Tightness of the maximum likelihood semidefinite relaxation for angular synchronization

Afonso S. Bandeira[*], Nicolas Boumal[†], Amit Singer[‡]

November 13, 2014

Abstract

Many maximum likelihood estimation problems are, in general, intractable optimization problems. As a result, it is common to approximate the maximum likelihood estimator (MLE) using convex relaxations. Semidefinite relaxations are among the most popular. Sometimes, the relaxations turn out to be tight. In this paper, we study such a phenomenon.

The angular synchronization problem consists in estimating a collection of n phases, given noisy measurements of some of the pairwise relative phases. The MLE for the angular synchronization problem is the solution of a (hard) non-bipartite Grothendieck problem over the complex numbers. It is known that its semidefinite relaxation enjoys worst-case approximation guarantees.

In this paper, we consider a stochastic model on the input of that semidefinite relaxation. We assume there is a planted signal (corresponding to a ground truth set of phases) and the measurements are corrupted with random noise. Even though the MLE does not coincide with the planted signal, we show that the relaxation is, with high probability, tight. This holds even for high levels of noise.

This analysis explains, for the interesting case of angular synchronization, a phenomenon which has been observed without explanation in many other settings. Namely, the fact that even when exact recovery of the ground truth is impossible, semidefinite relaxations for the MLE tend to be tight (in favorable noise regimes).

Keywords: Angular Synchronization; Semidefinite programming; Tightness of convex relaxation; Maximum likelihood estimation.

[*]Princeton University. ajsb@math.princeton.edu, http://www.math.princeton.edu/~ajsb/.
[†]INRIA and ENS Paris. nicolasboumal@gmail.com, http://www.nicolasboumal.net.
[‡]Princeton University. amits@math.princeton.edu, http://www.math.princeton.edu/~amits/.
Authors are listed alphabetically.

1 Introduction

Recovery problems in statistics and many other fields are commonly solved under the paradigm of maximum likelihood estimation, partly due to the rich theory it enjoys. Unfortunately, in many important applications, the parameter space is exponentially large and non-convex, often rendering the computation of the maximum likelihood estimator (MLE) intractable. It is then common to settle for heuristics, such as expectation-maximization algorithms to name but one example. However, it is also common for such iterative heuristics to get trapped in local optima. Furthermore, even when these methods do attain a global optimum, there is in general no way to verify this.

A now classic alternative to these heuristics is the use of convex relaxations. The idea is to maximize the likelihood in a larger, convex set that contains the parameter space of interest, as (well-behaved) convex optimization problems are generally well understood and can be solved in polynomial time. The downside is that the solution obtained might not be in the original feasible (acceptable) set. One is then forced to take an extra, potentially suboptimal, rounding step.

This line of thought is the basis for a wealth of modern *approximation algorithms* [54]. One preeminent example is Goemans and Williamson's treatment of Max-Cut [29], an NP-hard combinatorial problem that involves segmenting a graph in two clusters to maximize the number of edges between the two clusters. They first show that Max-Cut can be formulated as a semidefinite program (SDP)—a convex optimization problem where the variable is a positive semidefinite matrix—with the additional, non-convex constraint that the sought matrix be of rank one. Then, they propose to solve this SDP while relaxing (ignoring) the rank constraint. They show that the obtained solution, despite typically being of rank strictly larger than one, can be rounded to a (suboptimal) rank-one solution, and that it provides a guaranteed approximation of the optimal value of the hard problem. Results of the same nature abound in the recent theoretical computer science literature [30, 55, 17, 25, 11, 10, 48, 37, 55].

In essence, approximation algorithms insist on solving *all* instances of a given NP-hard problem in polynomial time, which, unless P = NP, must come at the price of accepting some degree of sub-optimality. This worst-case approach hinges on the fact that a problem is NP-hard as soon as every efficient algorithm for it can be hindered by at least one pathological instance.

