., PUN) is intended to bound the instability that comes with normalizing small values of (2.4). As we will show in the appendix, PUN can be bounded below by using the projective uniformity of Φ_V , and furthermore, a complex Gaussian Φ_V has projective uniformity with overwhelming probability.

After applying Algorithm 1, our graph will have slightly fewer vertices, but the remaining edges will correspond to reliable pieces of relative phase information. Recall that we plan to use this information on the edges to determine phases for the vertices, and we want to do this in a stable way. To understand when this is even possible, we first consider a few simple scenarios. Suppose that after removing vertices with Algorithm 1, the graph has a vertex of degree 0. Then we have no information about the phase of this vertex, and it should be removed accordingly. For a less extreme scenario, suppose the vertex has degree 1. Then any noise in the corresponding edge measurement would be passed directly to the vertex, which inherently lacks stability compared to the noise cancellation that would come with more edges. More generally, if the graph has a cut vertex (e.g., the neighbor of a degree-1 vertex), then we would need to rely on the correctness of this lone vertex to ensure consistency between the parts of the graph it connects—this scenario is also rather unstable. After considering these examples, it makes intuitive sense that stability necessitates a high level of connectivity in the graph, regardless of the algorithm used to extrapolate the vertex phases.

As such, we seek to remove a small proportion of vertices so that the remaining graph is very connected, i.e., has large spectral gap. To do this, we will iteratively remove sets of vertices that are poorly connected to the rest of the graph. These sets will be identified using *spectral clustering* (Algorithm 2), a process which is strongly motivated by an inequality in Riemannian geometry by Cheeger [15] and which has performance guarantees originating with Alon [2, 3]. The main idea of spectral clustering follows the intuition that a random walk on a graph tends to be trapped in sections of the graph which have few connections to the rest of the vertices (this intuition is made more explicit in [35, 40]). Moreover, the second eigenvector of the corresponding stochastic matrix tends to identify these sections.

Note that when implementing spectral clustering, the values of $E(S_i, S_i^c)$, $\text{vol}(S_i)$ and $\text{vol}(S_i^c)$ can be cheaply updated from their $(i - 1)$ st values since S_i differs from S_{i-1} in only one vertex. As such, the bottleneck in spectral clustering is computing an eigenvector. For our application, we will iteratively apply spectral clustering to identify small collections of vertices which are poorly connected to the rest of the graph and then remove them to enhance connectivity (Algorithm 3).

In the appendix, we show that for a particular choice of threshold τ , Algorithm 3 recovers a level of connectivity that may have been lost when pruning for reliability in Algorithm 1, and it does so by removing only a small proportion of the vertices.

At this point, we have pruned our graph so that the measured relative phases are reliable and the vertex phases can be stably reconstructed. Now we seek an efficient method to reconstruct these vertex phases from the measured relative phases. Before devising such a method, we first organize the information we have into a matrix. Given the graph output $G' = (V', E')$ of Algorithm 3, we take A_1 to be a $|V'| \times |V'|$ weighted

Algorithm 2: Spectral clustering

Input: Graph $G = (V, E)$

Output: Subset of vertices S

Take D to be the diagonal matrix of vertex degrees

Compute the Laplacian $L \leftarrow I - D^{-1/2}AD^{-1/2}$

Compute the eigenvector u corresponding to the second eigenvalue of L

for $i = 1$ **to** $|V|$ **do**

 Let S_i denote the vertices corresponding to the i smallest entries of $D^{-1/2}u$

 Let $E(S_i, S_i^c)$ denote the number of edges between S_i and S_i^c

$h_i \leftarrow E(S_i, S_i^c) / \min\{\sum_{v \in S_i} \deg(v), \sum_{v \in S_i^c} \deg(v)\}$

end

Take S to be the S_i of minimal h_i (or S_i^c if this has smaller size)

Algorithm 3: Pruning for connectivity

Input: Graph $G = (V, E)$, pruning parameter τ

Output: Subgraph H with spectral gap $\lambda_2(H) \geq \tau$

Initialize $H \leftarrow G$

while $\lambda_2(H) < \tau$ **do**

 Perform spectral clustering (Algorithm 2) to identify a small set of vertices S

$H \leftarrow H \setminus S$

end

adjacency matrix of G' . Specifically, for each $\{i, j\} \in E'$, let ε_{ij} denote the effective noise in the estimate of $\langle x, \varphi_i \rangle \langle x, \varphi_j \rangle$ using (2.4), and normalize this noisy estimate to get

$$A_1[i, j] = \frac{\langle x, \varphi_i \rangle \langle x, \varphi_j \rangle + \varepsilon_{ij}}{|\langle x, \varphi_i \rangle \langle x, \varphi_j \rangle + \varepsilon_{ij}|}. \quad (3.2)$$

Otherwise when $\{i, j\} \notin E'$, take $A_1[i, j] = 0$. Unlike the noiseless case, here, we account for both directions (i, j) and (j, i) whenever $\{i, j\} \in E'$, with the understanding that $\varepsilon_{ji} = \overline{\varepsilon_{ij}}$; this will simplify our analysis since this makes A_1 self-adjoint. Considering $A_1[i, j]$ is an approximation of the relative phase $\omega_i^{-1}\omega_j$, it seems reasonable to extrapolate the vertex phases $\omega := \{\omega_i\}_{i \in V'}$ from A_1 by minimizing the following quantity:

$$\sum_{\{i, j\} \in E'} |\omega_j - A_1[i, j]\omega_i|^2 = \sum_{\{i, j\} \in E'} \left(|\omega_j|^2 - 2 \operatorname{Re} \overline{\omega_j} A_1[i, j]\omega_i + |A_1[i, j]\omega_i|^2 \right) = \omega^*(D - A_1)\omega,$$

where D is the diagonal matrix of vertex degrees. Dividing by $\operatorname{vol}(G') = \omega^*D\omega$, which does not vary with ω , this is equivalent to minimizing

$$\frac{\omega^*(D - A_1)\omega}{\omega^*D\omega} = \frac{(D^{1/2}\omega)^*(I - D^{-1/2}A_1D^{-1/2})(D^{1/2}\omega)}{\|D^{1/2}\omega\|^2} \geq \lambda_1(I - D^{-1/2}A_1D^{-1/2}).$$

To be clear, the right-hand side above is the first eigenvalue of $L_1 := I - D^{-1/2}A_1D^{-1/2}$, which we call the *connection Laplacian*; note that this bears some resemblance to the Laplacian defined in the previous section. In minimizing the above quantity, it makes sense to consider the eigenvector u corresponding to the smallest eigenvalue of L_1 , but we require each coordinate of $D^{-1/2}u$ to have unit modulus. Provided u has no entries which are zero, we can normalize the entries to form an estimate of ω , and as we show in the appendix (using results from [9]), this estimate is stable provided the spectral gap of G' is sufficiently large. This spectral method is known in the literature as *angular synchronization* [46], and we summarize the procedure in Algorithm 4

Algorithm 4: Angular synchronization

Input: Graph $G' = (V', E')$, noisy versions of (2.4) for every $\{i, j\} \in E'$

Output: Vector of phases corresponding to vertex measurements

Let A_1 denote the matrix given by (3.2) whenever $\{i, j\} \in E'$, and otherwise $A_1[i, j] = 0$

Let D denote the diagonal matrix of vertex degrees

Compute the connection Laplacian $L_1 \leftarrow I - D^{-1/2}A_1D^{-1/2}$

Compute the eigenvector u corresponding to the smallest eigenvalue of L_1

Output the phases of the coordinates of u

To reiterate, Algorithm 4 will produce estimates for the phases of the inner products $\{\langle x, \varphi_i \rangle\}_{i \in V'}$. Also, we can take square roots of the vertex measurements $\{|\langle x, \varphi_i \rangle|^2 + \nu_i\}_{i \in V'}$ to estimate $\{|\langle x, \varphi_i \rangle|\}_{i \in V'}$. Then we can combine these to estimate $\{\langle x, \varphi_i \rangle\}_{i \in V'}$. However, note that the largest of these inner products will be most susceptible to noise in the corresponding phase estimate. As such, we remove a small fraction of these largest vertices so that the final collection of vertices V'' has size $\kappa|V|$, where V was the original vertex set, and κ is sufficiently close to 1.

Now that we have estimated the phases of $\{\langle x, \varphi_i \rangle\}_{i \in V''}$, we wish to reconstruct x by applying the Moore-Penrose pseudoinverse of $\{\varphi_i\}_{i \in V''}$. However, since V'' is likely a strict subset of V , it can be difficult in general to predict how stable the pseudoinverse will be. Fortunately, a recent theory of *numerically erasure-robust frames (NERFs)* makes this prediction possible: If the members of Φ_V are independent Gaussian vectors, then with high probability, every submatrix of columns $\Phi_{V''}$ with $\kappa = |V''|/|V|$ sufficiently large has a stable pseudoinverse [20]. This concludes the phase retrieval procedure, briefly outlined below together with the measurement design.

Measurement Design B (noisy case)

- Fix $d > 2$ even and $\varepsilon \in (0, d - 2\sqrt{d-1})$.
- Given M , pick some d -regular graph $G = (V, E)$ with spectral gap $\lambda_2 \geq \lambda' := 1 - \frac{2\sqrt{d-1} + \varepsilon}{d}$ and $|V| = cM \log M$ for c sufficiently large, and arbitrarily direct the edges.
- Design the measurements $\Phi := \Phi_V \cup \Phi_E$ by taking $\Phi_V := \{\varphi_i\}_{i \in V} \subseteq \mathbb{C}^M$ to have independent entries with distribution $\mathcal{CN}(0, \frac{1}{M})$ and $\Phi_E := \bigcup_{(i,j) \in E} \{\varphi_i + \zeta^k \varphi_j\}_{k=0}^2$.

Phase Retrieval Procedure B (noisy case)

- Given $\{|\langle x, \varphi_\ell \rangle|^2 + \nu_\ell\}_{\ell=1}^N$, prune the graph G , keeping only reliable vertices (Algorithm 1).
- Prune the remaining induced subgraph for connectivity, producing the vertex set V' (Algorithm 3).
- Estimate the phases of the vertex measurements using angular synchronization (Algorithm 4).
- Remove the vertices with the largest measurements, keeping only $|V''| = \kappa|V|$.
- Having estimates for $\{\langle x, \varphi_i \rangle\}_{i \in V''}$ up to a global phase factor, find the least-squares estimate of $[x]$ by applying the Moore-Penrose pseudoinverse of $\{\varphi_i\}_{i \in V''}$, see (2.2).

Having established our measurement design and phase retrieval procedure for the noisy case, we now present the following guarantee of stable performance:

THEOREM 3.3. *Pick $N \sim CM \log M$ with C sufficiently large, and take $\{\varphi_\ell\}_{\ell=1}^N = \Phi_V \cup \Phi_E$ defined in Measurement Design B. Then there exist constants $C', K > 0$ such that the following guarantee holds for all $x \in \mathbb{C}^M$ with overwhelming probability: Consider measurements of the form*

$$z_\ell := |\langle x, \varphi_\ell \rangle|^2 + \nu_\ell.$$

If the noise-to-signal ratio satisfies $\text{NSR} := \frac{\|\nu\|}{\|x\|^2} \leq \frac{C'}{\sqrt{M}}$, then Phase Retrieval Procedure B produces an

estimate \tilde{x} from $\{z_\ell\}_{\ell=1}^N$ with squared relative error

$$\frac{\|\tilde{x} - e^{i\theta}x\|^2}{\|x\|^2} \leq K \sqrt{\frac{M}{\log M}} \text{NSR}$$

for some phase $\theta \in [0, 2\pi)$.

