

Near-optimal bounds for phase synchronization

Yiqiao Zhong

ORFE

Princeton University

Nicolas Boumal

Mathematics Department and PACM

Princeton University

March 21, 2017

Abstract

The problem of phase synchronization is to estimate the phases (angles) of a complex unit-modulus vector z from their noisy pairwise relative measurements $C = zz^* + \sigma W$, where W is a complex-valued Gaussian random matrix. The maximum likelihood estimator (MLE) is a solution to a unit-modulus constrained quadratic programming problem, which is nonconvex. Existing works have proposed polynomial-time algorithms such as a semidefinite relaxation (SDP) approach or the generalized power method (GPM) to solve it. Numerical experiments suggest both of these methods succeed with high probability for σ up to $\tilde{\mathcal{O}}(n^{1/2})$, yet, existing analyses only confirm this observation for σ up to $\mathcal{O}(n^{1/4})$. In this paper, we bridge the gap, by proving SDP is tight for $\sigma = \mathcal{O}(\sqrt{n/\log n})$, and GPM converges to the global optimum under the same regime. Moreover, we establish a linear convergence rate for GPM, and derive a tighter ℓ_∞ bound for the MLE. A novel technique we develop in this paper is to track (theoretically) n closely related sequences of iterates, in addition to the sequence of iterates GPM actually produces. As a by-product, we obtain an ℓ_∞ perturbation bound for leading eigenvectors. Our result also confirms intuitions that use techniques from statistical mechanics.

Keywords: angular synchronization, nonconvex optimization, semidefinite relaxation, power method, maximum likelihood estimator, eigenvector perturbation bound.

1 Introduction

Phase synchronization is the problem of estimating n angles $\theta_1, \dots, \theta_n$ in $[0, 2\pi)$ based on noisy measurements of their differences $\theta_k - \theta_\ell \bmod 2\pi$. This is equivalent to estimating n phases $e^{i\theta_1}, \dots, e^{i\theta_n}$ from measurements of relative phases $e^{i(\theta_k - \theta_\ell)}$.

A typical noise model for this estimation problem is as follows. The target parameter (the signal) is the vector $z \in \mathbb{C}^n$ with entries $z_k = e^{i\theta_k}$. The measurements are stored

in a matrix $C \in \mathbb{C}^{n \times n}$ such that, for $k < \ell$,

$$C_{k\ell} = z_k \bar{z}_\ell + \sigma W_{k\ell}, \quad (1)$$

where $\sigma \geq 0$ is the noise level and $\{W_{k\ell}\}_{k < \ell}$ are independent standard complex Gaussian variables. Under this model, defining $W_{kk} = 0$ and $W_{\ell k} = \overline{W_{k\ell}}$ for consistency, the model is compactly written in matrix notation as

$$C = zz^* + \sigma W, \quad (2)$$

where both C and W are Hermitian. An easy derivation¹ shows that a maximum likelihood estimator (MLE) $\hat{x} \in \mathbb{C}^n$ for the signal z is a global optimum of the following quadratically constrained quadratic program (we define $[n] = \{1, \dots, n\}$):

$$\max_{x \in \mathbb{C}^n} x^* C x \quad \text{subject to } |x_k| = 1 \text{ for } k \in [n]. \quad (\text{P})$$

Problem (P) is non-convex and hard in general ([39, Prop. 3.5]). Yet, numerical experiments in [3, 2] suggest that, provided $\sigma = \tilde{\mathcal{O}}(\sqrt{n})$,² the following convex semidefinite relaxation for (P) admits $\hat{x}\hat{x}^*$ as its unique global optimum with high probability (more generally, if the problem below admits a solution of rank 1, $\hat{x}\hat{x}^*$, then the relaxation is said to be tight and \hat{x} is an optimum of (P)):

$$\max_{X \in \mathbb{C}^{n \times n}, X = X^*} \text{trace}(CX) \quad \text{subject to } \text{diag}(X) = \mathbf{1}, X \succeq 0. \quad (\text{SDP})$$

In this paper, we give a rigorous proof for this observation, improving on the previous best result which only handles $\sigma = \mathcal{O}(n^{1/4})$ [2]. Our result also provides some justification for the analytical prediction in [19] on optimality of the semidefinite relaxation approach.³

Theorem 1. *If $\sigma = \mathcal{O}\left(\sqrt{\frac{n}{\log n}}\right)$, with high probability for large n , the semidefinite program (SDP) admits a unique solution $\hat{x}\hat{x}^*$, where \hat{x} is a global optimum of (P) (unique up to phase.)*

In the theorem statement, “up to phase” refers to the fact that the measurements C are relative: the distribution of C does not change if z is replaced by $ze^{i\theta}$ for any angle θ , so that if \hat{x} is a solution of (P), then necessarily so is $\hat{x}e^{i\theta}$ for any θ , and z can only be recovered up to a global phase shift.

Theorem 1 shows that the non-convex problem (P) enjoys what is sometimes called *hidden convexity*, that is, in the proper noise regime, it is equivalent to a (tractable)

¹Since W is Gaussian, an MLE minimizes the squared Frobenius norm: $\|C - xx^*\|_F^2 = \|C\|_F^2 + \|xx^*\|_F^2 - 2x^*Cx$. Owing to $|x_k| = 1 \forall k$, this is equivalent to maximizing x^*Cx .

²The notation $\tilde{\mathcal{O}}$ suppresses potential log factors.

³To be precise, in [19], it is predicted—but not proved—via statistical mechanics arguments that the SDP relaxation is nearly optimal when $\sigma = \mathcal{O}(\sqrt{n})$. Instead of showing a solution of (SDP) has rank one, a rescaled leading eigenvector of its solution is used as an estimator.

Algorithm 1 Generalized power method (GPM) without shift

- 1: **Input:** Hermitian measurement matrix $C \in \mathbb{C}^{n \times n}$.
- 2: **Initialize:** Set x^0 to be a leading eigenvector of C with $\|x^0\|_2 = \sqrt{n}$.
- 3: **for** $t = 1, 2 \dots$ **do**
- 4: $x^t = \mathcal{P}(Cx^{t-1})$ \triangleright Apply entrywise $\mathcal{P}(a) = \begin{cases} a/|a| & \text{if } a \neq 0, \\ 1 & \text{otherwise.} \end{cases}$
- 5: **end for**

convex problem. As a consequence, it is not a hard problem in that regime, suggesting local solvers may be able to solve it in its natural dimension. This is desirable, since the relaxation (SDP), while convex, has the disadvantage of lifting the problem from $n - 1$ to $n(n - 1)$ dimensions.

And indeed, numerical experiments in [8] suggest that local optimization algorithms applied to (P) directly succeed in the same regime as (SDP). This was confirmed theoretically in [8] for $\sigma = \mathcal{O}(n^{1/6})$, using both a modification of the power method called the generalized power method (GPM) and local optimization algorithms acting directly on the search space of (P), which is a manifold. Results pertaining to GPM have been rapidly improved to allow for $\sigma = \mathcal{O}(n^{1/4})$ in [23].

In this paper, we consider a version of GPM listed as Algorithm 1 and prove that it works in the same regime as the semidefinite relaxation, thus better capturing the empirical observation. Note that GPM, as a local algorithm, is a more desirable approach versus semidefinite relaxation in practice. GPM and its variants are also considered in a number of related problems [21, 15, 31, 11], and can be seen as special cases of the conditional gradient algorithm [21, 24].

Theorem 2. *If $\sigma = \mathcal{O}\left(\sqrt{\frac{n}{\log n}}\right)$, with high probability for large n , Algorithm 1 converges at least linearly to the global optimum of (P) (unique up to phase.)*

To establish both results, we develop an original proof technique based on following $n + 1$ separate but closely related sequences of feasible points for (P), designed so that they will have suitable statistical independence properties. Furthermore, as a necessary step toward proving the main theorems, we prove an ℓ_∞ perturbation bound for eigenvectors, which is of independent interest.

It is worth noting that for $\sigma > \sqrt{n}$, it is impossible to reliably detect, with probability tending to 1, whether C is of the form $zz^* + \sigma W$ or if it is only of the form σW [28, Thm. 6.11], which suggests that $\sigma < \sqrt{n}$ is necessary in order for a good estimator to exist. This can be made precise by considering the simpler problem of \mathbb{Z}_2 synchronization,⁴ where we have the stronger knowledge that $z_k \in \{\pm 1\}$. For the \mathbb{Z}_2 synchronization problem, non-rigorous arguments that use techniques from statistical

⁴The problem is formulated as follows: $z \in \{-1, 1\}^n$, noise W is a real random matrix, e.g., a Gaussian Wigner matrix, and the goal is to recover z from $C = zz^T + \sigma W$.

mechanics show $\sigma = \sqrt{n}$ is the information-theoretic threshold for mean squared estimation error (MSE): when σ is above this threshold, no estimator is able to beat the trivial estimator $x = 0$ as $n \rightarrow \infty$ [19]. In [22], it was rigorously proved that $\sigma = \sqrt{n}$ is the threshold for a different notion of MSE. These results a fortiori imply, for phase synchronization, that $\sigma = \mathcal{O}(\sqrt{n})$ is necessary⁵ in order for an estimator to have non-trivial MSE (better than the trivial estimator $x = 0$). It is also known that both the eigenvector estimator and the MLE have nontrivial MSE as soon as $\sigma < \sqrt{n}$ [10, 5, 19]. Whether the extra logarithmic factor is necessary to compute the MLE efficiently up to the threshold remains to be determined.

To close this introduction, we state the relevance of the MLE \hat{x} as an estimator for z .

Theorem 3. *Let \hat{x} be a global optimum of (P), with global phase such that $z^* \hat{x} = |z^* \hat{x}|$. Then, deterministically,*

$$\|\hat{x} - z\|_2 \leq 4\sigma \frac{\|W\|_2}{\sqrt{n}}. \quad (3)$$

Furthermore, if $\sigma = \mathcal{O}(\sqrt{n/\log n})$, then with high probability for large n ,

$$\|\hat{x} - z\|_2 = \mathcal{O}(\sigma), \text{ and} \quad (4)$$

$$\|\hat{x} - z\|_\infty = \mathcal{O}(\sigma \sqrt{\log n/n}). \quad (5)$$

The bound on ℓ_2 error appears in [2, Lem. 4.1], while the bound on ℓ_∞ error improves on [2, Lem. 4.2] as a by-product of the results obtained here. We remark that $\sigma = \mathcal{O}(\sqrt{n/\log n})$ is necessary for a nontrivial ℓ_∞ error (smaller than 1, which is trivially attained by $x = 0$) due to [2, 1].⁶

It is important to state that the eigenvector estimator mentioned above is order-wise as good an estimator as the MLE, in that it satisfies the same error bounds as in Theorem 3 up to constants. From the perspective of optimization, the main merit of Theorems 1 and 2 is that they rigorously explain the empirically observed tractability of (P) despite non-convexity.

The difficulty: statistical dependence

As will be argued momentarily, the main difficulty in the analysis is proving a sharp bound for $\|W\hat{x}\|_\infty$, which involves two dependent random quantities: the noise matrix W and a solution \hat{x} of (P), which is a nontrivial function of W . While in the ℓ_2 -norm

⁵Suppose the prior is supported and uniformly distributed on $\{-1, 1\}^n$. By independence, the Bayes-optimal estimator for phase synchronization is a function of $\text{Re}(C)$, so we can use information-theoretic results about the \mathbb{Z}_2 synchronization problem.

⁶In [2, 1], it is suggested information-theoretically exact recovery with high probability is impossible for synchronization over \mathbb{Z}_2 if $\sigma > \sqrt{n/(2-\varepsilon)\log n}$. We can use this result to show $\mathcal{O}(\sqrt{n/\log n})$ is necessary for a nontrivial ℓ_∞ error by putting a uniform prior on $\{-1, 1\}^n$.

the simple bound $\|W\hat{x}\|_2 \leq \|W\|_2 \|\hat{x}\|_2$ is sharp, no such simple argument is known to bound the ℓ_∞ -norm. The need to study perturbations in ℓ_∞ -norm appears inescapable, as it arises from the entry-wise constraints of (P) and the aim to control $\|\hat{x} - z\|_\infty$ as well as $\|\hat{x} - z\|_2$.

This issue has already been raised in [2, §4], which focuses on the relaxation (SDP). Specifically, in [2, eq. (4.10)], it is shown that the relaxation is tight in particular if

$$n - 216\sigma^2 - 3\sigma\sqrt{n} - \sigma\|W\hat{x}\|_\infty > 0, \quad (6)$$

where \hat{x} is an optimum of (P) which is then unique up to phase. From this expression, it is apparent that if $\sigma = \mathcal{O}(\sqrt{n/\log n})$, it only remains to show that $\|W\hat{x}\|_\infty = \mathcal{O}(\sqrt{n \log n})$ to conclude that solving (SDP) is equivalent to solving (P). This reduces the task to that of carefully bounding this scalar, random variable:

$$\|W\hat{x}\|_\infty = \max_{k \in [n]} |w_k^* \hat{x}|, \quad (7)$$

where $w_1, \dots, w_n \in \mathbb{C}^n$ are the columns of the random noise matrix W . If W and \hat{x} were statistically independent, this would be bounded with high probability by $\mathcal{O}(\sqrt{n \log n})$, as desired. Indeed, since the vector \hat{x} contains only phases and since the Gaussian distribution is isotropic (the distribution is invariant under rotation in the complex plane), $w_k^* \hat{x}$ would be distributed identically to a sum of $n - 1$ independent standard complex Gaussians. The modulus of such a variable concentrates close to \sqrt{n} . Taking the maximum over $k \in [n]$ incurs an additional $\mathcal{O}(\sqrt{\log n})$ factor.

Unfortunately, the intricate dependence between W and \hat{x} has not been satisfactorily resolved in previous work, where only suboptimal bounds have been produced for $\|W\hat{x}\|_\infty$, eventually leading to suboptimal bounds on the acceptable noise levels σ [2, eq. (4.11)][8, Lemma 12][23, Proof of Thm. 2].

As a key step to overcome this difficulty, we (theoretically) introduce auxiliary problems to transform the question of controlling $\|W\hat{x}\|_\infty$ into one about the sensitivity of the optimum \hat{x} to perturbations of the data C . This is outlined next.