Alternatively, in a non-adversarial setting where "the data is not an enemy," one may find that such pathological cases are not prominent. As a result, the applied mathematics community has been more interested in identifying regimes for which the convex relaxations are tight, that is, admit a solution that is also admissible for the hard problem. When this is the case, no rounding is necessary and a truly optimal solution is found in reasonable time, together with a certificate of optimality. This is sometimes achieved by positing a probability distribution on the instances of the problem and asserting tightness *with high probability*, as for example in compressed sensing [19, 27, 50], matrix completion [18], Z2-synchronization [1] and inverse problems [23, 7]. In this approach, one surrenders the hope to solve all instances of the hard problem, in exchange for true optimality.

The results presented in this work are of that nature. As a concrete object of study, we consider the *angular synchronization problem* [45, 11], which consists in estimating a collection of n phases $e^{i\theta_1}, \ldots, e^{i\theta_n}$, given noisy measurements of pairwise relative phases $e^{i(\theta_k - \theta_\ell)}$—see Section 2 for a formal description. This problem notably comes up in time-synchronization of distributed networks [28], signal reconstruction from phaseless measurements [5, 9], and surface reconstruction problems in computer vision [4] and optics [42]. Angular synchronization serves as a model for the more general problem of synchronization of rotations in any dimension, which comes up in structure from motion [38, 32], surface reconstruction from 3D scans [53] and cryo-electron microscopy [46], to name a few.

The main contribution of the present paper is a proof that, even though the angular synchronization problem is NP-hard [55], its MLE in the face of Gaussian noise can often be computed (and certified) in polynomial time. This remains true even for entry-wise noise levels growing to infinity as the size of the problem (the number of phases) grows to infinity. The MLE is obtained as the solution of a semidefinite relaxation described in Section 2. This arguably striking phenomenon has been observed empirically before [12]—see also Figure 2, but not explained.

Computing the MLE for angular synchronization is equivalent to solving a non-bipartite Grothendieck problem. Semidefinite relaxations for Grothendieck problems have been thoroughly studied in theoretical computer science from the point of view of approximation ratios. The name is inspired by its close relation to an inequality of Grothendieck [31]. We direct readers to the survey by Pisier [41] for a thorough discussion. Our results show the surprising phenomenon that, in a randomized version of the Grothendieck problem, where there is a planted signal, there is no gap between the SDP relaxation and the original problem, with high probability.

The proposed result is qualitatively different from most tightness results available in the literature. Typical results establish either exact recovery of a planted signal [1, 33] (mostly in discrete settings), or exact recovery in the absence of noise, joint with stable (but not necessarily optimal) recovery when noise is present [21, 24, 20, 53, 25]. In contrast, this paper shows optimal recovery even though exact recovery is not possible. In particular, Demanet and Jugnon showed stable recovery for angular synchronization via semidefinite programming, in an adversarial noise setting [25]. We complement that analysis by showing tightness in a non-adversarial setting, meaning that the actual MLE is computed.

The proof relies on verifying that a certain candidate dual certificate is valid with high probability. The main difficulty comes from the fact that the dual certificate depends of the MLE, which does not coincide with the planted signal, and is a nontrivial function of the noise. We use necessary optimality conditions of the hard problem to both obtain an explicit expression for the candidate dual certificate, and to partly characterize the point whose optimality we aim to establish. This seems to be required since the MLE is not known in closed form.

In the context of sparse recovery, a result with similar flavor is support recovery guarantee [50], where the support of the estimated signal is shown to be contained in the support of the original signal. Due to the noise, exact recovery is also impossible in this setting. Another example is a recovery guarantee in the context of latent variable selection in graphical models [22].

Besides the relevance of angular synchronization in and of its own, we are confident this new insight will help uncover similar results in other applications where it has been observed that semidefinite relaxations can be tight even when the ground truth cannot be recovered. Notably, this appears to be the case for the Procrustes and multi-reference alignment problems (see Section 5 and [8, 12]).

The crux of our argument is to show that the SDP, with random data following a given distribution, admits a unique rank-one solution with high probability. We mention in passing that there are many other, deterministic results in the literature pertaining to the rank of solutions of SDP's. For example, Barvinok [14] and Pataki [40] both show that, in general, an SDP with only equality constraints admits a solution of rank at most (on the order of) the square root of the number of constraints. Furthermore, Sagnol [44] shows that under some conditions (that are not fulfilled in our case), certain SDP's related to packing problems always admit a rank-one solution. Sojoudi and Lavaei [49] study a class of SDP's on graphs which is related to ours and for which, under certain strong conditions on the topology of the graphs, the SDP's admit rank-one solutions—see also applications to power flow optimization [36].