The interested reader is directed to the appendix for a proof of this guarantee. Before concluding this section, we evaluate the result. Note that the norms of the φ_ℓ 's tend to be $\mathcal{O}(1)$, and so the noiseless measurements $|\langle x, \varphi_\ell \rangle|^2$ tend to be of size $\mathcal{O}(\|x\|^2/M)$. Also, in the worst-case scenario, the noise annihilates our measurements $\nu_\ell = -|\langle x, \varphi_\ell \rangle|^2$, rendering the signal x unrecoverable; in this case, $\|\nu\| = \mathcal{O}(\|x\|^2 \sqrt{(\log M)/M})$ since $N = CM \log M$. In other words, if we allowed the noise-to-signal ratio to scale slightly larger than C'/\sqrt{M} (i.e., by a log factor), then it would be impossible to perform phase retrieval in the worst case. As such, the above guarantee is optimal in some sense. Furthermore, since $\sqrt{M/\log M} \text{NSR} = \mathcal{O}(1/\sqrt{\log M})$ by assumption, the result indicates that our phase retrieval process exhibits more stability as M grows large.

In comparison to the state of the art, namely, the work of Candès, Strohmer and Voroninski [14], the most visible difference between our stability results is how we choose to scale the measurement vectors. Indeed, the measurement vectors of the present paper tend to have norm $\mathcal{O}(1)$, whereas the measurement vectors of Candès et al. are all scaled to have norm \sqrt{M} ; considering the statement of Theorem 3.3 is riddled with square roots of M , either choice of scaling is arguably natural. For the sake of a more substantial comparison, their result (Theorem 1.2) gives that if the φ_ℓ 's are uniformly sampled from the sphere of radius \sqrt{M} , and if $\|\nu\| \leq \epsilon$, then with high probability, there exists $C_0 > 0$ such that

$$\frac{\|\tilde{x} - e^{i\theta}x\|}{\|x\|} \leq C_0 \min \left\{ 1, \frac{\epsilon}{\|x\|^2} \right\}$$

for some phase $\theta \in [0, 2\pi)$; here, \tilde{x} is the estimate which comes from PhaseLift. (This result actually suffers from the subtlety that it does not hold for all signals $x \in \mathbb{C}^M$ simultaneously, but this was later rectified in a sequel [13].) Note that the 1 in the minimum takes effect when $\|x\|^2 < \epsilon$, meaning it is possible that $\nu_\ell = -|\langle x, \varphi_\ell \rangle|^2$ (corresponding to our worst-case scenario, above); as such, this part of the guarantee is not as interesting. The other part of the guarantee is particularly interesting: Ignoring the $\sqrt{M/\log M}$ factor in Theorem 3.3 and identifying $\epsilon/\|x\|^2$ with NSR, we see that the main difference between the two guarantees is that Candès et al. bound relative error in terms of NSR rather than bounding *squared* relative error. In this sense, their result is stronger. On the other hand, they take ϵ as an input to their reconstruction algorithm (PhaseLift), whereas our method is agnostic to the size of ν . In particular, our guarantee is continuous in the sense that relative error vanishes with ν .

4. Numerical results. In the previous sections, we described measurement designs and phase retrieval procedures for both the noiseless and noisy cases. This section presents results from numerical simulations to illustrate how well our phase retrieval procedures perform in practice. In particular, we will consider the noiseless and noisy cases separately.

4.1. The noiseless case. For this case, we consider a slightly different measurement design. Rather than drawing a d -regular graph of n vertices at random, we instead draw an Erdős-Rényi random graph. That is, for a fixed n and $c \leq n$, we take n vertices, and place an edge between each pair of vertices independently with probability $p := c/n$. Note that for this model, the mean degree of each vertex is $p(n-1) = c(n-1)/n \approx c$. As we will see, this slight change to the graph model will not adversely affect the quality of our methods.

In the simulation, we considered three values for the dimension $M \in \{16, 32, 64\}$. For each of these dimensions, we let n range from M to $4M$ in increments of $M/4$. In MATLAB® parlance, denoting $r := n/M$, we took $r = 1 : 0.25 : 4$. We also took $c = 1 : 0.5 : 2$. For each (r, c) pair, we constructed an Erdős-Rényi random graph $G = (V, E)$ with $n = rM$ and $p = c/n$. We then picked each member of Φ_V to have independent

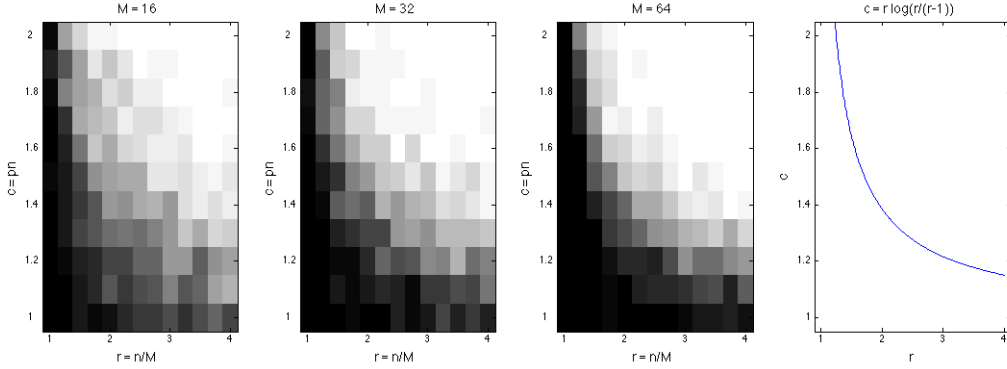


FIG. 4.1. The first three plots indicate the proportion of successful estimates of a random vector using Phase Retrieval Procedure A; the spatial dimensions for these plots are $M = 16, 32$ and 64 , respectively. Instead of using a random d -regular graph (as prescribed in Measurement Design A), the graphs for this simulation were Erdős-Rényi random graphs with $n = rM$ vertices and edge probability $p = c/n$. The horizontal axes take $r = 1 : 0.25 : 4$, while the vertical axes take $c = 1 : 0.5 : 2$. Darker cells indicate a smaller proportion of successful estimates for that particular choice of (r, c) . The plots appear to be approaching a phase transition, which is illustrated in the fourth plot; here, the curve is given by $c = r \log(r/(r-1))$, and being above or below this curve dictates whether the Erdős-Rényi random graph has a component of size at least M with high probability. Having this large component in the graph essentially determines whether the random vector will be successfully reconstructed.

$\mathcal{CN}(0, \frac{1}{M})$ entries, and we constructed $\Phi_E := \bigcup_{(i,j) \in E} \{\varphi_i + \zeta^k \varphi_j\}_{k=0}^2$. At this point, we drew a vector x with independent $\mathcal{CN}(0, \frac{1}{M})$ entries, and performed Phase Retrieval Procedure A to find an estimate \tilde{x} of x . To be clear, in implementing the first step of Phase Retrieval Procedure A, we set a tolerance for deletion: Given $\{|\langle x, \varphi \rangle|^2\}_{\varphi \in \Phi}$, we actually deleted vertices $i \in V$ with $|\langle x, \varphi_i \rangle|^2 < 10^{-4}$. After running the remainder of Phase Retrieval Procedure A, we considered the estimate to be “good” if the relative error was small:

$$\min_{\theta \in [0, 2\pi)} \frac{\|\tilde{x} - e^{i\theta} x\|}{\|x\|} < 10^{-5}.$$

For each (r, d) pair, we performed 30 trials of this process, and we populated the corresponding cell in Figure 4.1 with a shade of gray according to the proportion of good estimates (black indicates that none of the estimates were good, while white indicates that all of the estimates were good).

In observing the results of these simulations in Figure 4.1, we see that the proportion of good estimates appears to go to either 0 or 1 depending on whether the pair (r, d) lies above or below the curve given in the fourth plot. To explain this apparent phase transition, we first note that in the course of running Phase Retrieval Procedure A for the 12870 total trials, we only had 4 instances in which a vertex i was deleted due to $|\langle x, \varphi_i \rangle|^2 < 10^{-4}$. This is not terribly surprising, considering x and φ_i were both drawn to have independent $\mathcal{CN}(0, \frac{1}{M})$ entries, and therefore do not tend to be almost exactly orthogonal. What this suggests is that the dominant feature which determines whether reconstruction will be successful is whether the Erdős-Rényi random graph has a connected component of size at least M (the remainder of Phase Retrieval Procedure A accumulates only round-off errors, which are negligible here). At this point, we appeal to a result which appears in Section 5.2 of [30], namely that if $c > 1$ is held constant, then with high probability, the Erdős-Rényi random graph with n vertices and edge probability $p = c/n$ has a unique giant connected component of size $(\beta + o(1))n$, where β is the unique solution to the equation $\beta + e^{-\beta c} = 1$. Since in our case, we succeed if the giant component to has size $(\beta + o(1))n \geq M = \frac{n}{r}$, and we fail otherwise, the phase transition should correspond to $\beta = \frac{1}{r}$. Rearranging β 's defining equation then gives the formula for the curve in Figure 4.1, namely $c = r \log(r/(r-1))$.

Having established this phase transition, we can use it to minimize the number of measurements. In particular, the total number of edges in the graph tends to be around $\frac{nc}{2}$, and so the total number of measurements

is $N \approx (1 + \frac{3}{2}c)n$. As before, we seek to minimize redundancy:

$$\frac{N}{M} = \frac{rN}{n} \approx r \left(1 + \frac{3}{2}c \right) = r \left(1 + \frac{3}{2}r \log \frac{r}{r-1} \right).$$

The right-hand side above is minimized when $r \approx 1.28$, in which case we get a redundancy of $\frac{N}{M} \approx 5.02 \ll 236$. The reason for this disparity is simple: In the previous expander-graph-based analysis, we were chiefly concerned with ensuring that our measurement vectors Φ lend injective intensity measurements, so that we could reconstruct *any* given signal. On the other hand, the above analysis demonstrates that we can get away with far fewer measurement vectors if we only need to be able to reconstruct *almost every* signal. In this sense, these numerical simulations fail to capture the most challenging feature of measurement design for phase retrieval: injectivity (versus unique representation of almost every signal).

4.2. The noisy case. In this subsection, we use simulations to demonstrate the stability of Phase Retrieval Procedure B. Here, we perform two types of simulations. In the first simulation, we consider a slightly different noise model so as to compare with different “phase oracles.” To be clear, rather than considering noisy intensity measurements of the form $|\langle x, \varphi_\ell \rangle|^2 + \nu_\ell$, we instead add the noise *before* the modulus squared: $|\langle x, \varphi_\ell \rangle + \nu_\ell|^2$. With this change of noise model, we can compare the performance of Phase Retrieval Procedure B to (i) least-squares estimation from $\Phi_V^*x + \nu_V$ and (ii) least-squares estimation from $\Phi^*x + \nu$. This is particularly interesting because the bulk of Phase Retrieval Procedure B is dedicated to reconstructing (most of) the phases in $\Phi_V^*x + \nu_V$ which are lost in the modulus squared, and so the comparison with (i) will help us evaluate this portion of the procedure. Also, (ii) can be viewed as a theoretical limit of sorts, and so this comparison will indicate how much stability we lose with the phases.