Introducing auxiliary problems to reduce dependence

Since the main concern in controlling $\|W\hat{x}\|_\infty$ is the statistical dependence between W and \hat{x} , we introduce n new optimization problems of the form (P), where, for each value of m in $[n]$, the cost matrix C is replaced by

$$C^{(m)} = zz^* + \sigma W^{(m)}, \quad \text{with} \quad W_{k\ell}^{(m)} = W_{k\ell} \mathbf{1}_{\{k \neq m\}} \mathbf{1}_{\{\ell \neq m\}}, \quad (8)$$

where $\mathbf{1}$ is the indicator function. In other terms, $W^{(m)}$ is W with the m th row and column set to 0, so that $C^{(m)}$ is statistically independent from w_m . As a result, a global optimum $\hat{x}^{(m)}$ of (P) with C set to $C^{(m)}$ is also independent from w_m . This usefully informs the following observation, where the global phases of \hat{x} and $\hat{x}^{(m)}$ are chosen so

that $\hat{x}^* \hat{x}^{(m)} = |\hat{x}^* \hat{x}^{(m)}|$:

$$\begin{aligned} |(W\hat{x})_m| &= |w_m^* \hat{x}| \leq |w_m^* \hat{x}^{(m)}| + |w_m^* (\hat{x} - \hat{x}^{(m)})| \\ &\leq |w_m^* \hat{x}^{(m)}| + \|w_m\|_2 \|\hat{x} - \hat{x}^{(m)}\|_2. \end{aligned} \quad (9)$$

Crucially, independence of w_m and $\hat{x}^{(m)}$ implies the first term is $\mathcal{O}(\sqrt{n \log n})$ with high probability, by the argument laid out after eq. (7). In the second term, a standard concentration argument shows $\|w_m\|_2 = \mathcal{O}(\sqrt{n})$ with high probability—see Section 3. Hence, to control $\|W\hat{x}\|_\infty$, it is sufficient to show that, with high probability for all m , the solutions \hat{x} and $\hat{x}^{(m)}$ are within distance $\mathcal{O}(1)$ of each other, in the ℓ_2 sense.

This claim about the proximity of \hat{x} and $\hat{x}^{(m)}$ turns out to be a delicate statement about the sensitivity of the global optimum of (P) to perturbations of only measurements which involve the m th phase, z_m . To establish it, we need precise control of the properties of the optima of (P). To this end, we develop a strategy to track the properties of sequences which converge to \hat{x} as well as to $\hat{x}^{(m)}$ for each m .

To ease further discussion about ℓ_2 distances up to phase, consider the following distance-like function:

$$d_2(x, y) := \min_{\theta \in \mathbb{R}} \|xe^{i\theta} - y\|_2, \quad x, y \in \mathbb{C}^n.$$

Restricted to complex vectors of given ℓ_2 -norm or to complex vectors with unit-modulus entries, d_2 is a true metric on the quotient space induced by the equivalence relation \sim :

$$x \sim y \iff \exists \theta : x = ye^{i\theta}. \quad (10)$$

Thus, d_2 is appropriate as a distance between estimators for (P) and as a distance between candidate eigenvectors, being invariant under global phase shifts. Moreover, the quotient space is a complete metric space with d_2 . More details will follow.

Coming back to our problem, the core argument is an analysis of recursive error bounds of GPM, and this analysis leads to the proof that all iterates stay in $\mathcal{N} := \mathcal{N}_1 \cap \mathcal{N}_2$, where

$$\mathcal{N}_1 = \{x \in \mathbb{C}^n : \|Wx\|_\infty \leq \kappa_2 \sqrt{n \log n}\}, \quad (11)$$

$$\mathcal{N}_2 = \{x \in \mathbb{C}^n : d_2(x, z) \leq \kappa_3 \sqrt{n}\}, \quad (12)$$

and $\kappa_2, \kappa_3 > 0$ are some constants (determined in Section 4). On one hand, we show that, with high probability, the nonlinear mapping $\mathcal{T}x = \mathcal{P}(Cx)$ iterated by GPM is Lipschitz continuous over \mathcal{N} with constant $\rho \in (0, 1)$. On the other hand, we show that, with high probability, all iterates of GPM are in \mathcal{N} . Together, these two properties imply that, with high probability, \mathcal{T} is a contraction mapping over the set of iterates of GPM. By a completeness argument, this implies that the sequence of iterates of GPM converges in \mathcal{N} . The roadmap of our proof is the following:

1. $x^0 \in \mathcal{N}$ (this requires developing new ℓ_∞ bounds for eigenvector perturbation—see Theorem 8),

2. $x^t \in \mathcal{N} \implies x^{t+1} \in \mathcal{N}$ (this is done in two stages,⁷ for small and large t —see Theorems 15 and 17),
3. \mathcal{T} is ρ -Lipschitz with $\rho < 1$ on \mathcal{N} with respect to d_2 (by completeness, this implies $\lim_{t \rightarrow \infty} x^t = x^\infty \in \mathcal{N}$ —see Lemma 14), and
4. any fixed point x^∞ of \mathcal{T} in \mathcal{N} is a global optimum of (P) (see Lemma 18).

On top of securing results about GPM, this will imply that (P) admits a solution $\hat{x} = x^\infty$ which is in \mathcal{N} and hence, a fortiori, satisfies $\|W\hat{x}\|_\infty \leq \kappa_2 \sqrt{n \log n}$, yielding the announced results about the SDP relaxation as per (6).

As hinted above, we follow this reasoning not only for the sequence x^t which is expected to converge to \hat{x} , but also for auxiliary sequences $x^{t,m}$ expected to converge to $\hat{x}^{(m)}$. It is only through exploitation of the strong links between these sequences and reduction in statistical dependence they offer that we are able to go through with the proof program above.

Note that \mathcal{T} might not be a contraction mapping on all of \mathcal{N} since we do not show that $\mathcal{T}(\mathcal{N}) \subset \mathcal{N}$. Nevertheless, \mathcal{T} is a contraction on the iterates, which is sufficient for our purpose; henceforth, we say the mapping has the *local contraction property*.

We remark that, in the study of high-dimensional M -estimation [4], the idea of introducing auxiliary problems (and associated optimizers) is also used to tackle dependence, and it yields powerful analysis. While sharing similarity with that approach, our analysis relies on studying n auxiliary sequences of iterates, as will be discussed soon—also see Figure 1.

As a necessary and useful warm-up, we first focus on the task of showing that x^0 (a leading eigenvector of C) is in \mathcal{N} , via analysis of the related $x^{0,m}$ (leading eigenvectors of $C^{(m)}$). This requires sharp bounds for $d_2(x^0, x^{0,m})$. The outcome of this analysis is an eigenvector perturbation bound in the ℓ_∞ -norm, which is another motivation for the introduction of auxiliary problems.

First analysis: an ℓ_∞ perturbation bound for eigenvectors

As the initializer of Algorithm 1, the leading eigenvector of C has several good properties necessary for analysis, and we will discuss them in depth in Section 3. Theorems and lemmas in this direction are stated separately and proved first, because their proof is illustrative of the techniques deployed to prove results about (P). Most notably, we prove a sharp ℓ_∞ perturbation bound for leading eigenvectors.

Theorem 4. *If $\sigma = \mathcal{O}(\sqrt{n/\log n})$, then, with high probability for large n , a leading eigenvector x^0 of the data matrix C (2) scaled such that $\|x^0\|_2 = \sqrt{n}$ satisfies*

$$\|x^0 - z\|_\infty = \mathcal{O}(\sigma \sqrt{\log n/n}), \quad (13)$$

⁷To be precise, for large t , we consider a slightly larger \mathcal{N} (the constants in (11) and (12) are larger), and prove that all x^t stay in this larger region. This is a technical issue which does not affect the overall plan.

where the global phase of x^0 is suitably chosen (e.g., such that $z^*x^0 = |z^*x^0|$.)

The crux of the proof lies in a sharp bound on $d_2(x^0, x^{0,m})$. This is obtained by using (a suitable version of) the Davis–Kahan theorem (Lemma 11): when $\sigma \ll \sqrt{n}$,

$$\begin{aligned} d_2(x^0, x^{0,m}) &= \mathcal{O}\left(\frac{\sigma\|(W - W^{(m)})x^{0,m}\|_2}{n}\right) = \mathcal{O}(\sigma\sqrt{\log n/n}), \quad \text{whereas} \\ d_2(x^0, z) &= \mathcal{O}\left(\frac{\sigma\|Wz\|_2}{n}\right) = \mathcal{O}(\sigma). \end{aligned}$$

To reach the first conclusion, we view x^0 as the perturbed version of $x^{0,m}$ due to perturbation $\sigma(W - W^{(m)})$. Note that $W - W^{(m)}$ has nonzero entries only in the m th row and m th column, and they are independent of $x^{0,m}$. Compared to the full perturbation σW which perturbs the eigenvector z to x^0 , the matrix $\sigma(W - W^{(m)})$ results in a much smaller d_2 distance between $x^{0,m}$ and x^0 . Notice that, as will be detailed later, these results combined with the reasoning of (9) imply that x^0 is in \mathcal{N} , as desired.

Comparing Theorem 4 to Theorem 3 readily shows that the eigenvector x^0 is an excellent estimator for z (up to the fact that its entries are not necessarily unit-modulus, which can be easily corrected—see Theorem 8). Further efforts in this paper are dedicated to characterizing the performance and tractability of the MLE \hat{x} .

Analysis of iterations: tracking n auxiliary sequences

While analyzing the eigenvector x^0 is relatively straightforward, the optima of (P) are more difficult to tame due to the unit-modulus constraints. As hinted above, the novel idea we develop in this paper is to track the sequences $\{x^{t,m}\}_{t=0}^\infty$ produced by Algorithm 1 with inputs $C^{(m)}$ (8) instead of C , for each $m \in [n]$. These auxiliary sequences—which only serve for the analysis and are not (and could not be) computed in practice—enjoy the crucial proximity property desired in the previous subsection—see Figure 1.

Indeed, we will show by induction that there exist absolute constants $\kappa_1, \kappa_2, \kappa_3$ such that, for all m and for $t = 0, 1, 2, \dots$,

1. $d_2(x^t, x^{t,m}) \leq \kappa_1$, proximity property
 2. $x^t \in \mathcal{N}_1 = \{x \in \mathbb{C}^n : \|Wx\|_\infty \leq \kappa_2\sqrt{n \log n}\},$
 3. $x^t \in \mathcal{N}_2 = \{x \in \mathbb{C}^n : d_2(x, z) \leq \kappa_3\sqrt{n}\}.$
- } contraction region

The proximity and local contraction properties⁸ are both crucial and complementary for the analysis: the proximity property allows to control ℓ_∞ quantities in the presence of the random matrix W despite statistical dependence (as shown in (9)), and the local

⁸Formally, we only prove $x^t \in \mathcal{N}$; but proximity implies $x^{t,m}$ is also in a (slightly larger) contraction region, with different constants κ_2, κ_3 .




Figure 1: Sequence $\{x^t\}_{t=0}^\infty$ (in black) produced by Algorithm 1, and n auxiliary sequences $\{x^{t,m}\}_{t=0}^\infty$ (in red) produced (conceptually) by Algorithm 1 with modified inputs. Crucial properties along the paths: (i) proximity: x^t and $x^{t,m}$ stay close; (ii) local contraction: x^t and $x^{t,m}$ remain in the contraction region \mathcal{N} with high probability and converge in it.

contraction property is used to establish (14) below, making sure $d_2(x^t, x^{t,m})$ remains small.

For the high-level idea, consider the nonlinear operators \mathcal{T} and $\mathcal{T}^{(m)}$ implicitly defined by Algorithm 1 so that $x^{t+1} = \mathcal{T}x^t$ and $x^{t+1,m} = \mathcal{T}^{(m)}x^{t,m}$. If we can show that \mathcal{T} is ρ -Lipschitz with constant $\rho \in (0, 1)$ with respect to d_2 , then a recursive error bound follows:

$$\begin{aligned} d_2(x^{t+1}, x^{t+1,m}) &= d_2(\mathcal{T}x^t, \mathcal{T}^{(m)}x^{t,m}) \\ &\leq d_2(\mathcal{T}x^t, \mathcal{T}x^{t,m}) + d_2(\mathcal{T}x^{t,m}, \mathcal{T}^{(m)}x^{t,m}) \\ &\leq \rho \cdot d_2(x^t, x^{t,m}) + \text{discrepancy error}. \end{aligned} \quad (14)$$

This ensures $d_2(x^t, x^{t,m})$ does not accumulate with t , provided the discrepancy error—which is caused by the difference between \mathcal{T} and $\mathcal{T}^{(m)}$ —is small enough. This is assured with high probability, because $C - C^{(m)}$ is independent of $x^{t,m}$, causing the discrepancy error to be $\mathcal{O}(\sigma\sqrt{\log n/n})$ —considerably smaller than $d_2(x^{t+1}, z) = \mathcal{O}(\sigma)$. In spirit, this is the same argument as in the analysis of the eigenvector estimator.

The above recursive error bound hinges on the other important property, that is, x^t staying in the contraction region. Crucially, to establish $x^t \in \mathcal{N}_1$, we need a tight bound on $\|Wx^t\|_\infty$. Fortunately, we have seen how to control this quantity in (9): for any $m \in [n]$,

$$|(Wx^t)_m| \leq |w_m^* x^{t,m}| + \|w_m\|_2 \cdot d_2(x^t, x^{t,m}). \quad (15)$$

This, in turn, requires a proximity result for $d_2(x^t, x^{t,m})$. This insight naturally motivates an analysis of each iteration by induction.

There are two technical issues we briefly address before ending this introduction with a remark.

The first issue concerns the probabilistic argument in the proof. In (14) and (15), we invoke concentration inequalities to obtain tight bounds. However, since there is a (small) probability that such inequalities fail, we cannot use union bounds for $t = 1, 2, \dots$, which is infinite. To overcome this obstacle, we use concentration inequalities only for the first $T = \mathcal{O}(n^2)$ iterations, and resort to a deterministic analysis for iterations $t > T$. The critical observation is that, $d_2(x^{t+1}, x^t)$ decays exponentially for $t \leq T$ due to contraction, so the amount of update is tiny after T iterations. Using another inductive argument, we can secure exponential decay for $t > T$ as well. The rationale is that we already established good properties about x^T , and $d_2(x^t, x^T)$ is tiny for $t > T$, so we can easily relate x^t to x^T and show x^t also has good properties. Essentially, x^t remains in a contraction region with slightly larger constants.

The second issue is identifying the limit $x^\infty = \lim_{t \rightarrow \infty} x^t$ with a solution of (P). We will verify the optimality and uniqueness (up to phase) of x^∞ via a known dual optimality certificate $S = \text{Re}\{\text{ddiag}(Cx^\infty(x^\infty)^*)\} - C$ [2].

We close with a remark about the initializer x^0 (and $x^{0,m}$). Algorithm 1 uses the leading eigenvector of C for initialization, and our analysis relies on ℓ_∞ perturbation bounds to verify the base case of the induction. However, we point out that, even in the absence of such perturbation results, we could set $x^0 = x^{0,m} = z$ in theory, deduce that $\hat{x} \in \mathcal{N}$ and thus prove Theorem 1 about the SDP relaxation. In other words, even if we leave out the discussion about the initializer altogether, there is still enough material to secure tightness of the SDP relaxation. The proof augmented with the analysis of the eigenvector perturbation has the advantage of also providing a statement about Algorithm 1 which is actually runnable in practice.