1.1 Notation

For a complex scalar $a \in \mathbb{C}$, $\overline{a}$ denotes its complex conjugate and $|a| = \sqrt{a\overline{a}}$ its modulus. For a complex vector $v \in \mathbb{C}^n$, v^* denotes its conjugate transpose, $\|v\|^2 = \|v\|_2^2 = v^*v$ and $\|v\|_\infty = \max_i |v_i|$. $\mathbb{1}$ is a vector whose entries are all equal to 1. $\mathrm{diag}(v)$ is a diagonal matrix with diagonal formed by v. For a matrix M, $\|M\|_{\mathrm{op}}$ is the operator norm (maximum singular value), $\|M\|_{\mathrm{e},\infty} = \max_{i,j} |M_{ij}|$ is the entry-wise ℓ_∞ norm (largest entry in absolute value), $\Re(M)$ and $\Im(M)$ extract, respectively, the real and imaginary parts of a matrix (or a vector, or a scalar). For a square matrix M, $\mathrm{diag}(M)$ extracts the diagonal of M into a vector and $\mathrm{ddiag}(M)$ sets all off-diagonal entries of M to zero. For a symmetric matrix M, $\mathcal{L}(M)$ denotes its Laplacian (4.10) and $\lambda_{\mathcal{L}(M)}$ denotes the spectral gap of the Laplacian (4.12). $\mathbb{E}$ denotes mathematical expectation.

1.2 Organization of the paper

Section 2 gives a formal description of the angular synchronization problem and its semidefinite relaxation, for both the real and complex cases. Figures illustrate empirically the tightness of the relaxation. Our main theorem is stated for Gaussian noise, with a road-map of the proof. Section 3 states our main deterministic result. Section 4 hosts a proof of the main result, demonstrating both that the MLE is close to the sought signal and that the semidefinite relaxation is tight, under some assumptions. Section 5 discusses the applicability of these results to the Procrustes problem. Finally, Section 6 proposes some final thoughts on this work.

2 The Angular Synchronization problem

We focus on the problem of *angular synchronization* [45, 11], in which one wishes to estimate a collection of n phases ($n \geq 2$) based on measurements of pairwise phase differences. We will restrict our analysis to the case where a measurement is available for every pair of nodes. More precisely, we let $z \in \mathbb{C}^n$ be an unknown, complex vector with unit modulus entries, $|z_1| = \cdots = |z_n| = 1$, and we consider measurements of the form $C_{ij} = z_i\overline{z_j} + \varepsilon_{ij}$, where $\overline{z_j}$ denotes the complex conjugate of z_j and $\varepsilon_{ij} \in \mathbb{C}$ is noise affecting the measurement. By symmetry, we define $C_{ji} = \overline{C_{ij}}$ and $C_{ii} = 1$, so that the matrix $C \in \mathbb{C}^{n\times n}$ whose entries are given by the C_{ij}'s is a Hermitian matrix. Further letting the noise ε_{ij} be i.i.d. (complex) Gaussian variables for $i < j$, it follows directly that an MLE for the vector z is any vector of phases $x \in \mathbb{C}^n$ minimizing $\sum_{i,j} |C_{ij}x_j - x_i|^2$. Equivalently, an MLE is a solution of the following quadratically constrained quadratic program (sometimes called the complex constant-modulus QP [37, Table 2] in the optimization literature and non-bipartite Grothendieck problem in theoretical computer science):

$$\max_{x\in\mathbb{C}^n} \ x^*Cx, \quad \text{subject to } |x_1| = \cdots = |x_n| = 1, \tag{QP}$$

where x^* denotes the conjugate transpose of x. Of course, this problem can only be solved up to a global phase, since only relative information is available. Indeed, given any solution x, all vectors of the form $xe^{i\theta}$ are equivalent solutions, for arbitrary phase θ.