For the simulations, we again used the Erdős-Rényi random graph model instead of a random d -regular graph. Specifically, since 8-regular random graphs produced measurements with minimum redundancy for Measurement Design A, and since the corresponding ratio was $r := \frac{n}{M} \approx 3$, we decided to take $r = 3$ and $c = 8$ for these simulations. For each $M = 8 : 4 : 128$, we picked each of the $n = 3M$ members of Φ_V to have independent $\mathcal{CN}(0, \frac{1}{M})$ entries, and we constructed $\Phi_E := \bigcup_{(i,j) \in E} \{\varphi_i + \zeta^k \varphi_j\}_{k=0}^2$ in accordance with an Erdős-Rényi random graph of edge probability $p = \frac{c}{n} = \frac{8}{3M}$. Next, we picked the signal x to have independent $\mathcal{CN}(0, \frac{1}{M})$ entries and the noise vector ν to have independent $\mathcal{CN}(0, \frac{\sigma^2}{M})$ entries with $\sigma = 0.4$. We then applied Phase Retrieval Procedure B to $\{|\langle x, \varphi_\ell \rangle + \nu_\ell|^2\}_{\ell=1}^N$ (again, instead of $\{|\langle x, \varphi_\ell \rangle|^2 + \nu_\ell\}_{\ell=1}^N$). Specifically, we took $\alpha = 0.9925$ and $\tau = 0.1$ as our pruning parameters. We also performed least-squares estimation from $\Phi_V^*x + \nu_V$ and from $\Phi^*x + \nu$. We recorded the relative error and the computation time for each of these three estimates; see Figure 4.2.

Interestingly, Phase Retrieval Procedure B produces relative errors similar to those gotten by doing least-squares estimation from $\Phi_V^*x + \nu_V$. This suggests that the portion of Phase Retrieval Procedure B which reconstructs (most of) the phases in $\Phi_V^*x + \nu_V$ performs rather well. In fact, this shows that the polarization trick of using edges to estimate vertex phases is particularly successful. Notice that performing least-squares estimation from $\Phi^*x + \nu$ produces a much smaller relative error, as expected. As far as runtime is concerned, Phase Retrieval Procedure B is slower than the phase oracles because it takes some time to estimate vertex phases with angular synchronization; however, this is not a substantial difference in runtime, as our procedure still produces an estimate in less than one second.

We would like to point out that in pruning for connectivity, instead of directly applying Algorithm 3, it sufficed to find the largest surviving component, since in our trials, this component always had spectral gap larger than $\tau = 0.1$. We suspect that this is an artifact of the random graph, as this will certainly not happen in general.

The second simulation we ran to demonstrate stability was a comparison with the standard alternating projections routine: After assigning arbitrary (say, positive) phases to the intensity measurements, iteratively project onto the column space of Φ^* and then project onto the nonconvex set of vectors y whose entry magnitudes have the given intensity measurements. We applied both algorithms to $\{|\langle x, \varphi_\ell \rangle|^2 + \nu_\ell\}_{\ell=1}^N$,

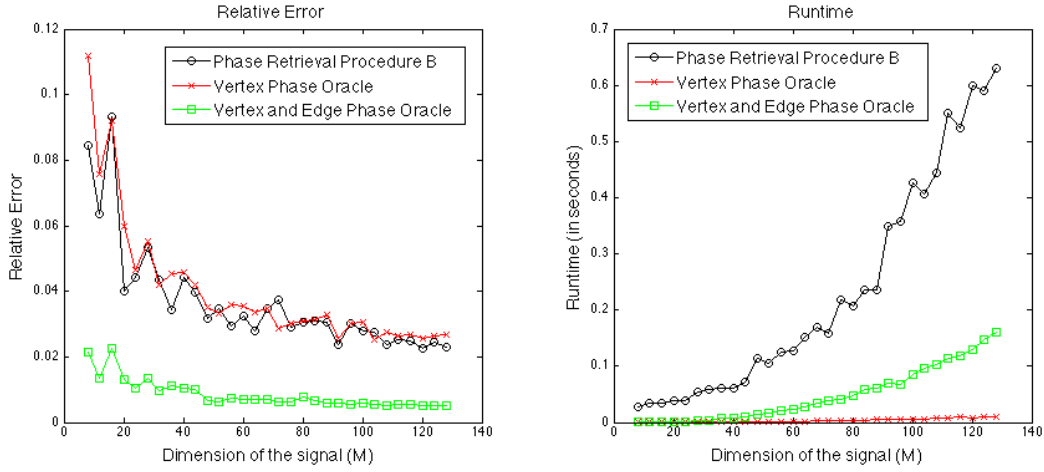


FIG. 4.2. Relative error and runtime (in seconds) for three different estimates. In both plots, the horizontal axis takes $M = 8 : 4 : 128$. The black curve denotes Phase Retrieval Procedure B applied to $\{|\langle x, \varphi_\ell \rangle + \nu_\ell|^2\}_{\ell=1}^N$, where ν has independent complex-Gaussian entries. The noise here is added before the modulus squared so that we can compare with two “phase oracles.” Specifically, the red curve denotes least-squares estimation from $\Phi_V^* x + \nu_V$, and the green curve denotes least-squares estimation from $\Phi^* x + \nu$. Overall, Phase Retrieval Procedure B does a good job of reconstructing (most of) the vertex phases (since the black and red curves are similar in the relative error plot), but this comes with a modest price in runtime.

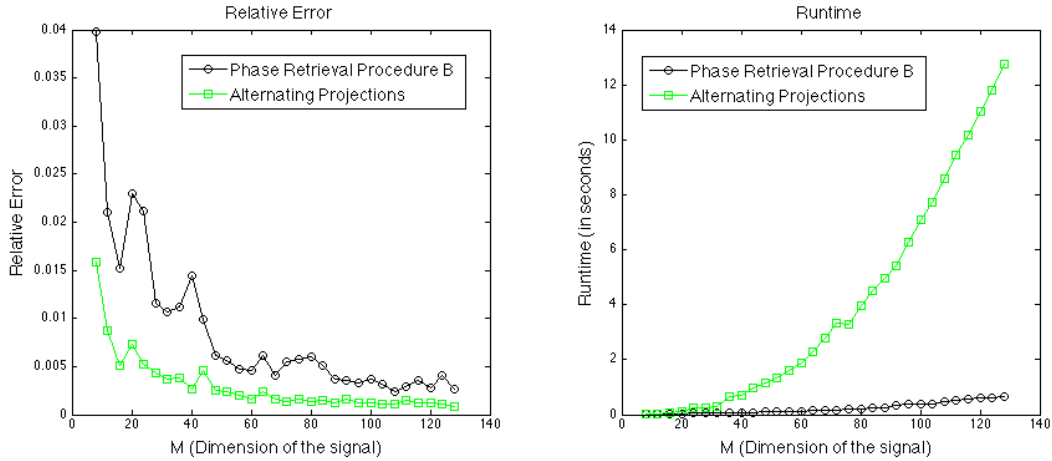


FIG. 4.3. Relative error and runtime (in seconds) for two different estimates. In both plots, the horizontal axis takes $M = 8 : 4 : 128$. The black curve denotes Phase Retrieval Procedure B applied to $\{|\langle x, \varphi_\ell \rangle|^2 + \nu_\ell\}_{\ell=1}^N$, where ν has independent Gaussian entries. The green curve denotes alternating projections applied to the same noisy intensity measurements. Overall, alternating projections does surprisingly well at doing phase retrieval in these instances, though requiring a bit more runtime. Note that the relative error of Phase Retrieval Procedure B is different here than in Figure 4.2 because in this case, the noise is added after the modulus squared.

where the ν_ℓ 's are independent $\mathcal{N}(0, \frac{\sigma^2}{M})$ entries with $\sigma = 0.4$. Other than the different phase retrieval algorithm and noise model, our simulation protocol was identical to the previous one, and the results are presented in Figure 4.3.

The most striking thing about this simulation is that alternating projections consistently produces a slightly better estimate than Phase Retrieval Procedure B. For comparison, we also ran alternating projections using only the $3M$ vertex measurement vectors Φ_V , and the relative errors were consistently on the order of 1, i.e., alternating projections consistently stalled in this case. This suggests that there is some fundamental

quality about the polarized measurement vectors Φ which makes them particularly well-suited for alternating projections, and we intend to study this in the future. Regardless, alternating projections took a lot longer to terminate; to be clear, we terminated the loop once applying both projections moved the estimate by less than 10^{-3} , or by the 100th iteration, whichever occurred first.

5. Concluding remarks. This paper provides a new way to perform phase retrieval, and our main result (Theorem 3.3) shows that our method is stable. In comparing with the stability result of [14], we note that neither result is completely satisfying when viewed from the perspective of application: In the real world, you are given a noise level and an acceptable level of estimate error, and you are asked to meet these specifications with signal processing techniques. For phase retrieval, the available guarantees fail to prescribe a measurement design that overcomes a given noise level—rather, they merely establish that with sufficiently many measurements, there exists some level of stability, i.e., K in Theorem 3.3 or C_0 in Theorem 1.2 of [14]. This reveals a gap in what is known about stability in phase retrieval, and we leave this for future work.

Admittedly, there are several gaps remaining between modern theory and application of phase retrieval. For example, throughout this paper, it is assumed that the user has complete knowledge of the measurement design, but this is not always possible in practice. This can be resolved in part with new stability results which account for “noise” in the measurement design (this is sometimes called mismatch error).

One might feel that our phase retrieval algorithms are slightly unsatisfying because we perform hard thresholds to remove vertices according to how small or large the corresponding measurements are. Alternatively, there could very well be a way to more smoothly weight these measurements according to our confidence in them, and such weightings are already accounted for in the theory of angular synchronization [9]. However, we decided to use hard thresholds because they greatly simplify the analysis of projective uniformity (though the analysis is still rather technical).

While the worst-case analysis we provide here is useful in many applications (and enables a comparison with the worst-case stability results of [14]), stochastic noise is a more appropriate model in other applications. We believe that the phase retrieval procedure of this paper will perform substantially better in the average case, but we leave this analysis for future work. Also, a notable distinction between our measurement designs in the noiseless and noisy cases is the presence of a log factor in the number of measurements used. However, we believe this factor is an artifact of our current analysis, and we intend to remove it in the future.

6. Appendix.

6.1. Graph pruning.

This section proves the following guarantee:

THEOREM 6.1. *Take proportions $p \geq q \geq \frac{2}{3}$, and consider a regular graph $G = (V, E)$ with spectral gap $\lambda_2 > g(p, q) := 1 - 2(q(1 - q) - (1 - p))$. After Algorithm 1 removes at most $(1 - p)|V|$ vertices from G , then setting $\tau = \frac{1}{8}(\lambda_2 - g(p, q))^2$, Algorithm 3 outputs a subgraph with at least $q|V|$ vertices.*

To prove this theorem, we will apply a graph version of the Cheeger inequality, which provides a guarantee for Algorithm 2:

THEOREM 6.2 (Constructive Cheeger inequality [16]). *Consider a graph $G = (V, E)$ with spectral gap λ_2 . Then Algorithm 2 outputs a set of vertices S such that $h(S) \leq \sqrt{2\lambda_2}$.*

Proof. [Proof of Theorem 6.1] First, Algorithm 1 removes a set of vertices, which we denote by S_0 . In applying Algorithm 3, the i th step of the while loop removes another set of vertices S_i . We claim this while loop will end with $|\bigcup_{i \geq 0} S_i| < (1 - q)|V|$. Supposing to the contrary, consider the first k for which $S := \bigcup_{i=0}^k S_i$ has at least $(1 - q)|V|$ vertices. Then since each iteration of the while loop removes at most half of the remaining vertices, we have $(1 - q)|V| \leq |S| \leq (1 - \frac{q}{2})|V|$.