2 Main results

In this section we will state our main theorems formally. The assumption on random noise W will also be relaxed to a broader class of random matrices. To begin with, let us first clarify the “up to phase” statements in Section 1.

The quotient space For any $\theta \in \mathbb{R}$, whether the true phases are z or $ze^{i\theta}$ does not affect the measurements C (2). As a result, the available data are insufficient to distinguish z from $ze^{i\theta}$. Clearly the program (P) is invariant to global phase shifts as well. It thus makes sense to ignore the global phase in defining distances between estimators. A reasonable notion of ℓ_2 error then becomes

$$d_2(x, y) = \min_{\theta \in \mathbb{R}} \|xe^{i\theta} - y\|_2 = \sqrt{\|x\|_2^2 + \|y\|_2^2 - 2|x^*y|}, \quad (16)$$

where the optimal phase $e^{i\theta}$ is the phase (Arg) of x^*y . Similarly, a notion of ℓ_∞ error can be defined:

$$d_\infty(x, z) = \min_{\theta \in \mathbb{R}} \|xe^{i\theta} - z\|_\infty. \quad (17)$$

Formally, one can partition all points in \mathbb{C}^n into equivalence classes via the equivalence relation \sim (10). The resulting quotient space \mathbb{C}^n / \sim contains the equivalence classes $[x] = \{xe^{i\theta} : \theta \in \mathbb{R}\}$ for all $x \in \mathbb{C}^n$. Specifically, the feasible set of (P),

$$\mathbb{C}_1^n := \{x \in \mathbb{C}^n : |x_1| = \dots = |x_n| = 1\}, \quad (18)$$

reduces to \mathbb{C}_1^n / \sim under this equivalence relation. It is easily verified that d_2 defines a distance on \mathbb{C}_1^n / \sim . In particular, it satisfies the triangular inequality (where $d_2(x, y)$ is understood to mean $d_2([x], [y])$):

$$\begin{aligned} \forall x, y, z \in \mathbb{C}_1^n, \quad d_2(x, y) &= \min_{\theta_1, \theta_2} \|xe^{i\theta_1} - z + z - ye^{i\theta_2}\|_2 \\ &\leq \min_{\theta_1} \|xe^{i\theta_1} - z\|_2 + \min_{\theta_2} \|z - ye^{i\theta_2}\|_2 = d_2(x, z) + d_2(z, y). \end{aligned}$$

Moreover, \mathbb{C}_1^n / \sim is a complete metric space under d_2 (see Theorem 17). Similarly, d_∞ is also a distance. As will be shown, the sequence $\{x^t\}_{t=1}^\infty$ described in Section 1 satisfies the local contraction property (see (14)) on the metric space $(\mathbb{C}_1^n / \sim, d_2)$, hence converges to a fixed point which is exactly \hat{x} (understood as $[\hat{x}]$).

The noise matrix In Section 1 we assume that W has independent standard complex Gaussian variables above its diagonal. However, this restricted assumption is only for expository convenience, and can be relaxed to the class of Hermitian Wigner matrices with sub-gaussian entries. Statements about Algorithm 1 and about tightness of the SDP relaxation continue to hold, although of course the solution of (P) now no longer necessarily corresponds to the MLE.

The class of sub-gaussian variables subsumes Gaussian variables, but has one defining feature similar to Gaussian variables, that is, the tail probability decaying no slower than Gaussian variables. In our model, each entry of W satisfies the tail bound

$$\mathbb{P}(|\xi| > t) \leq \exp(1 - t^2/K^2) \quad (19)$$

for both real and imaginary parts, where $K > 0$ is an absolute constant. Formally, we assume the Hermitian matrix W satisfies the following: $\{\text{Re}(W_{k\ell}), \text{Im}(W_{k\ell})\}_{k < \ell}$ are jointly independent, have zero mean, and satisfy the sub-gaussian tail bound (19); the diagonal elements are zero, and $W_{\ell k} = \overline{W_{k\ell}}$ for any $k < \ell$. Note there are equivalent definitions of sub-gaussian variables (up to constants) [36].

This random model is a much richer class of noise matrices, containing the Gaussian model introduced in Section 1 as a special case. Each random variable in $\{\text{Re}(W_{k\ell}), \text{Im}(W_{k\ell})\}_{k \leq \ell}$ can be, for example, a symmetric Bernoulli variable, any other centered and bounded variable, or simply zero.

SDP approach The SDP approach tries to solve (P) via its convex semidefinite relaxation (SDP). It is a relaxation of (P) in the following sense. For any feasible $x \in \mathbb{C}^n$, the corresponding matrix $X = xx^*$ is feasible for (SDP). Likewise, any feasible matrix X of rank 1 can be factored as $X = xx^*$ such that x is feasible for (P). Thus, the relaxation consists in allowing solutions of rank more than 1 in (SDP). Consequently, if (SDP) admits a solution of rank 1, $X = \hat{x}\hat{x}^*$, then the corresponding \hat{x} is a global optimum for (P). Furthermore, if the rank-1 solution of (SDP) is unique, it can be recovered in polynomial time. For this reason, the regime of interest is one where (SDP) admits a unique solution of rank 1.

The following theorem—a statement of Theorem 1 which holds in the broadened noise model—closes the gap in previous papers [3, 2, 23, 8].

Theorem 5. *There exists an absolute constant $c_0 > 0$ such that, if $\sigma \leq c_0 \sqrt{n/\log n}$, then, with probability $1 - \mathcal{O}(n^{-2})$, the semidefinite program (SDP) admits a unique solution $\hat{x}\hat{x}^*$ where \hat{x} is the unique global optimum of (P) up to phase.*

Note that the exponent 2 in the failure probability $\mathcal{O}(n^{-2})$ can be replaced by any positive numerical constant, only affecting other absolute constants in the theorem (and all other theorems).

GPM approach The generalized power method (Algorithm 1) is an iterative algorithm similar to the classical power method, but instead of projecting vectors onto a sphere after matrix-vector multiplication Cx^{t-1} , it extracts the phases from Cx^{t-1} , which is an entry-wise projection. It is much faster than SDP, and converges linearly to a limit, which is the optimum \hat{x} up to phase (optimality is stated in Theorem 7). The next theorem is a precise version of Theorem 2.

Theorem 6. *There exists an absolute constant $c_0 > 0$ such that, if $\sigma \leq c_0 \sqrt{n/\log n}$, then, with probability $1 - \mathcal{O}(n^{-2})$, the sequence $\{x^t\}$ produced by Algorithm 1 has a linear convergence rate to some $\hat{x} \in \mathbb{C}_1^n$ (up to phase):*

$$d_2(x^t, \hat{x}) \leq 2^{2^{-t}} \sqrt{n}, \quad t = 0, 1, 2, \dots$$

The proof is based on induction: in each iteration, we will establish the proximity property and the contraction property for x^t . The proof is simply a rigorous justification of the heuristics we discussed in Section 1. It also leads to ℓ_2 and ℓ_∞ error bounds for \hat{x} . The following theorem is a precise version of Theorem 3.

Theorem 7. *Under the same conditions and notation as in Theorem 6, with probability $1 - \mathcal{O}(n^{-2})$, \hat{x} is the unique optimum of (P) up to phase. If we choose the global phase of \hat{x} such that $\hat{x}^*z = |\hat{x}^*z|$, then*

$$d_2(\hat{x}, z) = \|\hat{x} - z\|_2 \leq C_0 \sigma, \text{ and } d_\infty(\hat{x}, z) = \|\hat{x} - z\|_\infty \leq C_0 \sigma \sqrt{\log n/n},$$

where $C_0 > 0$ is an absolute constant.

Eigenvector estimator We denote, henceforth, the leading eigenvector of C by \tilde{x} , and similarly the leading eigenvector of $C^{(m)}$ by $\tilde{x}^{(m)}$. Note that \tilde{x} and x^0 (similarly $\tilde{x}^{(m)}$ and $x^{0,m}$) are identical.⁹ We highlight the significance of the eigenvector estimator in the following theorem, which is a precise version of Theorem 4.

Theorem 8. *Let \tilde{x} be a leading eigenvector of C with scale and global phase chosen such that $\|\tilde{x}\|_2 = \sqrt{n}$ and $\tilde{x}^*z = |\tilde{x}^*z|$. There exists an absolute constant $c'_0 > 0$ such that, if $\sigma < c'_0\sqrt{n/\log n}$, then, with probability $1 - \mathcal{O}(n^{-2})$,*

$$d_2(\tilde{x}, z) = \|\tilde{x} - z\|_2 \leq C'_0\sigma, \text{ and } d_\infty(\tilde{x}, z) \leq \|\tilde{x} - z\|_\infty \leq C'_0\sigma\sqrt{\log n/n},$$

where $C'_0 > 0$ is some absolute constant. Moreover, the projected leading eigenvector, namely, $\mathcal{P}\tilde{x}$, satisfies the same bounds with C'_0 replaced by $2C'_0$.

The eigenvector estimator has been studied extensively in recent years, prominently in the statistics literature [27, 20], under the *spiked covariance model*. While the perturbation $\tilde{x} - z$ is usually studied under ℓ_2 norms, the ℓ_∞ norm received much less attention. A recent ℓ_∞ perturbation result appeared in [16], but it is a deterministic bound and would produce a suboptimal result here.

3 Proof organization for eigenvector perturbations

We begin with some concentration lemmas, which will also be useful in Section 4. Recall the definition of $W^{(m)}$ (8). We also define $\Delta W^{(m)} := W - W^{(m)}$, which has nonzero entries only in the m th row and m th column, given by w_m .

Concentration lemmas

The first concentration result is standard and is a direct consequence of, for example, Proposition 2.4 in [32].

Lemma 9. *With probability $1 - \mathcal{O}(n^{-2})$, the following holds for any $m \in [n]$:*

$$\|W\|_2 \leq C'_2\sqrt{n}, \quad \|W^{(m)}\|_2 \leq C'_2\sqrt{n}, \quad \|\Delta W^{(m)}\|_2 \leq C'_2\sqrt{n}, \quad \|w_m\|_2 \leq C'_2\sqrt{n}, \quad (20)$$

where $C'_2 > 0$ is an absolute constant.

Let \mathbb{S}^{n-1} be the set of unit vectors in \mathbb{C}^n . Suppose for each $m \in [n]$ we have a finite (random) set $\mathcal{U}^{(m)} \subset \sqrt{n}\mathbb{S}^{n-1}$ whose elements are independent of $\Delta W^{(m)}$, and the cardinality of $\mathcal{U}^{(m)}$ is not random. Concentration inequalities enable us to bound $\|\Delta W^{(m)}u\|_2$ uniformly over all $u \in \mathcal{U}^{(m)}$ with high probability. We state this formally in the next lemma.

⁹We use notation \tilde{x} (similarly $\tilde{x}^{(m)}$) to emphasize the leading eigenvector as an estimator, as opposed to merely an initializer.

Lemma 10. Suppose $|\mathcal{U}^{(m)}| \leq 3n^2$ for all $m \in [n]$. Let u_m be the m th entry of a vector $u \in \mathcal{U}^{(m)}$ and denote $M_m = \max_{u \in \mathcal{U}^{(m)}} |u_m|$. Then, with probability $1 - \mathcal{O}(n^{-2})$,

$$\max_{u \in \mathcal{U}^{(m)}} \|\Delta W^{(m)} u\|_2 \leq C'_1 \sqrt{n \log n} + C'_1 \sqrt{n} M_m, \quad \forall m \in [n], \quad (21)$$

where $C'_1 > 0$ is an absolute constant. In particular, we can choose $C'_1 > 0$ such that, if $\mathcal{U}^{(m)} \subset \mathbb{C}_1^n$ (defined in (18)), then, with probability $1 - \mathcal{O}(n^{-2})$,

$$\max_{m \in [n]} \max_{u \in \mathcal{U}^{(m)}} \|\Delta W^{(m)} u\|_2 \leq C'_1 \sqrt{n \log n}. \quad (22)$$

A straightforward application of Hoeffding's inequality for sub-gaussian variables [36] shows $\max_{u \in \mathcal{U}^{(m)}} |w_m^* u| = \mathcal{O}(\sqrt{n \log n})$ with probability $1 - \mathcal{O}(n^{-2})$. Lemma 10 is more general, because $|w_m^* u| \leq \|\Delta W^{(m)} u\|_2$, and it will be useful in later proofs.

For the eigenvector problem, we will choose $\mathcal{U}^{(m)} = \{\tilde{u}^{(m)}\}$ (a singleton) where $\tilde{u}^{(m)}$ is a leading eigenvector of $C^{(m)} = zz^* + \sigma W^{(m)}$ scaled to have norm \sqrt{n} . For problem (P), for each $m \in [n]$, the set $\mathcal{U}^{(m)}$ will be $\{x^{t,m}\}_{t=0}^T$, namely, the first $T + 1$ iterates of Algorithm 1 with input $C^{(m)}$, where $T := 3n^2 - 1$. By construction, elements of the set $\mathcal{U}^{(m)}$ are independent of $\Delta W^{(m)}$.

Introducing auxiliary eigenvector problems

As is well known, the leading eigenvectors of C are the solutions to the following optimization problem (note that this problem is a relaxation of (P)):

$$\begin{aligned} \max_{x \in \mathbb{C}^n} \quad & x^* C x \\ \text{s.t.} \quad & \|x\|_2 = \sqrt{n}. \end{aligned} \quad (\tilde{\text{P}})$$

We aim to show that a solution \tilde{x} of $(\tilde{\text{P}})$ is close to z in the sense of d_∞ . As before, the major difficulty of the analysis is obtaining a sharp bound on $\|W\tilde{x}\|_\infty$. This is apparent when we write $\lambda_1(C)$ for the leading eigenvalue of C and use $C\tilde{x} = \lambda_1(C)\tilde{x}$ to obtain (choosing the global phase of \tilde{x} such that $z^*\tilde{x} = |z^*\tilde{x}|$):

$$|\tilde{x}_m - z_m| = \left| \frac{(C\tilde{x})_m}{\lambda_1(C)} - z_m \right| \leq \left| \frac{|z^*\tilde{x}|}{\lambda_1(C)} - 1 \right| + \frac{|\sigma|(W\tilde{x})_m|}{\lambda_1(C)}, \quad \forall m \in [n].$$

While it is easy to analyze $\lambda_1(C)$ and $|z^*\tilde{x}|$, bounding $\|W\tilde{x}\|_\infty$ requires more work.

For $m \in [n]$, let $\tilde{x}^{(m)}$ be the solution to an auxiliary problem $(\tilde{\text{P}})$ in which C is replaced by $C^{(m)} = zz^* + \sigma W^{(m)}$ —thus, $\tilde{x}^{(m)}$ is equivalent to $x^{0,m}$. Following the same strategy as in (9), we can now split $W\tilde{x}$ into two terms and try to bound separately:

$$|(W\tilde{x})_m| = |w_m^* \tilde{x}| \leq |w_m^* \tilde{x}^{(m)}| + \|w_m\|_2 \cdot d_2(\tilde{x}, \tilde{x}^{(m)}) \quad (23)$$

where $|w_m^* \tilde{x}^{(m)}|$ is the dominant term, and can be easily bounded—see the paragraph below Lemma 10; and $\|w_m\|_2 \cdot d_2(\tilde{x}, \tilde{x}^{(m)}) = \mathcal{O}(\sqrt{n} d_2(\tilde{x}, \tilde{x}^{(m)}))$ is the higher-order discrepancy error, which is the price we pay for replacing \tilde{x} with $\tilde{x}^{(m)}$.