Solving (QP) is, in general, an NP-hard problem [55, Prop. 3.5]. It is thus unlikely that there exists an algorithm capable of solving (QP) in polynomial time for an arbitrary cost matrix C. In response, building upon now classical techniques, various authors [55, 48, 45, 1, 2, 10] study the following convex relaxation of (QP). For any admissible x, the Hermitian matrix $X = xx^* \in \mathbb{C}^{n\times n}$ is Hermitian positive semidefinite, has unit diagonal entries and is of rank one. Conversely, any such X may be written in the form $X = xx^*$ such that x is admissible for (QP). In this case, the cost function can be rewritten in linear form: $x^*Cx = \operatorname{Trace}(x^*Cx) = \operatorname{Trace}(CX)$. Dropping the rank constraint then yields the relaxation we set out to study:

$$\max_{X\in\mathbb{C}^{n\times n}} \ \operatorname{Trace}(CX), \quad \text{subject to } \operatorname{diag}(X) = \mathbb{1} \text{ and } X \succeq 0, \tag{SDP}$$

where $\mathbb{1} \in \mathbb{C}^n$ is the all-ones vector and $\operatorname{diag}\colon \mathbb{C}^{n\times n} \to \mathbb{C}^n$ extracts the diagonal entries of a matrix into a vector. Such relaxations *lift* the problem to higher dimensional spaces. Indeed, the search space of (QP) has dimension n (or $n-1$, discounting the global phase) whereas the search space of (SDP) has dimension $n(n-1)/2$. In general, increasing the dimension of an optimization problem may not be advisable. But in this case, the relaxed problem is a semidefinite program. Such optimization problems can be solved to global optimality up to arbitrary precision in polynomial time [51].

It is known that the solution of (SDP) can be rounded to an approximate solution of (QP), with a guaranteed approximation ratio [48, § 4]. But even better, when (SDP) admits an optimal solution X of rank one, then no rounding is necessary: the leading eigenvector x of $X = xx^*$ is a global optimum of (QP), meaning we have solved the original problem exactly. Elucidating when the semidefinite program admits a solution of rank one, i.e., when the relaxation is *tight*, is the focus of the present paper.

2.1 A detour through synchronization over $\mathbb{Z}_2 = \{\pm 1\}$

Problem (QP) is posed over the complex numbers. As a result, the individual variables x_i in (QP) live on a continuous search space (the unit circle). One effect of this is that even small noise on the data precludes exact recovery of the signal z (in general). This is the root of most of the complications that will arise in the developments hereafter. In order to first illustrate some of the ideas of the proof in a simpler context, this section proposes to take a detour through the real case. Besides this expository rationale, the real case is interesting in and of itself. It notably relates to correlation clustering [1] and the stochastic block model [2].

Let $z \in \{\pm 1\}^n$ be the signal to estimate and let $C = zz^\top + \sigma W$ contain the measurements, with $W = W^\top$ a Wigner matrix: its above-diagonal entries are i.i.d. (real) normal random variables, and its diagonal entries are zero. Each entry C_{ij} is a noisy measurement of the relative sign $z_i z_j$. For example, z_i could represent political preference of an agent (left or right wing) and C_{ij} could be a measurement of similarity between two agents' views [2]. Consider this MLE problem:

$$\max_{x \in \mathbb{R}^n} \; x^\top C x, \quad \text{subject to } |x_1| = \cdots = |x_n| = 1. \tag{QP$\mathbb{R}$}$$

Thus, $x_i \in \{\pm 1\}$ for each i. The corresponding relaxation reads:

$$\max_{X \in \mathbb{R}^{n \times n}} \; \operatorname{Trace}(CX), \quad \text{subject to } \operatorname{diag}(X) = \mathbb{1} \text{ and } X \succeq 0. \tag{SDP$\mathbb{R}$}$$

Certainly, if (SDP$\mathbb{R}$) admits a rank-one solution, then that solution reveals a global optimizer of (QP$\mathbb{R}$). Beyond rank recovery, because the variables x_i are discrete, it is expected that, if noise is sufficiently small, then the global optimizer of (QP$\mathbb{R}$) will be the true signal z. Figure 1 confirms this even for large noise, leveraging the tightness of the relaxation (SDP$\mathbb{R}$). Note that exact recovery is a strictly stronger requirement than rank recovery. We now investigate this exact recovery phenomenon for (SDP$\mathbb{R}$).