To derive a contradiction, we will find incompatible upper and lower bounds on $E(S, S^c)$. For the upper

bound, we apply the fact that G is d -regular along with the definition of $h(S_i)$ to get

$$E(S, S^c) = E(S_0, S^c) + \sum_{i=1}^k E(S_i, S^c) \leq \text{vol}(S_0) + \sum_{i=1}^k E(S_i, S_i^c) \leq d|S_0| + \sum_{i=1}^k h(S_i) \text{vol}(S_i).$$

Next, Theorem 6.2 bounds each $h(S_i)$ in terms of the spectral gap the remaining graph, which is necessarily less than τ by the condition of the while loop. Thus, we continue:

$$E(S, S^c) < d|S_0| + \sum_{i=1}^k \sqrt{2\tau}d|S_i| \leq d(1-p)|V| + \sqrt{2\tau}d|S| \leq d|V| \left((1-p) + \sqrt{2\tau} \left(1 - \frac{q}{2}\right) \right). \quad (6.1)$$

For the lower bound, we use the expander mixing lemma, which says

$$\left| E(S, S^c) - \frac{d}{|V|} |S||S^c| \right| \leq d(1-\lambda_2) \sqrt{|S||S^c|}.$$

Since the function $x \mapsto x(1-x)$ is concave down, we know the minimum subject to $1-q \leq x \leq 1 - \frac{q}{2}$ is achieved at an endpoint of this interval. Thus, evaluating the function at $|S|/|V|$ gives $|S||S^c|/|V|^2 \geq \min\{q(1-q), \frac{q}{2}(1 - \frac{q}{2})\} = q(1-q)$, where the last step follows from the fact that $q \geq \frac{2}{3}$. Applying this and $|S||S^c| \leq |V|^2/4$ to the expander mixing lemma then gives

$$E(S, S^c) \geq \frac{d}{|V|} |S||S^c| - d(1-\lambda_2) \sqrt{|S||S^c|} \geq d|V|q(1-q) - d(1-\lambda_2) \frac{|V|}{2} = d|V| \left(q(1-q) - \frac{1-\lambda_2}{2} \right). \quad (6.2)$$

Finally, we combine (6.1) and (6.2) to get

$$q(1-q) - \frac{1-\lambda_2}{2} < (1-p) + \sqrt{2\tau} \left(1 - \frac{q}{2}\right),$$

which, as substitution reveals, contradicts our choice for τ . \square

6.2. Angular synchronization.

This section proves the following guarantee:

THEOREM 6.3. *Consider a graph $G = (V, E)$ with spectral gap $\tau > 0$, and define $\|\theta\|_{\mathbb{T}} := \min_{k \in \mathbb{Z}} |\theta - 2\pi k|$ for all angles $\theta \in \mathbb{R}/2\pi\mathbb{Z}$. Given the weighted adjacency matrix A_1 (3.2), then Algorithm 4 outputs $u \in \mathbb{C}^{|V|}$ with unit-modulus entries such that, for some phase $\theta \in \mathbb{R}/2\pi\mathbb{Z}$,*

$$\sum_{i \in V} \left\| \arg(u_i) - \arg(\langle x, \varphi_i \rangle) - \theta \right\|_{\mathbb{T}}^2 \leq \frac{C \|\varepsilon\|^2}{\tau^2 P^2},$$

where $P := \min_{\{i,j\} \in E} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle + \varepsilon_{ij}|$ and C is a universal constant.

In terms of phases, we are interested in reconstructing $\omega^* : V \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ such that $\omega_i^* = \arg(\langle x, \varphi_i \rangle)$ for every $i \in V$. For some graph $G = (V, E)$, we are given a weighted adjacency matrix A_1 (3.2), which encodes edge measurements:

$$\rho_{ij} := \omega_j^* - \omega_i^* + e_{ij}$$

for every $(i, j) \in E$; here, $e_{ij} \in \mathbb{R}/2\pi\mathbb{Z}$ is angular noise, and we note that $e_{ji} = -e_{ij}$. We measure the size of $e = \{e_{ij}\}_{\{i,j\} \in E}$ in terms of its components: $\|e\|_{\mathbb{T}}^2 := \sum_{\{i,j\} \in E} \|e_{ij}\|_{\mathbb{T}}^2$.

We will prove Theorem 6.3 using a recent Cheeger inequality for the connection Laplacian [9]. In effect, this result provides a guarantee for Algorithm 4 in terms of a certain objective function. To be precise, given $\omega : V \rightarrow \mathbb{R}/2\pi\mathbb{Z}$, we define

$$\eta(\omega) = \sum_{\{i,j\} \in E} |e^{i((\omega_j - \omega_i) - \rho_{ij})} - 1|^2.$$

Note that one way to reconstruct ω^* is to minimize this quantity. Indeed, $\eta(\omega)$ is small when the angular differences $\omega_j - \omega_i$ are close to the measured differences ρ_{ij} , and in the noiseless case, $\eta(\omega) = 0$ precisely when $\omega = \omega^*$, provided the graph G is connected. In terms of this objective function, the following guarantee ensures that the output of Algorithm 4 is no worse than a constant multiple of optimal:

THEOREM 6.4 (Cheeger inequality for the connection Laplacian [9]). *Consider a graph $G = (V, E)$ with spectral gap $\tau > 0$. Given the weighted adjacency matrix A_1 (3.2), which encodes edge measurements $\rho : E \rightarrow \mathbb{R}/2\pi\mathbb{Z}$, then Algorithm 4 outputs $u \in \mathbb{C}^{|V|}$, which encodes $\hat{\omega} : V \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ such that*

$$\eta(\hat{\omega}) \leq \frac{C'}{\tau} \min_{\omega: V \rightarrow \mathbb{T}} \eta(\omega), \quad (6.3)$$

where C' is a universal constant.

We start with two lemmas:

LEMMA 6.5. *For every $a, b \in \mathbb{R}/2\pi\mathbb{Z}$, we have $\frac{1}{2}\|a\|_{\mathbb{T}}^2 - \|b\|_{\mathbb{T}}^2 \leq \|a - b\|_{\mathbb{T}}^2$.*

Proof. By abuse of notation, we identify a and b with their coset representatives in $[-\pi, \pi)$. Since

$$0 \leq \frac{1}{2}(a - 2b)^2 = (a - b)^2 - (\frac{1}{2}a^2 - b^2),$$

then rearranging gives the following inequality:

$$\frac{1}{2}\|a\|_{\mathbb{T}}^2 - \|b\|_{\mathbb{T}}^2 = \frac{1}{2}a^2 - b^2 \leq (a - b)^2 = \|a - b\|_{\mathbb{T}}^2,$$

where the last equality requires $a - b \in [-\pi, \pi]$. Otherwise, we note that

$$0 \leq \frac{1}{2}(a - 2b \pm 2\pi)^2 + 2\pi(\pi \pm a) = (a - b \pm 2\pi)^2 - (\frac{1}{2}a^2 - b^2),$$

and so rearranging gives

$$\frac{1}{2}\|a\|_{\mathbb{T}}^2 - \|b\|_{\mathbb{T}}^2 = \frac{1}{2}a^2 - b^2 \leq (a - b \pm 2\pi)^2,$$

one of which equals $\|a - b\|_{\mathbb{T}}^2$ depending on whether $a - b < -\pi$ or $a - b > \pi$. \square

LEMMA 6.6. *For every $a, b \in \mathbb{C}$ with $|a| = 1$, we have $|a - \frac{b}{|b|}| \leq 2|a - b|$.*

Proof. The reverse triangle inequality gives $|\frac{b}{|b|} - b| = |1 - |b|| = ||a| - |b|| \leq |a - b|$, and so the triangle inequality gives $|a - \frac{b}{|b|}| \leq |a - b| + |b - \frac{b}{|b|}| \leq 2|a - b|$. \square

Proof. [Proof of Theorem 6.3] Identifying $\theta \in \mathbb{R}/2\pi\mathbb{Z}$ with its coset representative in $[-\pi, \pi)$, we have $\|\theta\|_{\mathbb{T}} = |\theta|$, and a double-angle formula gives $|e^{i\theta} - 1| = 2 \sin \frac{|\theta|}{2}$. From these, it follows that

$$\frac{2}{\pi}\|\theta\|_{\mathbb{T}} \leq |e^{i\theta} - 1| \leq \|\theta\|_{\mathbb{T}}. \quad (6.4)$$

This relationship will allow us to apply Theorem 6.4. To this end, for notational convenience, we define $\gamma_i := \arg(u_i) - \arg(\langle x, \varphi_i \rangle)$ for every $i \in V$. The right-hand inequality of (6.4) and Lemma 6.5 together give

$$\frac{1}{2} \sum_{\{i,j\} \in E} |e^{i(\gamma_j - \gamma_i)} - 1|^2 - \|e\|_{\mathbb{T}}^2 \leq \sum_{\{i,j\} \in E} \left(\frac{1}{2}\|\gamma_j - \gamma_i\|_{\mathbb{T}}^2 - \|e_{ij}\|_{\mathbb{T}}^2 \right) \leq \sum_{\{i,j\} \in E} \|\gamma_j - \gamma_i - e_{ij}\|_{\mathbb{T}}^2. \quad (6.5)$$

Denoting $\hat{\omega}_i := \arg(u_i)$ and $\omega_i^* := \arg(\langle x, \varphi_i \rangle)$, then the definition of η gives

$$\eta(\hat{\omega}) = \sum_{\{i,j\} \in E} |e^{i(\hat{\omega}_j - \hat{\omega}_i - e_{ij})} - 1|^2.$$

With this, we continue (6.5) by applying the left-hand inequality of (6.4):

$$\frac{1}{2} \sum_{\{i,j\} \in E} |e^{i(\gamma_j - \gamma_i)} - 1|^2 - \|e\|_{\mathbb{T}}^2 \leq \frac{\pi^2}{4} \eta(\hat{\omega}) \leq \frac{\pi^2}{4} \frac{C'}{\tau} \min_{\omega: V \rightarrow \mathbb{T}} \eta(\omega) \leq \frac{\pi^2}{4} \frac{C'}{\tau} \eta(\omega^*), \quad (6.6)$$

where the second inequality follows from Theorem 6.4. Furthermore, the right-hand inequality of (6.4) gives

$$\eta(\omega^*) = \sum_{\{i,j\} \in E} |e^{-ie_{ij}} - 1|^2 \leq \|e\|_{\mathbb{T}}^2,$$

and so combining this with (6.6) gives

$$\sum_{\{i,j\} \in E} |e^{i(\gamma_j - \gamma_i)} - 1|^2 \leq 2\left(\frac{\pi^2}{4} \frac{C'}{\tau} + 1\right) \|e\|_{\mathbb{T}}^2. \quad (6.7)$$

At this point, take $\alpha := \frac{1}{\text{vol}(G)} \sum_{i \in V} \deg(i) e^{i\gamma_i}$, define v entrywise by $v_i := e^{i\gamma_i}$, and set $w = v - \alpha 1$. Then

$$1^* D w = \sum_{i \in V} \deg(i) (e^{i\gamma_i} - \alpha) = 0,$$

i.e., $D^{1/2} w$ is orthogonal to $D^{1/2} 1$. Also, since $(D - A)1 = 0$, we have

$$L(D^{1/2} 1) = (I - D^{-1/2} A D^{-1/2}) D^{1/2} 1 = D^{-1/2} (D - A) D^{-1/2} D^{1/2} 1 = 0,$$

and so $D^{1/2} w$ is orthogonal to a first eigenvector of L , considering L is positive semidefinite. Thus,

$$\frac{(D^{1/2} w)^* L (D^{1/2} w)}{w^* D w} \geq \min_{\substack{y \in \mathbb{C}^{|V|} \\ y \perp D^{1/2} 1}} \frac{y^* L y}{y^* y} = \tau.$$

Rearranging then gives

$$\tau w^* D w \leq (D^{1/2} w)^* L (D^{1/2} w) = (v - \alpha 1)^* (D - A) (v - \alpha 1) = v^* (D - A) v.$$