The crucial point is that $d_2(\tilde{x}, \tilde{x}^{(m)}) = \mathcal{O}(\sigma\sqrt{\log n/n})$, which is much smaller than $d_2(\tilde{x}, z) = \mathcal{O}(\sigma)$. This is because the difference between \tilde{x} and $\tilde{x}^{(m)}$ results from a sparse perturbation $\Delta W^{(m)}$, whose effect on the leading eigenvector is small. This point is formalized in the next lemma, which follows from [14].¹⁰

Lemma 11 (Davis–Kahan $\sin \Theta$ Theorem). *Suppose that $A, E \in \mathbb{C}^{n \times n}$ are Hermitian matrices, and $\tilde{A} = A + E$. Let $\delta := \lambda_1(A) - \lambda_2(A)$ be the gap between the top two eigenvalues of A , and u, \tilde{u} be leading eigenvectors of A and \tilde{A} respectively, normalized such that $\|u\|_2 = \|\tilde{u}\|_2 = \sqrt{n}$. If $\delta > \|E\|_2$, then*

$$d_2(\tilde{u}, u) \leq \frac{\sqrt{2} \|Eu\|_2}{\delta - \|E\|_2}. \quad (24)$$

The benefit of this perturbation result is pronounced when E is a sparse random matrix: if we set $A = C^{(m)}, E = \sigma \Delta W^{(m)}$, then the numerator in (24) becomes $\sqrt{2} \sigma \|\Delta W^{(m)} \tilde{x}^{(m)}\|_2 = \mathcal{O}(\sigma \sqrt{n \log n})$ with high probability (by Lemma 10 and a bound on M_m). If, however, we set $A = zz^*, E = \sigma W$, then the numerator is $\sqrt{2} \sigma \|Wz\|_2 = \mathcal{O}(\sigma n)$ with high probability. This is why $d_2(\tilde{x}, \tilde{x}^{(m)})$ is so small and (23) yields a tight bound.

We remark that in many later uses of perturbation results (e.g., [37, 29]), especially in statistics and theoretical computer science, it is common to invoke a variant of the Davis–Kahan theorem in which $\|Eu\|_2$ is replaced by $\|E\|_2$ in (24), which would lead to a suboptimal result here. This is because $\|Eu\|_2 \ll \|E\|_2$ with high probability when E is a random sparse matrix and $\|u\|_\infty$ is not large. Our analysis here is an example that shows the merit of using the more precise version of the Davis–Kahan theorem.

4 Proof organization for phase synchronization

We begin with some useful lemmas about the local contraction property. These will prove useful to establish the desired properties of the iterates x^t of Algorithm 1, by induction. These properties extend to the limit $x^\infty = \lim_{t \rightarrow \infty} x^t$ by continuity. Finally, we will use a known optimality certificate for (SDP) to validate x^∞ .

Local contraction lemmas

First let us denote the rescaled matrix C/n by \mathcal{L} , which can also be viewed as a linear operator in \mathbb{R}^n :

$$\mathcal{L}x := \frac{1}{n}Cx = \frac{z^*x}{n}z + \frac{\sigma}{n}Wx, \quad \forall x \in \mathbb{C}^n.$$

This is a linear combination of z (signal) and Wx (noise). When $\sigma\|W\|_2$ is small compared to n , \mathcal{L} is Lipschitz continuous on the sphere around z , with respect to d_2 .

¹⁰By Davis–Kahan Theorem and Weyl’s inequality, $\sin \theta(\tilde{u}, u) \leq \|Eu\|_2 / \sqrt{n} (\delta - \|E\|_2)$. Thus, the lemma follows from $d_2(\tilde{u}, u)^2/n = 2 - 2|\tilde{u}^*u|/n = 2 - 2\cos \theta(\tilde{u}, u) \leq 2\sin^2 \theta(\tilde{u}, u)$.

Lemma 12. Suppose $\varepsilon \in (0, 1/2)$, and $x, y \in \sqrt{n} \mathbb{S}^{n-1}$, with $d_2(x, z) \leq \varepsilon \sqrt{n}$, $d_2(y, z) \leq \varepsilon \sqrt{n}$. Then,

$$d_2(\mathcal{L}x, \mathcal{L}y) \leq \left(6\varepsilon + \frac{\sigma \|W\|_2}{n}\right) d_2(x, y). \quad (25)$$

This lemma is instrumental in establishing the key local contraction property (Lemma 14). It is related to the contraction mapping theorem, in which an iteratively defined sequence converges to a fixed point. We could use this lemma to easily show (using [23, Proof of Thm. 2] for example) that the normalized version $\bar{\mathcal{L}}x := \mathcal{L}x / \|\mathcal{L}x\|_2$ is a contraction mapping in a neighborhood of z on $\sqrt{n} \mathbb{S}^{n-1}$ — $\bar{\mathcal{L}}$ is the power method operator.

However, our problem is more complicated due to the unit-modulus constraints in (P) which call for the entry-wise operator $\mathcal{P}(\cdot)$ in Algorithm 1. Consequently, an analysis of Algorithm 1 requires entry-wise bounds on key quantities in each iteration, which are more involved than ℓ_2 bounds. In the next two lemmas, we will see which entry-wise bounds we need in order to establish the local contraction property.

Recall that $\mathcal{P} : \mathbb{C}^n \rightarrow \mathbb{C}_1^n$ maps each entry of a vector to the unit circle in the complex plane:

$$\forall k \in [n], \quad (\mathcal{P}(x))_k = \begin{cases} x_k / |x_k| & \text{if } x_k \neq 0, \\ 1 & \text{if } x_k = 0. \end{cases}$$

The case $x_k = 0$ will not appear in the proofs because it can be excluded with high probability. Henceforth, we also drop parentheses in $\mathcal{P}(x)$ for simplicity.

In the next two lemmas, we establish the local contraction property of $\mathcal{P}\mathcal{L}$ (the GPM operator) under the distance d_2 . Under certain conditions on the input points, $\mathcal{P}\mathcal{L}$ shrinks the d_2 distance between points by a ratio in $(0, 1)$. In a rigorous sense, this does not imply that $\mathcal{P}\mathcal{L}$ is a contraction mapping, because the output points do not necessarily satisfy the conditions themselves. However, for the sequences of interest, the conditions are satisfied with high probability and this is all we need to ascertain convergence—see Theorems 15 and 17.

Lemma 13. Suppose $\varepsilon \in [0, 1)$. For any $x, y \in \mathbb{C}$, if $|x| \geq 1 - \varepsilon$, $|y| \geq 1 - \varepsilon$, then

$$\left| \frac{x}{|x|} - \frac{y}{|y|} \right| \leq (1 - \varepsilon)^{-1} |x - y|. \quad (26)$$

As a consequence, for any $w, v \in \mathbb{C}^n$ with $\min_k |w_k| \geq 1 - \varepsilon$ and $\min_k |v_k| \geq 1 - \varepsilon$,

$$d_2(\mathcal{P}w, \mathcal{P}v) \leq (1 - \varepsilon)^{-1} d_2(w, v). \quad (27)$$

This lemma says \mathcal{P} is Lipschitz continuous (in the quotient space \mathbb{C}^n / \sim) in a region where $|w_k|$ and $|z_k|$ are uniformly lower bounded. The composition $\mathcal{P}\mathcal{L}$, will have the local contraction property as long as the contraction ratio in (25) is small enough, and the ε in (27) is not too close to 1. The next lemma formalizes this result. For later use in the proofs, we also introduce an additional notation: for a Hermitian matrix $W' \in \mathbb{C}^{n \times n}$, we let $\mathcal{L}' = (zz^* + \sigma W')/n$.

Lemma 14. Suppose $x, y \in \sqrt{n} \mathbb{S}^{n-1}$, $K_1, K_2 > 0$ and $\varepsilon_1, \varepsilon_2 \in (0, 1/2)$ with

$$d_2(x, z) \leq \varepsilon_1 \sqrt{n}, \quad d_2(y, z) \leq \varepsilon_2 \sqrt{n}, \quad (28)$$

$$\|Wx\|_\infty \leq K_1 \sqrt{n \log n}, \quad \|W'y\|_\infty \leq K_2 \sqrt{n \log n}. \quad (29)$$

Let $\varepsilon = \max\{\varepsilon_1^2/2 + K_1\sigma\sqrt{\log n/n}, \varepsilon_2^2/2 + K_2\sigma\sqrt{\log n/n}\}$. If $\varepsilon < 1$, then

$$d_2(\mathcal{P}\mathcal{L}x, \mathcal{P}\mathcal{L}'y) \leq (1 - \varepsilon)^{-1} d_2(\mathcal{L}x, \mathcal{L}'y). \quad (30)$$

In particular, when $\mathcal{L} = \mathcal{L}'$, if $\varepsilon < 1$ we have

$$d_2(\mathcal{P}\mathcal{L}x, \mathcal{P}\mathcal{L}y) \leq \rho \cdot d_2(x, y), \quad (31)$$

where $\rho = (1 - \varepsilon)^{-1}(6 \max\{\varepsilon_1, \varepsilon_2\} + \sigma\|W\|_2/n)$.

This deterministic lemma states that $\mathcal{T} := \mathcal{P}\mathcal{L}$ has the local contraction property in a region where (28) and (29) are satisfied, as long as $\rho < 1$. Note that we have to require ℓ_∞ bounds in (29) because of the entry-wise nature of \mathcal{P} . In the next subsection, we use this lemma to show that, with high probability, $d_2(\mathcal{T}x^t, \mathcal{T}x^{t,m})$ is controlled by $\rho \cdot d_2(x^t, x^{t,m})$ where the ratio ρ lies in $(0, 1)$.

Convergence analysis

Let us denote the rescaled matrix $C^{(m)}/n$ by $\mathcal{L}^{(m)}$, where $m \in [n]$. Also let $\mathcal{T}^{(m)} = \mathcal{P}\mathcal{L}^{(m)}$. Recall that x^0 is identical to the leading eigenvector \tilde{x} of \mathcal{L} , and $x^{0,m}$ is identical to the leading eigenvector $\tilde{x}^{(m)}$ of $\mathcal{L}^{(m)}$; and they are normalized such that $\|x^0\|_2 = \|x^{0,m}\|_2 = \sqrt{n}$. Algorithm 1 iterates $x^{t+1} = \mathcal{T}x^t$. The auxiliary sequences defined for theoretical analysis (not implemented in practice) follow a similar update rule: $x^{t+1,m} = \mathcal{T}^{(m)}x^{t,m}$ (Figure 1). Note that, for all $t \geq 1$, each entry of x^t and $x^{t,m}$ has unit modulus, i.e., $x^t, x^{t,m} \in \mathbb{C}_1^n$, but $x^0, x^{0,m}$ are not in \mathbb{C}_1^n in general.

As shown in Lemma 14, the local contraction property of \mathcal{T} hinges on the condition that the vectors to be updated are in the contraction region $\mathcal{N} = \mathcal{N}_1 \cap \mathcal{N}_2$, where \mathcal{N}_1 and \mathcal{N}_2 are defined in (11) and (12). The absolute constants $\kappa_2, \kappa_3 > 0$ in their definitions will be specified in Theorem 15.

In order to show the iterates x^t stay in \mathcal{N} , we analyze the dependence between the random quantities \mathcal{T} and x^t by making use of the auxiliary sequences $x^{t,m}$, as illustrated by eq. (15). As shown in the analysis of eigenvectors in Section 3, we know $d_2(x^0, x^{0,m})$ is small. Owing to the local contraction property (Lemma 14), we can prove the recursive error bound (14), which ensures that $d_2(x^{t,m}, x^t)$ is bounded throughout all iterations. The analysis is based on induction.

As discussed in Section 1, a technical issue is that we cannot use concentration results infinitely many times for all $t = 1, 2, \dots$, because we use the union bound to achieve a high probability result. We study first $T+1 := 3n^2$ iterates using concentration results, and resort to deterministic analysis for later iterations. In the following theorems, constants C'_1, C'_2 are those constants in Lemmas 9 and 10.

Theorem 15. Suppose $n \geq 2$ and σ satisfies

$$\sigma \leq \min \left\{ \frac{\sqrt{n}}{120\sqrt{2}C'_2}, \frac{1}{240\kappa_2} \sqrt{\frac{n}{\log n}} \right\}. \quad (32)$$

Then, with probability $1 - \mathcal{O}(n^{-2})$, for any t in $\{0, 1, 2, \dots, T\}$,

$$d_2(x^{t,m}, x^t) \leq \kappa_1, \quad \forall m \in [n], \quad (33)$$

$$\|Wx^t\|_\infty \leq \kappa_2 \sqrt{n \log n}, \quad (34)$$

$$d_2(x^t, z) \leq \kappa_3 \sqrt{n}. \quad (35)$$

Here $\kappa_1, \kappa_2, \kappa_3$ are absolute constants: $\kappa_1 = \kappa_3 = 1/60$ and $\kappa_2 = 4C'_1 + 2C'_2\kappa_1$.

Note that (34) and (35) guarantee $\{x^t\}_{t=0}^T \subset \mathcal{N}$. Considering (14) from our proof map,

$$d_2(x^{t+1,m}, x^{t+1}) \leq d_2(\mathcal{T}x^{t,m}, \mathcal{T}x^t) + d_2(\mathcal{T}^{(m)}x^{t,m}, \mathcal{T}x^{t,m}).$$

Assuming (34) and (35) for the case t , by the local contraction property (Lemma 14), the first term is bounded by $\rho \cdot d_2(x^{t,m}, x^t)$ where $\rho < 1$. By the concentration bounds (Lemma 10), the second term is bounded by $\mathcal{O}(\sigma \|\Delta W^{(m)}x^{t,m}\|_2/n) = \mathcal{O}(\sigma \sqrt{\log n/n}) = \mathcal{O}(1)$ with high probability. Therefore, it is expected that (33) continues to hold for the case $t+1$.

This proximity property (33) is crucial to show that x^{t+1} stays in the local region \mathcal{N} . To bound $\|Wx^{t+1}\|_\infty$, we use the concentration bounds (Lemma 10) and the proximity property (33) in the inequality (15).

To bound $d_2(x^{t+1}, z)$, we derive an entry-wise bound on $\mathcal{L}x^t$, then use Lemma 13. This is straightforward once we have a bound on $\|Wx^t\|_\infty$.

The next result says $d_2(x^{t+1}, x^t)$ decreases geometrically for $t = 0, \dots, T-1$, which is notably useful to analyze later iterations ($t \geq T$).

Theorem 16. Under the same assumption as in Theorem 15, with probability $1 - \mathcal{O}(n^{-2})$, we have

$$d_2(x^{t+1}, x^t) \leq \frac{1}{2} d_2(x^t, x^{t-1}), \quad \forall t \in \{1, 2, \dots, T-1\}, \quad (36)$$

and as a consequence, $d_2(x^T, x^{T-1}) \leq 2^{2-T} \sqrt{n}$.

This result is similar to the contraction mapping theorem (though we cannot prove $\mathcal{T}x \in \mathcal{N}$ for all $x \in \mathcal{N}$), which says the sequence produced by a contraction mapping is a Cauchy sequence and satisfies an inequality similar to (36). After $T = 3n^2 - 1$ iterations, the update from x^t to x^{t+1} is almost negligible. Although we no longer rely on concentration results, we will show, by induction again, that $d_2(x^t, x^T)$ remains very small for all $t \geq T$. This ensures that x^t stays in a slightly larger contraction region for all t (with larger constants). The next theorem depends crucially on the conclusions of Theorem 15 and 16.

Theorem 17. Suppose $\|W\|_2 \leq C'_2\sqrt{n}$, $\|Wx^{T-1}\|_\infty \leq \kappa_2\sqrt{n \log n}$, $d_2(x^{T-1}, z) \leq \kappa_3\sqrt{n}$ and σ is bounded as in Theorem 15. Also suppose $d_2(x^T, x^{T-1}) \leq \kappa_3/4$. Then,

$$d_2(x^{T+k}, x^{T+k-1}) \leq 2^{-k} d_2(x^T, x^{T-1}), \quad \forall k \geq 0. \quad (37)$$

Furthermore, in the quotient space \mathbb{C}_1^n / \sim equipped with distance d_2 , the sequence $[x^t]$ ($t \geq 1$) converges to a limit $[x^\infty]$ where $x^\infty \in \mathbb{C}_1^n$ is a fixed point of \mathcal{T} , i.e., $\mathcal{T}x^\infty = x^\infty$. This fixed point satisfies

$$d_2(x^\infty, z) \leq \frac{3}{2}\kappa_3\sqrt{n}, \text{ and } \|Wx^\infty\|_\infty \leq (\kappa_2 + C'_2\kappa_3)\sqrt{n \log n}. \quad (38)$$

Moreover,

$$Cx^\infty = \text{diag}(\mu)x^\infty, \quad (39)$$

where $\mu_k = |(Cx^\infty)_k|$.

Under the stated conditions, this theorem is a deterministic result. By Theorems 15 and 16, the conditions hold with high probability (note that $2^{2-T}\sqrt{n} \leq \kappa_3/4$ when $n \geq 2$). This theorem establishes the convergence of $[x^t]$ and, importantly, the bounds (38) extend to the limit $[x^\infty]$ by continuity. This strong characterization of the limit point puts us in a favorable position to verify optimality.

Verifying optimality

To verify optimality of x^∞ , it is convenient to use a known dual certificate for the SDP relaxation. The following is a combination of Lemmas 4.3 and 4.4 in [2].

Lemma 18. A feasible X for (SDP) is optimal if and only if

$$S(X) := \text{Re}\{\text{ddiag}(CX)\} - C \quad (40)$$

is positive semidefinite, where ddiag sets all off-diagonal entries to zero. If furthermore $\text{rank}(S) = n - 1$, then X is the unique solution of (SDP), it is of the form $X = \hat{x}\hat{x}^*$ and \hat{x} is the unique global optimum of (P) up to global phase.

To simplify notation, let $x = x^\infty$. Using the same developments as in [2, §4.4], we verify that $S = S(xx^*)$ is positive semidefinite and has rank $n - 1$ under condition (32) on σ and the conclusions of Theorem 17, namely, inequalities (38) and equation (39). By construction, $Sx = 0$. Hence, it is sufficient to verify that $u^*Su > 0$ for all $u \in \mathbb{C}^n$ with $\|u\|_2 = 1$ and $u^*x = 0$:

$$\begin{aligned} u^*Su &= \sum_{k \in [n]} |u_k|^2 \text{Re}\{(Cx)_k \overline{x_k}\} - u^*Cu \\ &\stackrel{(39)}{=} \sum_{k \in [n]} |u_k|^2 |(Cx)_k| - |u^*z|^2 - \sigma u^*Wu \\ &\geq |z^*x| - \sigma \|Wx\|_\infty - d_2^2(z, x) - \sigma \|W\|_2. \end{aligned}$$

(Owing to $u^*x = 0$, we used $|u^*z| = |u^*(z - xe^{i\theta})| \leq \|z - xe^{i\theta}\|_2 = d_2(z, x)$, with appropriate choice of θ .) Now using $|z^*x| = n - \frac{1}{2}d_2^2(z, x)$ and assuming $\|W\|_2 \leq C'_2\sqrt{n}$ and the bounds (38) hold, it follows that

$$\begin{aligned} u^*Su &\geq n - \frac{3}{2}d_2^2(z, x) - \sigma\|Wx\|_\infty - \sigma\|W\|_2 \\ &\geq n - \frac{27}{8}\kappa_3^2n - \sigma(\kappa_2 + C'_2\kappa_3)\sqrt{n\log n} - \sigma C'_2\sqrt{n}. \end{aligned}$$

Assume σ satisfies inequality (32). Then, using $\kappa_1 = \kappa_3 = 1/60$ and $\kappa_2 = 4C'_1 + 2C'_2\kappa_1 \geq 2C'_2\kappa_3$ as in Theorem 15,

$$u^*Su \geq n \left(1 - \frac{27}{8}\kappa_3^2 - \frac{\kappa_2 + \kappa_2/2}{240\kappa_2} - \frac{1}{120\sqrt{2}} \right) > 0$$

Thus, S is positive semidefinite and has rank $n - 1$ —which implies xx^* is the unique solution of (SDP) and x is the unique solution of (P) up to phase by Lemma 18—provided σ satisfies (32), $\|W\|_2 \leq C'_2\sqrt{n}$ and the conclusions of Theorem 17 hold. Theorem 5 follows directly; details for Theorems 6–7 are in the appendix.

5 Conclusions and perspectives

We proved that both semidefinite relaxation and the generalized power method are able to find the global optimum of (P) under the regime $\sigma = \mathcal{O}(\sqrt{n/\log n})$ with high probability. In other words, the maximum-likelihood estimator of phase synchronization is computationally feasible under noise level $\sigma = \mathcal{O}(\sqrt{n/\log n})$, which (nearly) matches the information-theoretic threshold, thus closing the gap in previous papers. We also derived ℓ_2 and ℓ_∞ bounds on the optimum \hat{x} , and the ℓ_∞ bound improves upon previous results. The proof is based on tracking n auxiliary sequences, which is a novel technique developed in this paper. As a by-product, we also proved an ℓ_∞ bound for the eigenvector estimator, which is of independent interest.

An interesting problem for future work is to prove (or disprove) that second-order necessary optimality conditions are sufficient for (P). If this is true, then any algorithm that finds a second-order critical point also solves the nonconvex problem (P). This was proved in [8] for $\sigma = \mathcal{O}(n^{1/6})$ then in [23] for $\sigma = \mathcal{O}(n^{1/4})$. Numerical experiments in [8] suggest that a local optimization method (namely, the Riemannian trust-region method) with random initialization finds the global optimum with $\sigma = \tilde{\mathcal{O}}(n^{1/2})$ and random initialization. The analysis presented here does not apply directly though, because it hinges on a characterization of the limit points of GPM: a priori, this does not allow to characterize all second-order critical points.

A natural extension of our work is to establish similar results for synchronization over $\text{SO}(d)$ [9, 38] and $\text{SE}(d)$ [26]. The general synchronization problem is to recover group element $g_1, \dots, g_n \in G$ from their noisy pairwise measurements $g_k^{-1}g_\ell$. Our work here addresses synchronization over the group $\text{SO}(2)$ (equivalently, the group $\text{U}(1)$), that is,

in-plane rotations (equivalently, points on the unit complex circle). Another important group in practice is the rotation group $\text{SO}(3)$, which is often used to describe the orientation of an object [25, 12, 33]. It is shown empirically in [7] that the Riemannian trust-region method performs well. The analysis may be complicated by the fact that $\text{SO}(3)$ is a non-commutative group.

Another important problem in practice is to handle incomplete measurement sets. In this paper, we suppose all entries of C are known. A more realistic setting is that some pairs of phase differences are measured, forming edges of a graph. This appears in many applications [35, 17], and is addressed in a number of papers [34, 13]. The effect of an incomplete measurement graph on fundamental bounds is well understood as being related to the Laplacian of the graph [6]. See [30] for robotics applications.

Finally, another problem of practical concern is robustness of estimation methods. Here, (P) minimizes the sum of squared errors. However, in practice, more robust methods may be required to deal with outliers. In [38] for example, the authors minimize a sum of unsquared errors. A common way to solve such problems is via iterative reweighted least squares (IRLS), which is widely used in statistics [18]. IRLS solves a weighed least squares problem in each step, where the weights depend on the current iterate. In this regard, our analysis could be a first step toward understanding robust methods with IRLS for synchronization problems.

Acknowledgements

The authors thank Afonso Bandeira, Amelia Perry, Alexander Wein, Amit Singer and Jianqing Fan for helpful discussions.

A Proofs

Proofs for Section 3

Proof of Lemma 9. Let M_u be the upper triangular part of W , i.e., $(M_u)_{ij} = W_{ij}\mathbf{1}_{i \leq j}$, and $M_l = W - M_u$. Then $\text{Re}(M_u), \text{Im}(M_u), \text{Re}(M_l), \text{Im}(M_l)$ are all matrices with independent and sub-gaussian entries, whose sub-gaussian moments are bounded by an absolute constant (see [36] for equivalent definitions of sub-gaussian variables). We can then apply Proposition 2.4 in [32] and obtain the desired concentration bound on W . For $W^{(m)}$, we take a union bound over choice of m . The bound on $\Delta W^{(m)}$ follows from those on W and $W^{(m)}$. The bound on $\|w_m\|_2$ follows from that on $\|W\|_2$, since

$$\|W\|_2 = \sup_{\|u\|_2=1} \|Wu\|_2 \geq \sup_{\|u\|_2=1} |w_m^* u| = \|w_m\|_2.$$

□

Proof of Lemma 10. We will prove this lemma in the case where $\mathcal{U}^{(m)}$ is a *deterministic* set. The case where $\mathcal{U}^{(m)}$ is random follows easily from the deterministic case, since

we can first condition on $\mathcal{U}^{(m)}$ and use independence between $\Delta W^{(m)}$ and $\mathcal{U}^{(m)}$. In the proof, notations $C_1, C'_1, C_2, c_1, c_2 > 0$ denote some absolute constants. For a fixed $m \in [n]$ and $u \in \mathcal{U}^{(m)}$,

$$\|\Delta W^{(m)}u\|_2^2 = |(\Delta W^{(m)}u)_m|^2 + \sum_{k \neq m} |(\Delta W^{(m)}u)_k|^2. \quad (41)$$

We will bound the two parts on the right-hand side separately. We can expand $(\Delta W^{(m)}u)_m$ into a sum: $(\Delta W^{(m)}u)_m = \sum_k W_{mk}u_k$. By assumption, $\text{Re}(W_{mk})$ is a sub-gaussian random variable. By Hoeffding's inequality for sub-gaussian variables [36],

$$\mathbb{P} \left(\left| \sum_{k=1}^n \text{Re}(W_{mk}) \text{Re}(u_k) \right| \geq t \right) \leq \exp \left(1 - \frac{c_1 t^2}{n} \right).$$

For sums of random variables of $\text{Re}(W_{mk}) \text{Im}(u_k)$, $\text{Im}(W_{mk}) \text{Re}(u_k)$ or $\text{Im}(W_{mk}) \text{Im}(u_k)$ over k , similar concentration results hold. Thus, we can set $t = C_1 \sqrt{n \log n}$, where $C_1 > 0$ is some large absolute constant such that $c_1 C_1^2 \geq 5$, and deduce that with probability at least $1 - 4e \cdot n^{-5}$,

$$|(\Delta W^{(m)}u)_m| \leq 4C_1 \sqrt{n \log n}.$$

Now let us bound the second term in the right-hand side of (41). Observe

$$\sum_{k \neq m} |(\Delta W^{(m)}u)_k|^2 = \sum_{k \neq m} |W_{km}u_m|^2 = \sum_{k \neq m} (\text{Re}(W_{km})^2 + \text{Im}(W_{km})^2) |u_m|^2. \quad (42)$$

Since $\text{Re}(W_{km})$ is sub-gaussian, it follows that $\text{Re}(W_{km})^2$ is sub-exponential with a bounded sub-exponential norm. So we can use Bernstein's inequality for sub-exponential random variables [36],

$$\mathbb{P} \left(\left| \sum_{k=1}^n (\text{Re}(W_{km})^2 - \mathbb{E}[\text{Re}(W_{km})^2]) \right| \geq t \right) \leq 2 \cdot \exp \left(-c_2 \min \left\{ \frac{t^2}{4n}, \frac{t}{2} \right\} \right).$$

A similar concentration bound holds for $\sum_k \text{Im}(W_{km})^2$. Setting $t = n$, we know that with probability at least $1 - 4e^{-c_2 n/4}$,

$$\left| \sum_{k=1}^n (|W_{km}|^2 - \mathbb{E}|W_{km}|^2) \right| \leq 2n.$$

From an equivalent definition of sub-gaussian variables (see [36]), we obtain $\sum_{k=1}^n \mathbb{E}|W_{km}|^2 \leq C_2 n$, so it follows that $\sum_{k \neq m} |(\Delta W^{(m)}u)_k|^2 \leq (2 + C_2)nM_m^2$, where M_m is a bound on $|u_m|$, uniform over $\mathcal{U}^{(m)}$. Therefore, combining the upper bounds for the two terms in (41), we deduce that for some large absolute constant $C'_1 > 0$,

$$\|\Delta W^{(m)}u\|_2 \leq (16C_1^2 n \log n + (2 + C_2)nM_m^2)^{1/2} \leq C'_1 \sqrt{n \log n} + C'_1 \sqrt{n} M_m$$

holds with probability $1 - 4en^{-5} - 4e^{-c_2n/4}$. Taking a union bound over the choice of $u \in \mathcal{U}^{(m)}$ and m , we conclude that (21) holds with probability $1 - 12en^{-2} - 12n^3e^{-c_2n/4}$, or equivalently $1 - \mathcal{O}(n^{-2})$. \square

Proof of Theorem 8. In this proof, we use $a(n) \lesssim b(n)$ to mean there exists an absolute constant C such that $a(n) \leq Cb(n)$. We also suppose $\sigma < c'_0\sqrt{n}$ where $c'_0 < (8C'_2)^{-1}$, and C'_2 is the absolute constant in Lemma 9. With probability $1 - \mathcal{O}(n^{-2})$, Lemma 9 and Lemma 10 hold, so we can safely use the concentration bounds. First we note the second part of Theorem 8 (about $\mathcal{P}\tilde{x}$) follows directly from [23, Prop. 1] once the first part is proved.