Continuing, we apply the definitions of D and A to get

$$\tau w^* D w \leq 2|E| - \sum_{i \in V} \sum_{j \in V} e^{-i\gamma_j} A[i, j] e^{i\gamma_i} = \sum_{\{i,j\} \in E} \left(2 - e^{i(\gamma_i - \gamma_j)} - e^{i(\gamma_j - \gamma_i)}\right) = \sum_{\{i,j\} \in E} |e^{i\gamma_j} - e^{i\gamma_i}|^2.$$

Next, we factor $e^{i\gamma_i}$ from the inside of each term and apply (6.7) to get

$$\tau \sum_{i \in V} \deg(i) |e^{i\gamma_i} - \alpha|^2 = \tau w^* D w \leq \sum_{\{i,j\} \in E} |e^{i(\gamma_j - \gamma_i)} - 1|^2 \leq 2\left(\frac{\pi^2}{4} \frac{C'}{\tau} + 1\right) \|e\|_{\mathbb{T}}^2. \quad (6.8)$$

From here, we proceed in two cases. First, when $\alpha \neq 0$, we may take $\theta := \arg(\alpha)$. Then the left-hand inequality of (6.4) and Lemma 6.6 give

$$\sum_{i \in V} \|\gamma_i - \theta\|_{\mathbb{T}}^2 \leq \frac{\pi^2}{4} \sum_{i \in V} |e^{i(\gamma_i - \theta)} - 1|^2 = \frac{\pi^2}{4} \sum_{i \in V} |e^{i\gamma_i} - \frac{\alpha}{|\alpha|}|^2 \leq \pi^2 \sum_{i \in V} |e^{i\gamma_i} - \alpha|^2 \leq \pi^2 \sum_{i \in V} \deg(i) |e^{i\gamma_i} - \alpha|^2, \quad (6.9)$$

where the last inequality uses the fact that $\deg(i) \geq 1$ for every $i \in V$, i.e., the graph G has no isolated vertex since G is connected, which follows from the fact that $\tau > 0$. In the case where $\alpha = 0$, we may arbitrarily take $\theta = 0$. Then similar analysis yields

$$\sum_{i \in V} \|\gamma_i - \theta\|_{\mathbb{T}}^2 \leq \frac{\pi^2}{4} \sum_{i \in V} |e^{i\gamma_i} - 1|^2 \leq \pi^2 |V| \leq \pi^2 \sum_{i \in V} \deg(i) = \pi^2 \sum_{i \in V} \deg(i) |e^{i\gamma_i} - \alpha|^2,$$

where in this case, the second inequality applies the triangle inequality to each term instead of Lemma 6.6. Note that both cases produce the same bound on $\sum_{i \in V} \|\gamma_i - \theta\|_{\mathbb{T}}^2$, which we now conclude by combining (6.8) and (6.9):

$$\sum_{i \in V} \|\gamma_i - \theta\|_{\mathbb{T}}^2 \leq \frac{2\pi^2}{\tau} \left(\frac{\pi^2}{4} \frac{C'}{\tau} + 1\right) \|e\|_{\mathbb{T}}^2 \leq \frac{2\pi^2}{\tau} \left(\frac{\pi^2}{4} \frac{C'}{\tau} + 1\right) \cdot \frac{\pi^2 \|e\|_{\mathbb{T}}^2}{P^2}.$$

The last inequality follows from Lemma 3.2, taking $z = \overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle + \varepsilon_{ij}$ and $\varepsilon = -\varepsilon_{ij}$. \square

6.3. Projective uniformity. This section is motivated by Theorem 6.3 of the previous section, which exhibits significant dependence on the size of P . Here, we show how Algorithm 1 ensures that P will not too small, and our guarantee will be in terms of the following noise-robust version of projective uniformity:

DEFINITION 6.7 (Projective uniformity with noise). *Consider a graph $G = (V, E)$ and $M \times |V|$ matrix Φ , and for some proportion $\alpha \in (0, 1)$, signal $x \in \mathbb{C}^M$ and noise $\varepsilon = \{\varepsilon_{ij}\}_{\{i,j\} \in E}$, let $\mathcal{J}(\alpha, x, \varepsilon)$ denote the set of vertices that remain after applying Algorithm 1. Then the projective uniformity with noise of Φ is*

$$\text{PUN}(\Phi; \alpha, \varepsilon) = \min_{\substack{x \in \mathbb{C}^M \\ \|x\|=1}} \min_{\{i,j\} \in \mathcal{J}(\alpha, x, \varepsilon)} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle + \varepsilon_{ij}|.$$

THEOREM 6.8. *Consider a graph $G = (V, E)$ with $|V| \sim CM \log M$ for some constant C , and draw the entries of an $M \times |V|$ matrix Φ independently from $\mathcal{CN}(0, \frac{1}{M})$. Then for each proportion $\alpha < 1 - \frac{1}{2C}$, there exists a constant $C' > 0$ such that, with overwhelming probability,*

$$\max \left\{ \|\varepsilon\|_2, \sqrt{M} \text{PUN}(\Phi; \alpha, \varepsilon) \right\} \geq \frac{C'}{\sqrt{M}}$$

for every noise vector $\varepsilon = \{\varepsilon_{ij}\}_{\{i,j\} \in E}$.

To prove this theorem, we apply the following lemma, which follows from concentration-of-measure arguments that we provide later:

LEMMA 6.9. *Take $n \sim CM \log M$ for some constant C , and draw the entries of an $M \times n$ matrix Φ independently from $\mathcal{CN}(0, \frac{1}{M})$. Then for each proportion $\alpha < 1 - \frac{1}{2C}$, there exists a constant $C' > 0$ such that*

$$\text{PU}(\Phi; \alpha) \geq \frac{C'}{M}$$

with overwhelming probability.

Proof. [Proof of Theorem 6.8] Denote $n := |V|$. Note that if

$$\|\varepsilon\|_2 > \sqrt{\frac{1-\alpha}{2}n} \cdot \frac{1}{2} \text{PU}(\Phi; \alpha + \frac{1+\alpha}{2}), \quad (6.10)$$

then by Lemma 6.9, we have $\|\varepsilon\|_2 \geq \frac{C'}{\sqrt{M}}$ with overwhelming probability for some constant $C' > 0$. It remains to consider the case where (6.10) does not hold. To this end, define

$$\mathcal{J}_1(\alpha, x, \varepsilon) := \operatorname{argmax}_{\substack{\mathcal{J} \subseteq \{1, \dots, n\} \\ |\mathcal{J}| \geq \alpha n}} \min_{\{i,j\} \in \mathcal{J}} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle + \varepsilon_{ij}|.$$

Recalling the definition of $\mathcal{J}(\alpha, x, \varepsilon)$ in Definition 6.7, we claim that $\mathcal{J}(\alpha, x, \varepsilon) \subseteq \mathcal{J}_1(\alpha, x, \varepsilon)$. To see this, label each edge $\{i, j\} \in E$ with $|\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle + \varepsilon_{ij}|$. Then by definition, $V \setminus \mathcal{J}_1(\alpha, x, \varepsilon)$ neighbors more of the smallest edges than any other collection of $n - \lceil \alpha n \rceil = \lfloor (1 - \alpha)n \rfloor$ vertices. By comparison, Algorithm 1 effectively deletes the smallest edge $\{i, j\}$ by deleting both of its incident vertices i and j , one of which must be in $\mathcal{J}_1(\alpha, x, \varepsilon)$. The smallest edge in the remaining graph is guaranteed to not touch either i or j , but rather some $k \in \mathcal{J}_1(\alpha, x, \varepsilon)$, provided $\mathcal{J}_1(\alpha, x, \varepsilon)$ has not yet been completely removed from the graph. After $\lfloor (1 - \alpha)n \rfloor$ iterations, then by the pigeonhole principle, all vertices in $V \setminus \mathcal{J}_1(\alpha, x, \varepsilon)$ will be removed, meaning $\mathcal{J}(\alpha, x, \varepsilon) \subseteq \mathcal{J}_1(\alpha, x, \varepsilon)$, as claimed. This implies that

$$\min_{\{i,j\} \in \mathcal{J}(\alpha, x, \varepsilon)} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle + \varepsilon_{ij}| \geq \min_{\{i,j\} \in \mathcal{J}_1(\alpha, x, \varepsilon)} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle + \varepsilon_{ij}|,$$

and so minimizing over all unit vectors x gives

$$\text{PUN}(\Phi; \alpha, \varepsilon) \geq \min_{\substack{x \in \mathbb{C}^M \\ \|x\|=1}} \max_{\substack{\mathcal{J} \subseteq \{1, \dots, n\} \\ |\mathcal{J}| \geq \alpha n}} \min_{\{i,j\} \in \mathcal{J}} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle + \varepsilon_{ij}|. \quad (6.11)$$

Next, define

$$\mathcal{J}_2(\alpha, x, \varepsilon) := \operatorname{argmax}_{\substack{\mathcal{J} \subseteq \{1, \dots, n\} \\ |\mathcal{J}| \geq (\alpha + \frac{1-\alpha}{2})n}} \min_{\{i, j\} \in \mathcal{J}} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle|.$$

Since in this case (6.10) does not hold, there are at most $\frac{1-\alpha}{2}n$ edges $\{i, j\} \in E$ with $|\varepsilon_{ij}| \geq \frac{1}{2}\operatorname{PU}(\Phi; \alpha + \frac{1+\alpha}{2})$. Some of these edges are induced by $\mathcal{J}_2(\alpha, x, \varepsilon)$; as such, delete at most $\frac{1-\alpha}{2}n$ vertices from $\mathcal{J}_2(\alpha, x, \varepsilon)$ which are incident to all of these edges, and denote the remaining vertices by $\mathcal{J}_3(\alpha, x, \varepsilon)$. By construction, we have $|\mathcal{J}_3(\alpha, x, \varepsilon)| \geq \alpha n$, and also the edges $\{i, j\}$ induced by $\mathcal{J}_3(\alpha, x, \varepsilon)$ satisfy $|\varepsilon_{ij}| < \frac{1}{2}\operatorname{PU}(\Phi; \alpha + \frac{1+\alpha}{2})$. These facts combined with the triangle inequality then give

$$\begin{aligned} \max_{\substack{\mathcal{J} \subseteq \{1, \dots, n\} \\ |\mathcal{J}| \geq \alpha n}} \min_{\{i, j\} \in \mathcal{J}} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle + \varepsilon_{ij}| &\geq \min_{\{i, j\} \in \mathcal{J}_3(\alpha, x, \varepsilon)} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle + \varepsilon_{ij}| \\ &\geq \min_{\{i, j\} \in \mathcal{J}_3(\alpha, x, \varepsilon)} \left(|\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle| - |\varepsilon_{ij}| \right) \\ &\geq \min_{\{i, j\} \in \mathcal{J}_3(\alpha, x, \varepsilon)} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle| - \frac{1}{2}\operatorname{PU}(\Phi; \alpha + \frac{1+\alpha}{2}). \end{aligned} \quad (6.12)$$

Continuing, we use the fact that $\mathcal{J}_3(\alpha, x, \varepsilon) \subseteq \mathcal{J}_2(\alpha, x, \varepsilon)$ to get

$$\begin{aligned} \min_{\{i, j\} \in \mathcal{J}_3(\alpha, x, \varepsilon)} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle| &\geq \min_{\{i, j\} \in \mathcal{J}_2(\alpha, x, \varepsilon)} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle| \\ &\geq \min_{\substack{x \in \mathbb{C}^M \\ \|x\|=1}} \max_{\substack{\mathcal{J} \subseteq \{1, \dots, n\} \\ |\mathcal{J}| \geq (\alpha + \frac{1-\alpha}{2})n}} \min_{\{i, j\} \in \mathcal{J}} |\overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle| \\ &\geq \operatorname{PU}(\Phi; \alpha + \frac{1+\alpha}{2}). \end{aligned} \quad (6.13)$$