It is also straightforward to bound $d_2(\tilde{x}, z)$. Since zz^* is a rank-1 matrix, the eigen-gap $\delta(zz^*) = \lambda_1(zz^*) - \lambda_2(zz^*)$ is simply its leading eigenvalue, which is n . From Lemma 11, clearly

$$d_2(\tilde{x}, z) \leq \frac{\sqrt{2}\sigma\|Wz\|_2}{n - \sigma\|W\|_2} \leq \frac{\sqrt{2}\sigma C'_2 n}{n - C'_2\sigma\sqrt{n}} \leq \frac{8\sqrt{2}}{7}C'_2\sigma \lesssim \sigma.$$

This leads to the first claim of the theorem.

Now let us consider bounding $\|W\tilde{x}\|_\infty$. Following the inequality (23), we only need to bound the two parts separately in (23). By Weyl's inequality, $\delta(C^{(m)}) = \lambda_1(C^{(m)}) - \lambda_2(C^{(m)}) \geq \lambda_1(zz^*) - 2\sigma\|W^{(m)}\|_2$, so Lemma 11 implies

$$d_2(\tilde{x}, \tilde{x}^{(m)}) \leq \frac{\sqrt{2}\sigma\|\Delta W^{(m)}\tilde{x}^{(m)}\|_2}{\delta(C^{(m)}) - \sigma\|\Delta W^{(m)}\|_2} \leq \frac{8\sqrt{2}c'_0}{5\sqrt{n}}\|\Delta W^{(m)}\tilde{x}^{(m)}\|_2 \quad (43)$$

Therefore, from (23) we have $\|W\tilde{x}\|_\infty \lesssim \max_m (|w_m^*\tilde{x}^{(m)}| + \|\Delta W^{(m)}\tilde{x}^{(m)}\|_2)$. Note the first term within the parenthesis is dominated by the second, since $w_m^*\tilde{x}^{(m)}$ is exactly the m th coordinate of $\Delta W^{(m)}\tilde{x}^{(m)}$. Lemma 10 implies that with probability $1 - \mathcal{O}(n^{-2})$,

$$\|W\tilde{x}\|_\infty \lesssim \max_{1 \leq m \leq n} \|\Delta W^{(m)}\tilde{x}^{(m)}\|_2 \lesssim \sqrt{n \log n} + \sqrt{n} \max_{1 \leq m \leq n} |\tilde{x}_m^{(m)}|. \quad (44)$$

We claim that $\max_m |\tilde{x}_m^{(m)}| < 8/7$. Indeed, for any $m \in [n]$, by definition of $\tilde{x}^{(m)}$,

$$\lambda_1(C^{(m)})\tilde{x}^{(m)} = C^{(m)}\tilde{x}^{(m)} = (z^*\tilde{x}^{(m)})z + \sigma W^{(m)}\tilde{x}^{(m)}.$$

Since $W^{(m)}$ has zero entries in its m th row, the m th coordinate of $W^{(m)}\tilde{x}^{(m)}$ vanishes, and we deduce

$$|\tilde{x}_m^{(m)}| = \frac{|(z^*\tilde{x}^{(m)})z_m|}{\lambda_1(C^{(m)})} = \frac{|z^*\tilde{x}^{(m)}|}{\lambda_1(C^{(m)})} \leq \frac{|z^*\tilde{x}^{(m)}|}{n - \sigma\|W^{(m)}\|_2} \leq \frac{n}{n - C'_2\sigma\sqrt{n}} < 8/7, \quad (45)$$

where we used Weyl's inequality $\lambda_1(C^{(m)}) \geq \lambda_1(zz^*) - \sigma\|W^{(m)}\|_2$. This leads to $\max_m |\tilde{x}_m^{(m)}| < 8/7$, and therefore $\|W\tilde{x}\|_\infty \lesssim \sqrt{n \log n}$. This ℓ_∞ bound is directly related to $d_\infty(\tilde{x}, z)$. We choose the global phase of \tilde{x} such that $\tilde{x}^*z = |\tilde{x}^*z|$, and thus $d_2(\tilde{x}, z)^2 = 2(n - |\tilde{x}^*z|)$. For any $m \in [n]$, $(C\tilde{x})_m = (z^*\tilde{x})z_m + \sigma(W\tilde{x})_m$, so

$$|\tilde{x}_m - z_m| = \left| \frac{(C\tilde{x})_m}{\lambda_1(C)} - z_m \right| \leq \left| \frac{z^*\tilde{x}}{\lambda_1(C)} - 1 \right| + \frac{\sigma\|W\tilde{x}\|_\infty}{\lambda_1(C)}.$$

From $d_2(\tilde{x}, z) \lesssim \sigma$ and the identity $d_2(\tilde{x}, z)^2 = 2(n - |\tilde{x}^* z|)$, we have $n \geq |z^* \tilde{x}| \geq n - \mathcal{O}(\sigma^2)$. By Weyl's inequality $|\lambda_1(C) - n| \leq \sigma \|W\|_2 \lesssim \sigma \sqrt{n}$, and therefore,

$$\|\tilde{x} - z\|_\infty \lesssim \frac{\sigma^2 + \sigma \sqrt{n}}{n} + \frac{\sigma \sqrt{n \log n}}{n} \lesssim \sigma \sqrt{\log n / n}.$$

□

Proofs for Section 4

Proof of Lemma 12. Let us decompose x, y into two parts that are orthogonal:

$$x = az + \sqrt{n} \alpha, \quad y = bz + \sqrt{n} \beta, \quad (46)$$

where $a, b \in \mathbb{C}; \alpha, \beta \in \mathbb{C}^n$, and $\alpha^* z = \beta^* z = 0$. Without loss of generality, we assume a, b are real and $a, b \geq 0$, since we can freely choose the global phases of x and y . Also suppose we choose $\theta \in \mathbb{R}$ such that $\|e^{i\theta} x - y\|_2$ is minimized, i.e., $\|e^{i\theta} x - y\|_2 = d_2(x, y)$. The key part of the proof is to show:

$$\frac{1}{n} |z^*(e^{i\theta} x - y)| = |e^{i\theta} a - b| \leq 6\varepsilon \|e^{i\theta} \alpha - \beta\|_2. \quad (47)$$

While the above equality is easily obtained by expressing x, y according to (46), the difficulty lies in showing the inequality. Once proved, this immediately leads to the desired inequality (25), because (using $\|e^{i\theta} \alpha - \beta\|_2 \leq d_2(x, y)/\sqrt{n}$)

$$\begin{aligned} \|e^{i\theta} \mathcal{L}x - \mathcal{L}y\|_2 &\leq \frac{1}{\sqrt{n}} |z^*(e^{i\theta} x - y)| + \frac{\sigma}{n} \|W(e^{i\theta} x) - Wy\|_2 \\ &\leq 6\varepsilon \sqrt{n} \|e^{i\theta} \alpha - \beta\|_2 + \frac{\sigma \|W\|_2}{n} \|e^{i\theta} x - y\|_2 \\ &= \left(6\varepsilon + \frac{\sigma \|W\|_2}{n}\right) d_2(x, y). \end{aligned}$$

Since $d_2(\mathcal{L}x, \mathcal{L}y) \leq \|e^{i\theta} \mathcal{L}x - \mathcal{L}y\|_2$, the proof will be complete. Note that using Cauchy-Schwarz inequality is futile as it leads to $|z^*(e^{i\theta} x - y)| \leq \sqrt{n} d_2(x, y)$, which cannot be used to show that \mathcal{L} is Lipschitz with a constant $\rho < 1$. An important intuition is that, if x and y are close to z , and $e^{i\theta}$ is close to 1, then $e^{i\theta} x - y$ cannot be aligned too much with z . The rest of the proof is devoted to showing (47).

Since $d_2(x, z)^2 = 2n(1 - a)$, $d_2(y, z)^2 = 2n(1 - b)$, the conditions $d_2(x, z) \leq \varepsilon \sqrt{n}$ and $d_2(y, z) \leq \varepsilon \sqrt{n}$ are equivalent to

$$a \geq 1 - \varepsilon^2/2, \quad b \geq 1 - \varepsilon^2/2. \quad (48)$$

Since $\varepsilon < 1/2$, we must have $a + b > 2(1 - 1/8) > 1$, and thus $|a - b| \leq (a + b)|a - b| = |a^2 - b^2|$. Since $\|x\|_2 = \|y\|_2 = \sqrt{n}$, we know $a^2 + \|\alpha\|_2^2 = 1$ and $b^2 + \|\beta\|_2^2 = 1$, so we

have $a^2 - b^2 = \|\beta\|_2^2 - \|\alpha\|_2^2$. Using the inequalities $\|\alpha\|_2 \leq d_2(x, z)/\sqrt{n} \leq \varepsilon$, $\|\beta\|_2 \leq d_2(y, z)/\sqrt{n} \leq \varepsilon$, we derive

$$|a - b| \leq (\|\alpha\|_2 + \|\beta\|_2) \left| \|\alpha\|_2 - \|\beta\|_2 \right| \leq 2\varepsilon \|e^{i\theta}\alpha - \beta\|_2, \quad (49)$$

where we used the triangular inequality in the second inequality. By the choice of θ , we have $\|e^{i\theta}x - y\|_2^2 \leq \|x - y\|_2^2$, or equivalently

$$|e^{i\theta}a - b|^2 + \|e^{i\theta}\alpha - \beta\|_2^2 \leq |a - b|^2 + \|\alpha - \beta\|_2^2.$$

Since $|e^{i\theta}a - b|^2 - |a - b|^2 = 2ab(1 - \cos \theta)$, and by the triangular inequality, $\|\alpha - \beta\|_2 - \|e^{i\theta}\alpha - \beta\|_2 \leq |1 - e^{i\theta}| \|\alpha\|_2 \leq \varepsilon|1 - e^{i\theta}|$, we obtain

$$\begin{aligned} 2ab(1 - \cos \theta) &\leq \|\alpha - \beta\|_2^2 - \|e^{i\theta}\alpha - \beta\|_2^2 \\ &\leq \varepsilon|1 - e^{i\theta}| (\|\alpha - \beta\|_2 + \|e^{i\theta}\alpha - \beta\|_2) \\ &\leq \varepsilon|1 - e^{i\theta}| (\varepsilon|1 - e^{i\theta}| + 2\|e^{i\theta}\alpha - \beta\|_2). \end{aligned}$$

Notice that $2(1 - \cos \theta) = |1 - e^{i\theta}|^2$, and thus we derive

$$(\varepsilon^{-1}ab - \varepsilon)|1 - e^{i\theta}| \leq 2\|e^{i\theta}\alpha - \beta\|_2.$$

From (48), we have $ab \geq (1 - \varepsilon^2/2)^2 > 1 - \varepsilon^2$, so

$$\varepsilon^{-1}ab - \varepsilon \geq \varepsilon^{-1}(1 - \varepsilon^2) - \varepsilon = \varepsilon^{-1}(1 - 2\varepsilon^2) > (2\varepsilon)^{-1},$$

where the last inequality is due to the condition $\varepsilon < 1/2$. Therefore, we deduce $|1 - e^{i\theta}| \leq 4\varepsilon\|e^{i\theta}\alpha - \beta\|_2$. Combining with (49), we obtain (recall that $a \leq 1$)

$$|e^{i\theta}a - b| = |e^{i\theta}a - a + a - b| \leq |e^{i\theta} - 1|a + |a - b| \leq 6\varepsilon\|e^{i\theta}\alpha - \beta\|_2.$$

This verifies (47) and completes the proof. \square

Proof of Lemma 13. By the cosine formula of triangles,

$$|x - y|^2 = |x|^2 + |y|^2 - 2|x||y|\cos \theta,$$

where $\theta = |\text{Arg}(x) - \text{Arg}(y)|$ is the angle formed by x and y . Using the AM-GM inequality (i.e., $|x|^2 + |y|^2 \geq 2|x||y|$ which follows from $||x| - |y||^2 \geq 0$), we have

$$|x - y|^2 \geq 2|x||y| - 2|x||y|\cos \theta \geq 2(1 - \varepsilon)^2(1 - \cos \theta). \quad (50)$$

Using the same cosine formula, we also have

$$\left| \frac{x}{|x|} - \frac{y}{|y|} \right|^2 = 2 - 2\cos \theta.$$

This equality, together with (50), leads to the desired inequality (26). Using (26) for all $k \in [n]$, it follows that $\|\mathcal{P}w - \mathcal{P}v\|_2 \leq (1 - \varepsilon)^{-1}\|w - v\|_2$. The same is true if we replace w by $we^{i\theta}$ with an arbitrary $\theta \in \mathbb{R}$. The second inequality (27) is obtained by minimizing $\|we^{i\theta} - v\|_2$ over θ . \square

Proof of Lemma 14. We will first show that the moduli of $(\mathcal{L}x)_k$ and $(\mathcal{L}'y)_k$ are uniformly lower bounded for all $k \in [n]$, so that we can use Lemma 13. Since $d_2(z, x)^2 = 2(n - |z^*x|)$ and $d_2(z, x)^2 \leq n\varepsilon_1^2$, it follows that $2|z^*x| \geq n(2 - \varepsilon_1^2)$. Thus, for any $k \in [n]$,

$$|(\mathcal{L}x)_k| = \frac{1}{n} |(z^*x)z_k + \sigma(Wx)_k| \geq \frac{1}{n} |z^*x| - \frac{\sigma}{n} |(Wx)_k| \geq 1 - \frac{1}{2}\varepsilon_1^2 - K_1\sigma\sqrt{\frac{\log n}{n}}.$$

Similarly, we have $|(\mathcal{L}'y)_k| \geq 1 - \varepsilon_2^2/2 - K_2\sigma\sqrt{\log n/n}$. Now that $|(\mathcal{L}x)_k|$ and $|(\mathcal{L}'y)_k|$ are uniformly lower bounded by $1 - \varepsilon$, where $\varepsilon := \max\{\varepsilon_1^2/2 + K_1\sigma\sqrt{\log n/n}, \varepsilon_2^2/2 + K_2\sigma\sqrt{\log n/n}\}$ is assumed smaller than 1, we are able to invoke Lemma 13, and deduce

$$d_2(\mathcal{P}\mathcal{L}x, \mathcal{P}\mathcal{L}'y) \leq (1 - \varepsilon)^{-1} d_2(\mathcal{L}x, \mathcal{L}'y).$$

This proves the first claim of the lemma; and the second claim follows directly from Lemma 12. \square

Proof of Theorem 15. Denote $C_1'' = 4C_1'$. We note that $x^{0,m}$ requires special treatment in the proof. Indeed, for $t \geq 1$, $x^{t,m}$ has unit-modulus entries, whereas $x^{0,m}$ does not (in general). From Lemma 9, we deduce that with probability $1 - \mathcal{O}(n^{-2})$, the matrix norm bounds (20) hold. In Lemma 10, we set $\mathcal{U}^{(m)} = \{x^{t,m} \in \mathbb{C}^n : t = 0, \dots, T\}$. By construction, $x^{t,m}$ is a function of $W^{(m)}$, and is independent of $\Delta W^{(m)}$. Besides, we have $|\mathcal{U}^{(m)}| = T + 1 \leq 3n^2$. This means the conditions in Lemma 10 are satisfied, so with probability $1 - \mathcal{O}(n^{-2})$,

$$\max_{0 \leq t \leq T} |w_m^* x^{t,m}| \leq \max_{0 \leq t \leq T} \|\Delta W^{(m)} x^{t,m}\|_2 \leq C_1'' \sqrt{n \log n}, \quad \forall m \in [n]. \quad (51)$$

This is because for $t \geq 1$, we can use (22) directly; and for $t = 0$, we use (21) and derive $\|\Delta W^{(m)} x^{0,m}\|_2 \leq C_1' \sqrt{n \log n} + 2C_1' \sqrt{n} \leq C_1'' \sqrt{n \log n}$ due to $\max_{m \in [n]} M_m \leq 2$ (derived in (45)).

First we verify (33)–(35) for $t = 0$. The initializers x^0 and $x^{0,m}$ are simply eigenvectors of C and $C^{(m)}$, and their bounds have been studied in Section 3. Recall that $x^{0,m}$ is independent of $\sigma \Delta W^{(m)}$, so applying Lemma 11 (Davis–Kahan) yields

$$d_2(x^{0,m}, x^0) \leq \frac{\sqrt{2}\sigma \|\Delta W^{(m)} x^{0,m}\|_2}{\delta(C^{(m)}) - \sigma \|\Delta W^{(m)}\|_2} \leq \frac{\sqrt{2}C_1''\sigma\sqrt{n \log n}}{n - 3\sigma \|\Delta W^{(m)}\|_2} \leq \frac{\sqrt{2}\kappa_2\sigma\sqrt{n \log n}}{n/2} < \kappa_1,$$

Here, in the second inequality, we used the bounds $\|\Delta W^{(m)} x^{0,m}\|_2 \leq C_1'' \sqrt{n \log n}$ by (51), $\|\Delta W^{(m)}\|_2 \leq C_2' \sqrt{n}$, and $\delta(C^{(m)}) = \lambda_1(C^{(m)}) - \lambda_2(C^{(m)}) \geq n - 2\sigma \|\Delta W^{(m)}\|_2$ by Weyl's inequality. In the third inequality, we used $C_1'' \leq \kappa_2$ and $3\sigma \|\Delta W^{(m)}\|_2 \leq 3C_2'\sigma\sqrt{n} < n/2$. The final inequality is due to the condition (32). This verifies (33) for the base case $t = 0$.

This immediately leads to (34), since by the reasoning in (23), we derive

$$\begin{aligned}
\|Wx^0\|_\infty &\leq \max_{1 \leq m \leq n} [|w_m^* x^{0,m}| + \|w_m\|_2 \cdot d_2(x^{0,m}, x^0)] \\
&\leq C_1'' \sqrt{n \log n} + C_2' \sqrt{n} \cdot d_2(x^{0,m}, x^0) \\
&\leq (C_1'' + 2C_2' \kappa_1) \sqrt{n \log n} \\
&= \kappa_2 \sqrt{n \log n}.
\end{aligned}$$

Finally, applying Lemma 11 (Davis–Kahan) again, we obtain (35):

$$d_2(x^0, z) \leq \frac{\sqrt{2} \sigma \|Wz\|_2}{n - \sigma \|W\|_2} \leq \frac{\sqrt{2} \sigma \cdot C_2' n}{n - C_2' \sqrt{n} \sigma} \leq \frac{n^{3/2}/120}{n/2} = \kappa_3 \sqrt{n}.$$

Now suppose (33)–(35) are true for t , and our goal is to derive the same inequalities for $t+1$, until t reaches T .

(i) Verifying (33) for $t+1$: notice that

$$d_2(x^{t+1,m}, x^{t+1}) \leq d_2(\mathcal{T}^{(m)} x^{t,m}, \mathcal{T} x^{t,m}) + d_2(\mathcal{T} x^{t,m}, \mathcal{T} x^t).$$

We shall use Lemma 14 to bound the above two terms. Before doing so, let us examine the conditions required in Lemma 14. In order to bound the first term, notice that

$$d_2(x^{t,m}, z) \leq d_2(x^{t,m}, x^t) + d_2(x^t, z) \leq \kappa_1 + \kappa_3 \sqrt{n} \leq (\kappa_1 + \kappa_3) \sqrt{n}.$$

Furthermore, let us check

$$\begin{aligned}
\|Wx^{t,m}\|_\infty &\leq \min_\theta \|W(e^{i\theta} x^{t,m} - x^t)\|_\infty + \|Wx^t\|_\infty \\
&\leq \|W\|_2 \cdot d_2(x^{t,m}, x^t) + \|Wx^t\|_\infty \\
&\leq C_2' \kappa_1 \sqrt{n} + \kappa_2 \sqrt{n \log n},
\end{aligned}$$

where we used the trivial inequality $\|\cdot\|_\infty \leq \|\cdot\|_2$. Moreover, using (51),

$$\begin{aligned}
\|W^{(m)} x^{t,m}\|_\infty &\leq \|\Delta W^{(m)} x^{t,m}\|_\infty + \|Wx^{t,m}\|_\infty \\
&\leq \|\Delta W^{(m)} x^{t,m}\|_2 + \|Wx^{t,m}\|_\infty \\
&\leq C_1'' \sqrt{n \log n} + C_2' \kappa_1 \sqrt{n} + \kappa_2 \sqrt{n \log n} \\
&\leq (C_1'' + 2C_2' \kappa_1 + \kappa_2) \sqrt{n \log n} = 2\kappa_2 \sqrt{n \log n}.
\end{aligned}$$

where we used the definition $\kappa_2 = C_1'' + 2C_2' \kappa_1$. Applying Lemma 14, we obtain, by (51)

$$\begin{aligned}
d_2(\mathcal{T}^{(m)} x^{t,m}, \mathcal{T} x^{t,m}) &\leq (1 - \eta_1)^{-1} d_2(\mathcal{L}^{(m)} x^{t,m}, \mathcal{L} x^{t,m}) \\
&\leq (1 - \eta_1)^{-1} \|(\mathcal{L}^{(m)} - \mathcal{L}) x^{t,m}\|_2 \\
&= (1 - \eta_1)^{-1} \sigma \|\Delta W^{(m)} x^{t,m}\|_2 / n \\
&\leq (1 - \eta_1)^{-1} \sigma C_1'' \sqrt{\log n / n},
\end{aligned}$$

where by condition (32),

$$\eta_1 := (\kappa_1 + \kappa_3)^2/2 + 2\kappa_2\sigma\sqrt{\log n/n} \leq \frac{1}{1800} + \frac{1}{120} < 1/2.$$

This now leads to

$$d_2(\mathcal{T}^{(m)}x^{t,m}, \mathcal{T}x^{t,m}) \leq 2\sigma C_1''\sqrt{\log n/n} \leq 2\sigma\kappa_2\sqrt{\log n/n} \leq \frac{1}{120} = \frac{\kappa_1}{2}, \quad (52)$$

by (32) again. We can bound $d_2(\mathcal{T}x^{t,m}, \mathcal{T}x^t)$ similarly. From the inequalities

$$\begin{aligned} d_2(x^t, z) &\leq \kappa_3\sqrt{n}, & d_2(x^{t,m}, z) &\leq \kappa_3\sqrt{n} + \kappa_1 \leq (\kappa_1 + \kappa_3)\sqrt{n}, \\ \|Wx^t\|_\infty &\leq \kappa_2\sqrt{n \log n}, & \|Wx^{t,m}\|_\infty &\leq C_2'\kappa_1\sqrt{n} + \kappa_2\sqrt{n \log n}, \end{aligned}$$

and Lemma 14, we derive

$$\begin{aligned} d_2(\mathcal{T}x^{t,m}, \mathcal{T}x^t) &\leq (1 - \eta_2)^{-1}(6(\kappa_1 + \kappa_3) + \sigma\|W\|_2/n) \cdot d_2(x^{t,m}, x^t) \\ &\leq 2 \cdot (0.2 + \sigma\|W\|_2/n)\kappa_1, \end{aligned}$$

where, since $2C_2'\kappa_1 < C_1'' + 2C_2'\kappa_1 = \kappa_2$,

$$\eta_2 := (\kappa_1 + \kappa_3)^2/2 + (2C_2'\kappa_1 + \kappa_2)\sigma\sqrt{\log n/n} < \eta_1 \leq 1/2.$$

Since $0.2 + \sigma\|W\|_2/n \leq 0.2 + \sigma C_2/\sqrt{n} < 1/4$ by condition (32), it follows that

$$d_2(\mathcal{T}x^{t,m}, \mathcal{T}x^t) \leq \kappa_1/2. \quad (53)$$

Therefore, bounds (52) and (53) imply that $d_2(x^{t+1,m}, x^{t+1}) \leq \kappa_1$.

(ii) Verifying (34) for $t + 1$: Following (15), for any $m \in [n]$, we obtain

$$\begin{aligned} |(Wx^{t+1})_m| &\leq |w_m^*x^{t+1,m}| + \|w_m\|_2 \cdot d_2(x^{t+1}, x^{t+1,m}) \\ &\leq C_1''\sqrt{n \log n} + C_2'\sqrt{n} \cdot \kappa_1, \end{aligned}$$

where we used $d_2(x^{t+1}, x^{t+1,m}) \leq \kappa_1$ (proved in (i)). Thus,

$$\|Wx^{t+1}\|_\infty \leq (C_1'' + 2C_2'\kappa_1)\sqrt{n \log n} = \kappa_2\sqrt{n \log n}.$$

(iii) Verifying (35) for $t + 1$: since $\mathcal{P}z = z$, by Lemma 13, we deduce

$$d_2(x^{t+1}, z) = d_2(\mathcal{P}\mathcal{L}x^t, \mathcal{P}z) \leq (1 - \varepsilon_0)^{-1}d_2(\mathcal{L}x^t, z), \quad (54)$$

where $\varepsilon_0 = 1 - \min_k |(\mathcal{L}x^t)_k|$. This is only informative if $\min_k |(\mathcal{L}x^t)_k| > 0$. Observe

$$\begin{aligned} \min_\theta \|e^{i\theta}\mathcal{L}x^t - z\|_\infty &\leq \min_\theta \|(e^{i\theta}(z^*x^t)/n - 1)z\|_\infty + \sigma\|Wx^t\|_\infty/n, \\ &\leq \min_\theta |e^{i\theta}(z^*x^t)/n - 1| + \kappa_2\sigma\sqrt{\log n/n}, \end{aligned}$$

where we used $\|Wx^t\|_\infty \leq \kappa_2\sqrt{n\log n}$ (by the induction hypothesis). Since

$$\min_{\theta} |e^{i\theta}(z^*x^t)/n - 1| = 1 - |z^*x^t|/n = \frac{d_2(x^t, z)^2}{2n} \leq \kappa_3^2/2,$$

and $\kappa_2\sigma\sqrt{\log n/n} \leq 1/240 = \kappa_3/4$, and also $\kappa_3 = 1/60$, it follows that

$$d_\infty(\mathcal{L}x^t, z) = \min_{\theta} \|e^{i\theta}\mathcal{L}x^t - z\|_\infty \leq \kappa_3/120 + \kappa_3/4 < \kappa_3/2.$$

This a fortiori verifies $\min_k |(\mathcal{L}x^t)_k| \geq 1 - \kappa_3/2 > 1/2$, and thus from (54),

$$d_2(x^{t+1}, z) \leq 2d_2(\mathcal{L}x^t, z) \leq 2\sqrt{n}d_\infty(\mathcal{L}x^t, z) < \kappa_3\sqrt{n}. \quad (55)$$

Here, the inequality $d_2(w, v) \leq \sqrt{n}d_\infty(w, v)$ ($\forall w, v \in \mathbb{C}_1^n$) is obtained by considering $d_2(w, v) \leq \|e^{i\theta}w - v\|_2 \leq \sqrt{n}\|e^{i\theta}w - v\|_\infty = d_\infty(w, v)$ for the choice of θ determined by d_∞ .

Finally, we use induction and deduce that, for any $t = 0, 1, \dots, T$, the desired bounds (33)–(35) hold with probability $1 - \mathcal{O}(n^{-2})$. \square

Proof of Theorem 16. In Theorem 15, we showed that with probability $1 - \mathcal{O}(n^{-2})$, $x^t \in \mathcal{N}$ for $t \leq T$. Now we can use Lemma 14, which yields

$$d_2(\mathcal{T}x^t, \mathcal{T}x^{t-1}) \leq \rho \cdot d_2(x^t, x^{t-1}), \quad \forall 1 \leq t \leq T-1,$$

where $\rho = (1 - \varepsilon)^{-1}(6\kappa_3 + \sigma\|W\|_2/n)$, and $\varepsilon = \kappa_3^2/2 + \kappa_2\sigma\sqrt{\log n/n} < 1/2$. Thus, $\rho \leq 2 \cdot (0.1 + \sigma C'_2/\sqrt{n}) \leq 2 \cdot (0.1 + 1/120\sqrt{2}) < 1/2$. This leads to

$$d_2(x^{t+1}, x^t) \leq \frac{1}{2} d_2(x^t, x^{t-1}), \quad \forall 1 \leq t \leq T-1,$$

and thus $d_2(x^T, x^{T-1}) \leq 2^{1-T}d_2(x^1, x^0) \leq 2^{1-T}(\|x^1\|_2 + \|x^0\|_2) = 2^{2-T}\sqrt{n}$. \square

Proof of Theorem 17. Consider the first part of the theorem. Let us use induction on k to establish (37). Trivially, (37) holds for $k = 0$. Now suppose for all $\ell \leq k$, $d_2(x^{T+\ell}, x^{T+\ell-1}) \leq 2^{-\ell}d_2(x^T, x^{T-1})$. We now show the same for $k+1$. First,

$$d_2(x^{T+k}, x^{T-1}) \leq \sum_{\ell=0}^k d_2(x^{T+\ell}, x^{T+\ell-1}) \leq \frac{\kappa_3}{4} \sum_{\ell=0}^k 2^{-\ell} < \frac{\kappa_3}{2}.$$

Similarly, $d_2(x^{T+k-1}, x^{T-1}) < \kappa_3/2$. By assumption, $d_2(x^{T-1}, z) \leq \kappa_3\sqrt{n}$, so

$$d_2(x^{T+k}, z) \leq 3\kappa_3\sqrt{n}/2, \quad d_2(x^{T+k-1}, z) \leq 3\kappa_3\sqrt{n}/2. \quad (56)$$

Moreover,

$$\begin{aligned} \|Wx^{T+k}\|_\infty &\leq \|Wx^{T-1}\|_\infty + \min_{\theta} \|W(e^{i\theta}x^{T+k} - x^{T-1})\|_2 \\ &\leq \kappa_2\sqrt{n\log n} + \|W\|_2 \cdot d_2(x^{T+k}, x^{T-1}) \\ &\leq (\kappa_2 + C'_2\kappa_3)\sqrt{n\log n}, \end{aligned} \quad (57)$$

and similarly $\|Wx^{T+k-1}\|_\infty \leq (\kappa_2 + C'_2\kappa_3)\sqrt{n \log n}$. Using Lemma 14, we have

$$\begin{aligned} d_2(\mathcal{T}x^{T+k}, \mathcal{T}x^{T+k-1}) &\leq (1 - \varepsilon)^{-1}(9\kappa_3 + \sigma\|W\|_2/n) \cdot d_2(x^{T+k}, x^{T+k-1}) \\ &\leq 2 \left(\frac{9}{60} + \frac{1}{120\sqrt{2}} \right) \cdot d_2(x^{T+k}, x^{T+k-1}) \\ &\leq \frac{1}{2}d_2(x^{T+k}, x^{T+k-1}), \end{aligned} \tag{58}$$

where $\varepsilon \leq (\frac{3}{2}\kappa_3)^2/2 + (\kappa_2 + C'_2\kappa_3)\sigma\sqrt{\log n/n} < 1/2$. By the induction hypothesis, $d_2(x^{T+k}, x^{T+k-1}) \leq 2^{-k}d_2(x^T, x^{T-1})$, so $d_2(\mathcal{T}x^{T+k}, \mathcal{T}x^{T+k-1}) \leq 2^{-k-1}d_2(x^T, x^{T-1})$. This completes the induction and proves (37).