Combining (6.11), (6.12) and (6.13) then gives $\operatorname{PUN}(\Phi; \alpha, \varepsilon) \geq \frac{1}{2}\operatorname{PU}(\Phi; \alpha + \frac{1+\alpha}{2})$, and so the result follows directly from Lemma 6.9. \square

We conclude this section with the rather technical proof of Lemma 6.9:

Proof. [Proof of Lemma 6.9] We seek a bound on the probability that $\operatorname{PU}(\Phi; \alpha) < \frac{C'}{M}$. To do so, we will cover this event with smaller failure events \mathcal{E}_k which correspond to members of a δ -net. First, define

$$G_\delta(v) := \{\varphi \in \mathbb{C}^M : |\langle v, \varphi \rangle| \geq 3\delta \text{ and } \|\varphi\| \leq 2\}.$$

Then for each member v of a given δ -net \mathcal{N}_δ of the unit sphere in \mathbb{C}^M , define the failure event

$$\mathcal{E}_v := \left\{ \text{less than } \alpha n \text{ columns of } \Phi \text{ lie in } G_\delta(v) \right\}.$$

We claim that $\{\operatorname{PU}(\Phi; \alpha) < \delta^2\} \subseteq \bigcup_{v \in \mathcal{N}_\delta} \mathcal{E}_v$. To see this, take an outcome $\omega \in (\bigcup_{v \in \mathcal{N}_\delta} \mathcal{E}_v)^c$. By the definition of \mathcal{N}_δ , we have that for every unit vector $x \in \mathbb{C}^M$, there exists $v \in \mathcal{N}_\delta$ such that $\|v - x\| \leq \delta$. This gives a mapping $x \mapsto v_x$, and we let \mathcal{I}_x denote the indices of columns of Φ which lie in $G_\delta(v_x)$. Since $\omega \in \bigcap_{v \in \mathcal{N}_\delta} \mathcal{E}_v^c$, we have $|\mathcal{I}_x| \geq \alpha n$ for every x . Moreover, by the triangle inequality, Cauchy-Schwarz inequality, and the definition of G_δ , every $i \in \mathcal{I}_x$ satisfies

$$|\langle x, \varphi_i \rangle| \geq |\langle v_x, \varphi_i \rangle| - |\langle v_x - x, \varphi_i \rangle| \geq |\langle v_x, \varphi_i \rangle| - \|v_x - x\| \|\varphi_i\| \geq 3\delta - 2\|v_x - x\| \geq \delta.$$

Applying this inequality then gives

$$\operatorname{PU}(\Phi; \alpha) = \min_{\substack{x \in \mathbb{C}^M \\ \|x\|=1}} \max_{\substack{\mathcal{I} \subseteq \{1, \dots, n\} \\ |\mathcal{I}| \geq \alpha n}} \min_{i \in \mathcal{I}} |\langle x, \varphi_i \rangle|^2 \geq \min_{\substack{x \in \mathbb{C}^M \\ \|x\|=1}} \min_{i \in \mathcal{I}_x} |\langle x, \varphi_i \rangle|^2 \geq \delta^2.$$

Thus $\omega \in \{\operatorname{PU}(\Phi; \alpha) < \delta^2\}^c$, thereby proving our claim. Continuing, the union bound gives

$$\operatorname{Pr}\left(\operatorname{PU}(\Phi; \alpha) < \delta^2\right) \leq \sum_{v \in \mathcal{N}_\delta} \operatorname{Pr}(\mathcal{E}_v) = |\mathcal{N}_\delta| \cdot \operatorname{Pr}(\mathcal{E}_v),$$

where the equality follows from the fact that the columns of Φ are independent with rotationally symmetric probability distributions. It therefore suffices to bound both $|\mathcal{N}_\delta|$ and $\Pr(\mathcal{E}_v)$.

To bound $|\mathcal{N}_\delta|$, we follow a standard argument, found in the proof of Lemma 5.2 in [47]. Let \mathcal{N}_δ be a maximal δ -packing of points on the unit sphere in \mathbb{C}^M . Since the packing is maximal, it follows that \mathcal{N}_δ is a δ -net. To count these points, map them into \mathbb{R}^{2M} according to $f: v \mapsto (\operatorname{Re} v, \operatorname{Im} v)$. Consider the open balls of radius $\frac{\delta}{2}$ centered at each $f(v)$. Note that the (disjoint) union of these balls is contained in the ball of radius $1 + \frac{\delta}{2}$ centered at the origin. Thus, a volume comparison gives

$$|\mathcal{N}_\delta| \cdot C\left(\frac{\delta}{2}\right)^{2M} = \operatorname{Vol}\left(\bigsqcup_{v \in \mathcal{N}_\delta} B(f(v), \frac{\delta}{2})\right) \leq \operatorname{Vol}\left(B(0, 1 + \frac{\delta}{2})\right) = C\left(1 + \frac{\delta}{2}\right)^{2M},$$

thereby implying $|\mathcal{N}_\delta| \leq \left(\frac{2}{\delta} + 1\right)^{2M}$.

To bound $\Pr(\mathcal{E}_v)$, note that

$$\Pr(\mathcal{E}_v) = \Pr\left(\sum_{i=1}^n \mathbf{1}_{\{\varphi_i \notin G_\delta(v)\}} \geq (1 - \alpha)n\right). \quad (6.14)$$

As such, we first consider the success probability of these Bernoulli random variables:

$$\Pr(\varphi_i \notin G_\delta(v)) \leq \Pr(|\langle v, \varphi_i \rangle| < 3\delta) + \Pr(\|\varphi_i\| > 2). \quad (6.15)$$

By the rotational symmetry of the distribution of φ_i , we may take v to be the first identity basis element e_1 without changing the probability. Next, since $\varphi_i = \sum_{m=1}^M (a_m + ib_m)e_m$ with the a_m 's and b_m 's independent with distribution $\mathcal{N}(0, \frac{1}{2M})$, we have

$$\Pr(|\langle v, \varphi_i \rangle| < 3\delta) = \Pr\left(\sqrt{a_1^2 + b_1^2} < 3\delta\right) \leq \Pr\left(\min\{a_1^2, b_1^2\} < \frac{9}{2}\delta^2\right) \leq 2\Pr\left(a_1^2 < \frac{9}{2}\delta^2\right),$$

where the last step is by the union bound. Since a_1 's density function is $\leq \sqrt{\frac{M}{\pi}}$, it follows that

$$\Pr(|\langle v, \varphi_i \rangle| < 3\delta) \leq 2\Pr\left(|a_1| < \frac{3}{\sqrt{2}}\delta\right) \leq 12\delta\sqrt{\frac{M}{2\pi}}. \quad (6.16)$$

For the other term in (6.15), note that $2M\|\varphi_i\|^2$ is a sum of independent standard Gaussian random variables. Applying Lemma 1 of [32] then gives that for every $t > 0$,

$$\Pr\left(2M\|\varphi_i\|^2 \geq \sqrt{8Mt} + 2t + 2M\right) \leq e^{-t}.$$

Thus, taking $t = \frac{M}{2}$ gives

$$\Pr(\|\varphi_i\| > 2) = \Pr\left(2M\|\varphi_i\|^2 > 8M\right) \leq \Pr\left(2M\|\varphi_i\|^2 \geq 5M\right) \leq e^{-M/2}. \quad (6.17)$$

Substituting (6.16) and (6.17) into (6.15) then gives

$$\Pr(\varphi_i \notin G_\delta(v)) \leq 12\delta\sqrt{\frac{M}{2\pi}} + e^{-M/2}. \quad (6.18)$$

Now, to bound (6.14), we will apply Hoeffding's inequality [28], which says that the tail probability of a sum of independent Bernoulli random variables X_i , each with success probability p , has the following bound:

$$\Pr\left(\sum_{i=1}^n X_i \geq n(p + t)\right) \leq e^{-2nt}.$$

Also, note that replacing p in the left-hand side above with some $p' \geq p$ will not increase the probability. As such, taking p' to be the right-hand side of (6.18) and $t = 1 - \alpha - p'$, we have

$$\Pr(\mathcal{E}_v) = \Pr\left(\sum_{i=1}^n 1_{\{\varphi_i \notin \mathcal{G}_\delta(v)\}} \geq (1 - \alpha)n\right) \leq \exp\left(-2n\left((1 - \alpha) - \left(12\delta\sqrt{\frac{M}{2\pi}} + e^{-M/2}\right)\right)\right).$$

Now that we have bounds on both $|\mathcal{N}_\delta|$ and $\Pr(\mathcal{E}_v)$, we note that we specifically wish to show that

$$\Pr\left(\text{PU}(\Phi; \alpha) < \frac{C'}{M}\right) < c_0 e^{-c_1 M}$$

for some constants $c_0, c_1 > 0$. Considering the above analysis and taking logarithms, it suffices to have

$$\log |\mathcal{N}_\delta| + \log \Pr(\mathcal{E}_v) \leq 2M \log\left(\frac{2}{\delta} + 1\right) - 2n(1 - \alpha) + 2n\left(12\delta\sqrt{\frac{M}{2\pi}} + e^{-M/2}\right) < \log c_0 - c_1 M$$

with $\delta^2 = \frac{C'}{M}$. Since $\alpha < 1 - \frac{1}{2C}$ by assumption, picking $C' < \frac{\pi}{72}\left(1 - \frac{1}{2C} - \alpha\right)^2$ will make $2n(1 - \alpha)$ the dominant term in the above inequality, thereby proving the result. \square

6.4. Removing large vertices. In this section, we prove how well we can remove the vertices with the largest noisy intensity measurements. We start with a lemma:

LEMMA 6.10. *Pick $C > 0$ and $n \geq e^{C/8} M \log M$. Draw the entries of an $M \times n$ matrix $\Phi = \{\varphi_i\}_{i=1}^n$ independently from $\mathcal{CN}(0, \frac{1}{M})$. Taking $\beta := 3e^{-C/8}$, then with overwhelming probability,*

$$\#\left\{i : |\langle x, \varphi_i \rangle|^2 > \frac{C}{M}\right\} < \beta n \quad (6.19)$$

for every unit norm $x \in \mathbb{C}^M$.