Now consider the second part of the theorem. The existence of the limit follows if we can show the metric space \mathbb{C}_1^n/\sim is complete. For any Cauchy sequence $\{[u^t]\}_{t=1}^\infty$ in \mathbb{C}_1^n/\sim , we can choose a subsequence $\{[u^{t_k}]\}_{k=1}^\infty$ such that $d_2([u^{t_k}], [u^{t_{k+1}}]) \leq 2^{-k}$. Now we define a sequence $\{y^k\}_{k=1}^\infty$ in \mathbb{C}_1^n by fixing the phases of $[u^{t_k}]$ iteratively. Given an arbitrary representative u^{t_1} in the equivalence class $[u^{t_1}]$, we set $y^1 = u^{t_1}$. Given y^k , we can choose a suitable $\theta \in \mathbb{R}$ such that $\|e^{i\theta}u^{t_{k+1}} - y^k\|_2$ attains $d_2(u^{t_{k+1}}, y^k)$. Then we set $y^{k+1} = e^{i\theta}u^{t_{k+1}}$. This produces a Cauchy sequence $\{y^k\}_{k=1}^\infty$ in \mathbb{C}_1^n , so it admits a limit $y^\infty = \lim_k y^k \in \mathbb{C}_1^n$. It follows, thus, that $[y^\infty]$ is the limit of $\{[u^{t_k}]\}_{k=1}^\infty$ in \mathbb{C}_1^n/\sim . Moreover, it is easy to see that $[y^\infty]$ is also the limit of $\{[u^t]\}_{t=1}^\infty$.

Hence, $[x^t]$ converges to a limit $[x^\infty]$ in \mathbb{C}_1^n/\sim . The inequalities in (38) follows from (56), (57) and the continuity argument.

Finally, we will show that x^∞ is a fixed point of \mathcal{T} . Because all iterates (including x^∞) satisfy (56) and (57), we can use Lemma 14 and derive $d_2(\mathcal{T}x^\infty, \mathcal{T}x^{T+k-1}) \leq d_2(x^\infty, x^{T+k-1})/2$ (similarly as in (58)). The right-hand side is vanishing as $k \rightarrow \infty$, so

$$d_2(\mathcal{T}x^\infty, x^\infty) \leq \lim_{k \rightarrow \infty} (d_2(\mathcal{T}x^\infty, x^{T+k}) + d_2(x^{T+k}, x^\infty)) = 0.$$

This implies that there exists some $\theta \in \mathbb{R}$ such that $e^{i\theta}\mathcal{T}x^\infty = x^\infty$. Since $\min_k |(\mathcal{L}x^\infty)_k| > 1 - \varepsilon > 0$ (as in (58)), we can rewrite it as

$$e^{i\theta}\mathcal{L}x^\infty = n^{-1}\text{diag}(\mu)x^\infty,$$

where $\mu_k = n|(\mathcal{L}x^\infty)_k|$. The above identity implies $e^{i\theta}(x^\infty)^*\mathcal{L}x^\infty = \sum_{k=1}^n \mu_k/n > 0$. Since \mathcal{L} is Hermitian, the left-hand side must be real, so $e^{i\theta} \in \{\pm 1\}$. Note that $e^{i\theta}$ must be 1, since

$$\begin{aligned} (x^\infty)^*\mathcal{L}x^\infty &= \frac{1}{n} |(x^\infty)^*z|^2 + \frac{1}{n}\sigma x^*Wx \geq \frac{1}{n}(n - d_2(x^\infty, z)^2/2)^2 - \sigma\|W\|_2 \\ &\geq n(1 - (3\kappa_3/2)^2/2)^2 - C'_2\sigma\sqrt{n} = n(1 - \frac{1}{3200})^2 - C'_2\sigma\sqrt{n}, \end{aligned}$$

which is positive under condition (32). Replacing \mathcal{L} with C/n , we finish the proof. \square

Proofs for Section 2

This subsection presents the proofs of the main results Theorems 6–7. Note Theorem 8 is already proved in Section 3, and Theorem 5 is proved in Section 4.

Proof of Theorem 6. With probability $1 - \mathcal{O}(n^{-2})$, (36) and (37) hold. With the trivial inequality $d_2(x^1, x^0) \leq 2\sqrt{n}$, we have $d_2(x^t, x^{t-1}) \leq 2^{2-t}\sqrt{n}$ for all t . By the triangular inequality,

$$d_2(x^t, x^\infty) \leq \sum_{k=t}^{\infty} d_2(x^{k+1}, x^k) \leq \sum_{k=t}^{\infty} 2^{1-k}\sqrt{n} = 2^{2-t}\sqrt{n}.$$

This proves Theorem 6. \square

Proof of Theorem 7. Without loss of generality, suppose $n \geq 2$ and fix the global phase of x^∞ such that $z^*x^\infty = |z^*x^\infty|$. Also suppose the constant c_0 satisfies (32).

(i) We first show $\|x^\infty - z\|_2 \leq 4C'_2\sigma$. This follows from the same argument in [2, Lemma 4.1]. Note that due to the sub-gaussian assumption in this paper, a difference is the bound on $\|W\|_2$: $\|W\|_2 \leq C'_2\sqrt{n}$ by Lemma 9 (whereas $\|W\|_2 \leq 3\sqrt{n}$ in [2]).

(ii) In the last subsection “Verifying optimality” of Section 4, we showed, using the conclusions of Theorem 17, that x^∞ is the unique optimum of (P) up to phase with probability $1 - \mathcal{O}(n^{-2})$.

(iii) Next we show $\|x^\infty - z\|_\infty \leq C\sigma\sqrt{\log n/n}$ where $C > 0$ is some constant. From (i), $|z^*x^\infty| = n - \|x^\infty - z\|_2^2/2 \geq n - 8(C'_2)^2\sigma^2$. In the proof of Lemma 4.2 in [2], it is shown that

$$|z^*x^\infty|\|x^\infty - z\|_\infty \leq 2\sigma\|Wx^\infty\|_\infty.$$

Therefore, it follows that

$$\begin{aligned} \|x^\infty - z\|_\infty &\leq 2n^{-1}\sigma\|Wx^\infty\|_\infty + 8(C'_2)^2n^{-1}\sigma^2\|x^\infty - z\|_\infty \\ &\leq 2n^{-1}\sigma(\kappa_2 + C'_2\kappa_3)\sqrt{n\log n} + 16(C'_2)^2n^{-1}\sigma^2, \end{aligned}$$

where we used (38) in Theorem 17 and a trivial bound $\|x^\infty - z\|_\infty \leq 2$. Since $\sigma = \mathcal{O}(\sqrt{n/\log n})$, we conclude that there exists a constant C such that $\|x^\infty - z\|_\infty \leq C\sigma\sqrt{\log n/n}$. \square

References

- [1] A.S. Bandeira. Random Laplacian matrices and convex relaxations. *arXiv preprint arXiv:1504.03987*, 2015.
- [2] A.S. Bandeira, N. Boumal, and A. Singer. Tightness of the maximum likelihood semidefinite relaxation for angular synchronization. *Mathematical Programming*, pages 1–23, 2016.
- [3] A.S. Bandeira, Y. Khoo, and A. Singer. Open problem: Tightness of maximum likelihood semidefinite relaxations. In Csaba Szepesvári Maria Florina Balcan, Vitaly Feldman, editor, *Proceedings of the 27th Conference on Learning Theory*, volume 35 of *JMLR W&CP*, pages 1265–1267, 2014.

- [4] Derek Bean, Peter J Bickel, Noureddine El Karoui, and Bin Yu. Optimal m-estimation in high-dimensional regression. *Proceedings of the National Academy of Sciences*, 110(36):14563–14568, 2013.
- [5] Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.
- [6] N. Boumal. On intrinsic Cramér-Rao bounds for Riemannian submanifolds and quotient manifolds. *Signal Processing, IEEE Transactions on*, 61(7):1809–1821, 2013.
- [7] N. Boumal. A Riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints. *arXiv preprint arXiv:1506.00575*, 2015.
- [8] N. Boumal. Nonconvex phase synchronization. *SIAM Journal on Optimization*, 26(4):2355–2377, 2016.
- [9] N. Boumal, A. Singer, and P.-A. Absil. Robust estimation of rotations from relative measurements by maximum likelihood. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 1156–1161, Dec 2013.
- [10] M. Capitaine, C. Donati-Martin, and D. Féral. The largest eigenvalues of finite rank deformation of large Wigner matrices: convergence and nonuniversality of the fluctuations. *The Annals of Probability*, 37(1):1–47, 2009.
- [11] Yuxin Chen and Emmanuel Candes. The projected power method: An efficient algorithm for joint alignment from pairwise differences. *arXiv preprint arXiv:1609.05820*, 2016.
- [12] M. Cucuringu, A. Singer, and D. Cowburn. Eigenvector synchronization, graph rigidity and the molecule problem. *Information and Inference, a Journal of the IMA*, 1(1):21–67, 2012.
- [13] Mihai Cucuringu, Yaron Lipman, and Amit Singer. Sensor network localization by eigenvector synchronization over the euclidean group. *ACM Transactions on Sensor Networks (TOSN)*, 8(3):19, 2012.
- [14] Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.
- [15] Y. Deshpande, A. Montanari, and E. Richard. Cone-constrained principal component analysis. In Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2717–2725. Curran Associates, Inc., 2014.
- [16] Jianqing Fan, Weichen Wang, and Yiqiao Zhong. An ℓ^∞ eigenvector perturbation bound and its application to robust covariance estimation. *arXiv preprint arXiv:1603.03516*, 2016.
- [17] A. Giridhar and P.R. Kumar. Distributed clock synchronization over wireless networks: Algorithms and analysis. In *Decision and Control, 2006 45th IEEE Conference on*, pages 4915–4920. IEEE, 2006.
- [18] Paul W Holland and Roy E Welsch. Robust regression using iteratively reweighted least-squares. *Communications in Statistics-theory and Methods*, 6(9):813–827, 1977.
- [19] Adel Javanmard, Andrea Montanari, and Federico Ricci-Tersenghi. Phase transitions in semidefinite relaxations. *Proceedings of the National Academy of Sciences*, 113(16):E2218–E2223, 2016.

- [20] Iain M Johnstone and Arthur Yu Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 2012.
- [21] M. Journée, Y. Nesterov, P. Richtárik, and R. Sepulchre. Generalized power method for sparse principal component analysis. *The Journal of Machine Learning Research*, 11:517–553, 2010.
- [22] Marc Lelarge and Léo Miolane. Fundamental limits of symmetric low-rank matrix estimation. *arXiv preprint arXiv:1611.03888*, 2016.
- [23] H. Liu, M.-C. Yue, and A. M.-C. So. On the estimation performance and convergence rate of the generalized power method for phase synchronization. *arXiv preprint arXiv:1603.00211*, 2016.
- [24] R. Luss and M. Teboulle. Conditional gradient algorithms for rank-one matrix approximations with a sparsity constraint. *SIAM Review*, 55(1):65–98, 2013.
- [25] Daniel Martinec and Tomas Pajdla. Robust rotation and translation estimation in multiview reconstruction. In *Computer Vision and Pattern Recognition, 2007. CVPR’07. IEEE Conference on*, pages 1–8. IEEE, 2007.
- [26] Onur Ozyesil, Nir Sharon, and Amit Singer. Synchronization over Cartan motion groups via contraction. *arXiv preprint arXiv:1612.00059*, 2016.
- [27] Debashis Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, pages 1617–1642, 2007.
- [28] A. Perry, A.S. Wein, A.S. Bandeira, and A. Moitra. Optimality and sub-optimality of PCA for spiked random matrices and synchronization. *arXiv preprint arXiv:1609.05573*, 2016.
- [29] Karl Rohe, Sourav Chatterjee, and Bin Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, pages 1878–1915, 2011.
- [30] David M Rosen, Luca Carlone, Afonso S Bandeira, and John J Leonard. A certifiably correct algorithm for synchronization over the special euclidean group. *arXiv preprint arXiv:1611.00128*, 2016.
- [31] V. Roulet, N. Boumal, and A. d’Aspremont. Renegar’s condition number and compressed sensing performance. *arXiv preprint arXiv:1506.03295*, 2015.
- [32] Mark Rudelson and Roman Vershynin. Non-asymptotic theory of random matrices: extreme singular values. *arXiv preprint arXiv:1003.2990*, 2010.
- [33] Yoel Shkolnisky and Amit Singer. Viewing direction estimation in cryo-em using synchronization. *SIAM journal on imaging sciences*, 5(3):1088–1110, 2012.
- [34] A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1):20–36, 2011.
- [35] X Yu Stella. Angular embedding: from jarring intensity differences to perceived luminance. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 2302–2309. IEEE, 2009.
- [36] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.

- [37] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- [38] L. Wang and A. Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference*, 2(2):145–193, 2013.
- [39] S. Zhang and Y. Huang. Complex quadratic optimization and semidefinite programming. *SIAM Journal on Optimization*, 16(3):871–890, 2006.