Proof. Take a δ -net \mathcal{N}_δ of the unit sphere in \mathbb{C}^M . Then for every unit vector x , there is a $v_x \in \mathcal{N}_\delta$ such that

$$|\langle x, \varphi_i \rangle| \leq |\langle v_x, \varphi_i \rangle| + |\langle x - v_x, \varphi_i \rangle| \leq |\langle v_x, \varphi_i \rangle| + \|x - v_x\| \leq |\langle v_x, \varphi_i \rangle| + \delta.$$

Taking $\delta = \frac{1}{2}\sqrt{\frac{C}{M}}$, we then have that $|\langle x, \varphi_i \rangle|^2 > \frac{C}{M}$ implies $|\langle v_x, \varphi_i \rangle|^2 > \frac{C}{4M}$. Recall that we wish to bound the probability of the event \mathcal{E} that (6.19) is violated for some unit vector x . To this end, the above implication allows us to focus on a finite set of points:

$$\Pr(\mathcal{E}) \leq \Pr\left(\exists v \in \mathcal{N}_\delta \text{ s.t. } \sum_{i=1}^n 1_{\{|\langle v, \varphi_i \rangle|^2 > \frac{C}{4M}\}} \geq \beta n\right).$$

As discussed in the proof of Lemma 6.9, we may take $|\mathcal{N}_\delta| \leq \left(\frac{2}{\delta} + 1\right)^{2M}$, and so the union bound and the symmetric distribution of Φ both give

$$\Pr(\mathcal{E}) \leq \left(4\sqrt{\frac{M}{C}} + 1\right)^{2M} \Pr\left(\sum_{i=1}^n 1_{\{|\langle v, \varphi_i \rangle|^2 > \frac{C}{4M}\}} \geq \beta n\right). \quad (6.20)$$

Similar to the proof of Lemma 6.9, we will bound the probability on the right-hand side using Hoeffding's inequality. First, we note that the symmetric distribution of $\varphi_i = \sum_{m=1}^M (a_m + ib_m)e_m$ gives that

$$\Pr\left(|\langle v, \varphi_i \rangle|^2 > \frac{C}{4M}\right) = \Pr\left(|\langle e_1, \varphi_i \rangle|^2 > \frac{C}{4M}\right) = \Pr\left(a_1^2 + b_1^2 > \frac{C}{4M}\right) \leq \Pr\left(a_1^2 > \frac{C}{8M}\right) + \Pr\left(b_1^2 > \frac{C}{8M}\right),$$

where the last step is by the union bound. We continue, using the fact that a_1 and b_1 are both distributed as $\mathcal{N}(0, \frac{1}{2M})$:

$$\Pr\left(|\langle v, \varphi_i \rangle|^2 > \frac{C}{4M}\right) \leq 2\Pr\left(a_1^2 > \frac{C}{8M}\right) = 2\Pr\left(z^2 > \frac{C}{4}\right) \leq 2e^{-C/8},$$

where z is a standard Gaussian random variable. With this, we now apply Hoeffding's inequality to (6.20):

$$\Pr(\mathcal{E}) \leq \left(4\sqrt{\frac{M}{C}} + 1\right)^{2M} \exp\left(-2n(\beta - 2e^{-C/8})\right) = \exp\left(2M \log\left(4\sqrt{\frac{M}{C}} + 1\right) - 2ne^{-C/8}\right).$$

Since $n \geq e^{C/8}M \log M$, then $2ne^{-C/8}$ is the dominant term above, thereby proving the result. \square

THEOREM 6.11. *Pick $C > 0$ and $n \geq e^{C/8}M \log M$. Draw the entries of an $M \times n$ matrix $\Phi = \{\varphi_i\}_{i=1}^n$ independently from $\mathcal{CN}(0, \frac{1}{M})$. Given a vector $x \in \mathbb{C}^M$, consider*

$$z_i := |\langle x, \varphi_i \rangle|^2 + \nu_i,$$

where $\nu = \{\nu_i\}_{i=1}^n$ satisfies $\frac{\|\nu\|}{\|x\|^2} \leq \frac{C'}{\sqrt{M}}$. Taking $\beta := 3e^{-C/8}$, we have

$$\#\left\{i : z_i > \frac{2C}{M}\|x\|^2\right\} < 2\beta n \quad (6.21)$$

with overwhelming probability.

Proof. In counting large z_i 's, we identify which come from large or small inner products $|\langle x, \varphi_i \rangle|^2$. First,

$$\left\{i : z_i > \frac{2C}{M}\|x\|^2\right\} \cap \left\{i : |\langle x, \varphi_i \rangle|^2 > \frac{C}{M}\|x\|^2\right\}$$

is of size $< \beta n$ with overwhelming probability by Lemma 6.10. The rest of the large z_i 's have indices in

$$\mathcal{K} := \left\{i : z_i > \frac{2C}{M}\|x\|^2\right\} \cap \left\{i : |\langle x, \varphi_i \rangle|^2 \leq \frac{C}{M}\|x\|^2\right\}.$$

To count these, note that $\nu_i = z_i - |\langle x, \varphi_i \rangle|^2 \geq \frac{C}{M}\|x\|^2$, and so

$$|\mathcal{K}| \frac{C^2}{M^2} \|x\|^4 \leq \sum_{i \in \mathcal{K}} |\nu_i|^2 \leq \|\nu\|^2 \leq \frac{(C')^2}{M} \|x\|^4.$$

Rearranging then reveals that $|\mathcal{K}| = \mathcal{O}(M)$, meaning \mathcal{K} has fewer than $\beta n \geq 3M \log M$ members when M is sufficiently large. \square

6.5. Main result. This section proves the main result of the paper, which we restate here:

THEOREM 6.12. *Pick $N \sim CM \log M$ with C sufficiently large, and take $\{\varphi_\ell\}_{\ell=1}^N = \Phi_V \cup \Phi_E$ defined in the measurement design of Section 3. Then there exist constants $C', K > 0$ such that the following guarantee holds for all $x \in \mathbb{C}^M$ with overwhelming probability: Consider measurements of the form*

$$z_\ell := |\langle x, \varphi_\ell \rangle|^2 + \nu_\ell.$$

If the noise-to-signal ratio satisfies $\text{NSR} := \frac{\|\nu\|}{\|x\|^2} \leq \frac{C'}{\sqrt{M}}$, then the phase retrieval procedure of Section 3 produces an estimate \tilde{x} from $\{z_\ell\}_{\ell=1}^N$ with squared relative error

$$\frac{\|\tilde{x} - e^{i\theta}x\|^2}{\|x\|^2} \leq K \sqrt{\frac{M}{\log M}} \text{NSR}$$

for some phase $\theta \in [0, 2\pi)$.

Proof. We will prove the result by considering the steps of our phase retrieval process in reverse order. In the last step, we have the following estimates of $\langle x, \varphi_i \rangle$ for every vertex $i \in V'' \subseteq V$ which survives our graph-pruning and large-vertex-removing processes:

$$y_i := e^{i\theta} \langle x, \varphi_i \rangle + \delta_i = (e^{i\theta} \Phi_{V''}^* x + \delta)_i.$$

Here, θ is a global phase which is calculated in the proof of Theorem 6.3. From these estimates, we reconstruct by finding the least-squares estimate of x :

$$\tilde{x} := (\Phi_{V''} \Phi_{V''}^*)^{-1} \Phi_{V''} y = e^{i\theta} x + (\Phi_{V''} \Phi_{V''}^*)^{-1} \Phi_{V''} \delta.$$

As such, we have the following bound on the reconstruction error:

$$\|\tilde{x} - e^{i\theta} x\| = \|(\Phi_{V''} \Phi_{V''}^*)^{-1} \Phi_{V''} \delta\| \leq \|(\Phi_{V''} \Phi_{V''}^*)^{-1} \Phi_{V''}\|_2 \|\delta\| = \frac{\|\delta\|}{\sigma_{\min}(\Phi_{V''}^*)}, \quad (6.22)$$

where the last equality holds with probability 1, specifically, in the event that $\sigma_{\min}(\Phi_{V''}^*) > 0$.

To continue this bound, recall that the large-vertex-removing process ensured $|V''| = \kappa n$. We claim there exists a constant $c_0 > 0$ such that with overwhelming probability, $\sigma_{\min}(\Phi_{V''}^*) \geq c_0 \sqrt{\log M}$ for every $V'' \subseteq V$ with $|V''| = \kappa n$. To prove this, we leverage the complex version of Corollary 5.35 in [47], which gives that for a fixed $V'' \subseteq V$, there exist $c, c' > 0$ such that

$$\Pr\left(\sigma_{\min}(\Phi_{V''}^*) < \frac{1}{\sqrt{M}}(\sqrt{\kappa n} - \sqrt{M} - \sqrt{tn})\right) \leq ce^{-c'tn}.$$

Performing a union bound over the $\binom{n}{\kappa n} = \binom{n}{(1-\kappa)n} \leq \left(\frac{e}{1-\kappa}\right)^{(1-\kappa)n}$ choices for V'' then gives

$$\Pr\left(\exists V'' \subseteq V, |V''| = \kappa n \text{ s.t. } \sigma_{\min}(\Phi_{V''}^*) < \frac{1}{\sqrt{M}}(\sqrt{\kappa n} - \sqrt{M} - \sqrt{tn})\right) \leq c \exp\left(c'n\left(-\frac{t}{2} + (1-\kappa)\log\left(\frac{e}{1-\kappa}\right)\right)\right).$$

Provided $\frac{\kappa}{2} > (1-\kappa)\log\left(\frac{e}{1-\kappa}\right)$, which occurs whenever $\kappa \geq 0.86$, then taking $t := \frac{\kappa}{2} + (1-\kappa)\log\left(\frac{e}{1-\kappa}\right)$ will ensure both $t < \kappa$ and $\frac{t}{2} > (1-\kappa)\log\left(\frac{e}{1-\kappa}\right)$. Therefore, this choice for t coupled with the fact that $n = CM \log M$ proves the claim. Continuing (6.22) then gives

$$\|\tilde{x} - e^{i\theta} x\| \leq \frac{\|\delta\|}{c_0 \sqrt{\log M}}. \quad (6.23)$$

Next, we wish to bound $\|\delta\|$. By definition, we have

$$\delta_i = y_i - e^{i\theta} \langle x, \varphi_i \rangle = \sqrt{z_i} e^{i \arg(u_i)} - e^{i\theta} \langle x, \varphi_i \rangle.$$

Note that the above square root operates under the assumption that $z_i \geq 0$ for each $i \in V''$, which is ensured when we prune for reliability. Denote $\xi_i := \sqrt{z_i} - |\langle x, \varphi_i \rangle|$. Then by the triangle inequality, we have

$$|\delta_i| = \left| \sqrt{z_i} e^{i \arg(u_i)} - \sqrt{z_i} e^{i(\theta + \arg(\langle x, \varphi_i \rangle))} + \xi_i e^{i(\theta + \arg(\langle x, \varphi_i \rangle))} \right| \leq \sqrt{z_i} \left| e^{i \arg(u_i)} - e^{i(\theta + \arg(\langle x, \varphi_i \rangle))} \right| + |\xi_i|.$$

Next, factoring $e^{i(\theta + \arg(\langle x, \varphi_i \rangle))}$ and applying (6.4) yields

$$|\delta_i| \leq \sqrt{z_i} \left| \arg(u_i) - \arg(\langle x, \varphi_i \rangle) - \theta \right|_{\mathbb{T}} + |\xi_i|.$$

For any $a, b \geq 0$, then since $0 \leq (a-b)^2 = a^2 - 2ab + b^2$, we have $(a+b)^2 = a^2 + 2ab + b^2 \leq 2(a^2 + b^2)$. Applying this inequality to the right-hand side above then gives

$$|\delta_i|^2 \leq 2z_i \left| \arg(u_i) - \arg(\langle x, \varphi_i \rangle) - \theta \right|_{\mathbb{T}}^2 + 2\xi_i^2.$$

Similarly, for any $a, b \geq 0$, then since $ab \geq \min\{a^2, b^2\}$, we have $(a-b)^2 = a^2 - 2ab + b^2 \leq |a^2 - b^2|$. Applying this to $\xi_i^2 = (\sqrt{z_i} - |\langle x, \varphi_i \rangle|)^2$ then gives $\xi_i^2 \leq |\nu_i|$. Also by Theorem 6.11, our large-vertex-removing process ensures that $z_i \leq \frac{2c_1}{M} \|x\|^2$ for every $i \in V''$. Combined, these facts imply

$$\|\delta\|^2 \leq \frac{4c_1}{M} \|x\|^2 \sum_{i \in V''} \left| \arg(u_i) - \arg(\langle x, \varphi_i \rangle) - \theta \right|_{\mathbb{T}}^2 + 2\|\nu_V\|_1. \quad (6.24)$$

Applying Theorem 6.3, Definition 6.7 and Theorem 6.8 further gives

$$\sum_{i \in V''} \|\arg(u_i) - \arg(\langle x, \varphi_i \rangle) - \theta\|_{\mathbb{T}}^2 \leq \frac{c_2 \|\varepsilon\|^2}{\tau^2 P^2} \leq \frac{c_2 \|\varepsilon\|^2}{\tau^2 \|x\|^4 \text{PUN}(\Phi_V; \alpha, \frac{\varepsilon}{\|x\|^2})^2} \leq \frac{c_2 M^2 \|\varepsilon\|^2}{c_3^2 \tau^2 \|x\|^4}. \quad (6.25)$$

Recalling the definition of ε , we have

$$\varepsilon_{ij} = \left(\frac{1}{3} \sum_{k=0}^2 \zeta^k \left(|\langle x, \varphi_i + \zeta^k \varphi_j \rangle|^2 + \nu_{ijk} \right) \right) - \overline{\langle x, \varphi_i \rangle} \langle x, \varphi_j \rangle = \frac{1}{3} \sum_{k=0}^2 \zeta^k \nu_{ijk},$$

and so, letting E' denote the edges which remained after pruning for connectivity, the Cauchy-Schwarz inequality gives

$$\|\varepsilon\|^2 = \sum_{\{i,j\} \in E'} |\varepsilon_{ij}|^2 = \sum_{\{i,j\} \in E'} \left| \frac{1}{3} \sum_{k=0}^2 \zeta^k \nu_{ijk} \right|^2 \leq \frac{1}{9} \sum_{\{i,j\} \in E'} \left(\sum_{k=0}^2 |\zeta^{-k}|^2 \right) \left(\sum_{k=0}^2 |\nu_{ijk}|^2 \right) \leq \frac{1}{3} \|\nu\|^2. \quad (6.26)$$

Finally, we combine (6.23), (6.24), (6.25) and (6.26), along with $\|\nu_V\|_1 \leq \sqrt{n} \|\nu_V\| \leq \sqrt{2CM \log M} \|\nu\|$:

$$\frac{\|\tilde{x} - e^{i\theta} x\|^2}{\|x\|^2} \leq \frac{4c_1 c_2}{3c_0^2 c_3^2 \tau^2} \frac{M}{\log M} \frac{\|\nu\|^2}{\|x\|^4} + \frac{2\sqrt{2C}}{c_0^2} \sqrt{\frac{M}{\log M}} \frac{\|\nu\|}{\|x\|^2} \leq K \sqrt{\frac{M}{\log M}} \frac{\|\nu\|}{\|x\|^2},$$

for some constant K . \square

Acknowledgments. The authors thank the anonymous referees for providing thoughtful suggestions that led to a more complete discussion of the context of our results. The authors also thank Prof. Amit Singer for insightful discussions. B. Alexeev was supported by the NSF Graduate Research Fellowship under Grant No. DGE-0646086, A.S. Bandeira was supported by NSF Grant No. DMS-0914892, M. Fickus was supported by NSF Grant No. DMS-1042701 and AFOSR Grant Nos. F1ATA01103J001 and F1ATA00183G003, and D.G. Mixon was supported by the A.B. Krongard Fellowship. The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

REFERENCES

- [1] B. Alexeev, J. Cahill, D.G. Mixon, Full spark frames, arXiv:1110.3548
- [2] N. Alon, Eigenvalues and expanders, *Combinatorica* 6 (1986) 83–96.
- [3] N. Alon, V.D. Milman, λ_1 , isoperimetric inequalities for graphs, and superconcentrators, *J. Comb. Theory B* 38 (1985) 73–88.
- [4] M.F. Atiyah, F. Hirzebruch, Quelques théorème de non-plongement pour les variétés différentiables, *Bull. Soc. Math. Fr.* 87 (1959) 383–396.
- [5] R. Balan, B.G. Bodmann, P.G. Casazza, D. Edidin, Painless reconstruction from magnitudes of frame coefficients, *J. Fourier Anal. Appl.* 15 (2009) 488–501.
- [6] R. Balan, P.G. Casazza, D. Edidin, On signal reconstruction without phase, *Appl. Comput. Harmon. Anal.* 20 (2006) 345–356.
- [7] A.S. Bandeira, J. Cahill, D.G. Mixon, A.A. Nelson, Saving phase: Injectivity and stability for phase retrieval, arXiv:1302.4618
- [8] A.S. Bandeira, Y. Chen, D.G. Mixon, Phase retrieval from power spectra of masked signals, arXiv:1303.4458
- [9] A.S. Bandeira, A. Singer, D.A. Spielman, A Cheeger inequality for the graph connection Laplacian, arXiv:1204.3873
- [10] B.G. Bodmann, N. Hammen, Stable phase retrieval with low-redundancy frames, arXiv:1302.5487
- [11] O. Bunk, A. Diaz, F. Pfeiffer, C. David, B. Schmitt, D.K. Satapathy, J.F. van der Veen, Diffractive imaging for periodic samples: retrieving one-dimensional concentration profiles across microfluidic channels, *Acta Cryst.* A63 (2007) 306–314.
- [12] E.J. Candès, Y.C. Eldar, T. Strohmer, V. Voroninski, Phase retrieval via matrix completion, arXiv:1109.0573
- [13] E.J. Candès, X. Li, Solving quadratic equations via PhaseLift when there are about as many equations as unknowns, arXiv:1208.6247

- [14] E.J. Candès, T. Strohmer, V. Voroninski, PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming, arXiv:1109.4499
- [15] J. Cheeger, A lower bound for the smallest eigenvalue of the Laplacian, In: Problems in analysis (Papers dedicated to Salomon Bochner, 1969), Princeton Univ. Press, Princeton, 1970, 195–199.
- [16] F. Chung, Four proofs for the Cheeger inequality and graph partition algorithms, Fourth International Congress of Chinese Mathematicians (2010) 331–349.
- [17] J.C. Dainty, J.R. Fienup, Phase retrieval and image reconstruction for astronomy, In: H. Stark, ed., Image Recovery: Theory and Application, Academic Press, New York, 1987.
- [18] L. Demanet, P. Hand, Stable optimizationless recovery from phaseless linear measurements, arXiv:1208.1803
- [19] H. Duadi, O. Margalit, V. Mico, J.A. Rodrigo, T. Alieva, J. Garcia, Z. Zalevsky, Digital holography and phase retrieval, in Holography, Research and Technologies, J. Rosen, ed., InTech, 2011; available online at <http://www.intechopen.com/books/holography-research-and-technologies/digital-holography-and-phase-retrieval>.
- [20] M. Fickus, D.G. Mixon, Numerically erasure-robust frames, Linear Algebra Appl. 437 (2012) 1394–1407.
- [21] J.R. Fienup, Phase retrieval algorithms: a comparison, Appl. Optics 21 (1982) 2758–2769.
- [22] J. Friedman, A proof of Alon’s second eigenvalue conjecture and related problems, Mem. Amer. Math. Soc. 195 (2008).
- [23] R. Gerchberg, W. Saxton, A practical algorithm for the determination of phase from image and diffraction plane pictures, Optik 35 (1972) 237–246.
- [24] D. Griffin, J. Lim, Signal estimation from modified short-time Fourier transform, IEEE Trans. Acoustics Speech Signal Proc. 32 (1984) 236–243.
- [25] R.W. Harrison, Phase problem in crystallography, J. Opt. Soc. Am. A 10 (1993) 1046–1055.
- [26] P. Harsha, A. Barth, Lecture 5: Derandomization (Part II), <http://www.tcs.tifr.res.in/~prahladh/teaching/05spring/lectures/lec5.pdf>
- [27] T. Heinosaari, L. Mazzarella, M.M. Wolf, Quantum tomography under prior information, arXiv:1109.5478
- [28] W. Hoeffding, Probability inequalities for sums of bounded random variables, J. Am. Statist. Assoc. 58 (1963) 13–30.
- [29] C. Iaconis, I.A. Walmsley, Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses, Opt. Lett. 23 (1998) 792–794.
- [30] S. Janson, T. Luczak, and A. Ruciński, Random graphs, John Wiley, 2000.
- [31] D. Keusters, H.-S. Tan, P. O’Shea, E. Zeek, R. Trebino, W.S. Warren, Relative-phase ambiguities in measurements of ultrashort pulses with well-separated multiple frequency components, J. Opt. Soc. Am. B 20 (2003) 2226–2237.
- [32] B. Laurent, P. Massart, Adaptive estimation of a quadratic functional by model selection, Ann. Statist. 28 (2000) 1302–1338.
- [33] A. Lubotzky, R. Phillips, P. Sarnak, Ramanujan graphs, Combinatorica 8 (1988) 261–277.
- [34] A.M. Maiden, J.M. Rodenburg, An improved ptychographical phase retrieval algorithm for diffractive imaging, Ultramicroscopy 109 (2009) 1256–1262.
- [35] M. Maila, J. Shi, A random walks view of spectral segmentation, AI and STATISTICS (AISTATS) 2001, 2001.
- [36] S. Marchesini, A unified evaluation of iterative projection algorithms for phase retrieval, Rev. Sci. Inst. 78 (2007) 011301
- [37] J. Miao, T. Ishikawa, Q. Shen, T. Earnest, Extending X-ray crystallography to allow the imaging of noncrystalline materials, cells, and single protein complexes, Annu. Rev. Phys. Chem. 59 (2008) 387–410.
- [38] R.P. Millane, Phase retrieval in crystallography and optics, J. Opt. Soc. Am. A 7 (1990) 394–411.
- [39] D. Mondragon, V. Voroninski, Determination of all pure quantum states from a minimal number of observables, arXiv:1306.1214
- [40] B. Nadler, S. Lafon, R.R. Coifman, I.G. Kevrekidis, Diffusion maps, spectral clustering and eigenfunctions of Fokker-Planck operators, Adv. Neural. Inf. Process. Syst. 18 (2005) 955–962.
- [41] M. Püschel, J. Kovačević, Real, tight frames with maximal robustness to erasures, Proc. Data Compr. Conf. (2005) 63–72.
- [42] O. Raz, N. Dudovich, B. Nadler, Vectorial phase retrieval of 1-D signals, IEEE Trans. Signal Process. 61 (2013) 1632–1643.
- [43] O. Raz, O. Schwartz, D. Austin, A.S. Wyatt, A. Schiavi, O. Smirnova, B. Nadler, I.A. Walmsley, D. Oron, N. Dudovich, Vectorial phase retrieval for linear characterization of attosecond pulses, Phys. Rev. Lett. 107 (2011) 133902/1–5.
- [44] J.M. Rodenburg, Ptychography and related diffractive imaging methods, Adv. Imag. Elect. Phys. 150 (2008) 87–184.
- [45] H. Sahinoglou, S.D. Cabrera, On phase retrieval of finite-length sequences using the initial time sample, IEEE Trans. Circuits Syst. 38 (1991) 954–958.
- [46] A. Singer, Angular synchronization by eigenvectors and semidefinite programming, Appl. Comput. Harmon. Anal. 30 (2011) 20–36.
- [47] R. Vershynin, Introduction to the non-asymptotic analysis of random matrices, arXiv:1011.3027
- [48] I. Waldspurger, A. d’Aspremont, S. Mallat, Phase recovery, MaxCut and complex semidefinite programming, arXiv:1206.0102
- [49] A. Walther, The question of phase retrieval in optics, Opt. Acta 10 (1963) 41–49.
- [50] Y. Wang, Minimal frames for phase retrieval, Workshop on Phaseless Reconstruction, U. Maryland, February 24, 2013.