The aim is to show that $X = zz^\top$ is a solution of (SDP$\mathbb{R}$). Semidefinite programs admit a dual problem [51], which for (SDP$\mathbb{R}$) reads:

$$\min_{S \in \mathbb{R}^{n \times n}} \; \operatorname{Trace}(S + C), \text{ s.t. } S + C \text{ is diagonal and } S \succeq 0. \tag{DSDP$\mathbb{R}$}$$

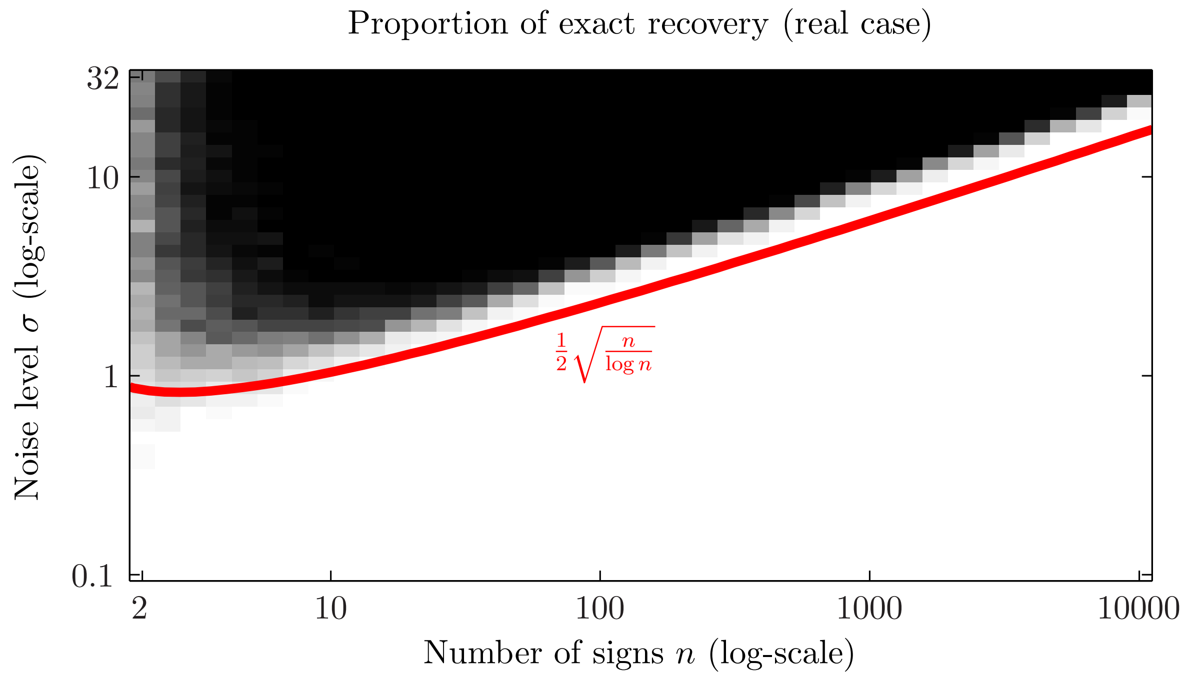

Figure 1: In the real case (QPℝ), one can hope to recover the signal $z \in \{\pm 1\}^n$ exactly from the pairwise sign comparisons $zz^\top + \sigma W$. This figure shows how frequently the semidefinite relaxation (SDPℝ) returns the correct signal z (or $-z$). For each pair (n, σ), 100 realizations of the noise W are generated independently and the dual certificate for the true signal (2.1) is verified (it is declared numerically positive semidefinite if its smallest eigenvalue exceeds $-10^{-14}n$). The frequency of success is coded by intensity (bright for 100% success, dark for 0% success). The results are in excellent agreement with the theoretical predictions (2.2).

Strong duality holds, which implies that a given feasible X is optimal if and only if there exists a dual feasible matrix S such that $\mathrm{Trace}\,(CX) = \mathrm{Trace}\,(S + C)$, or, in other words, such that $\mathrm{Trace}\,(SX) = 0$ (since $S + C$ is diagonal and the diagonal of X is $\mathbb{1}$). Since both S and C are positive semidefinite, this is equivalent to requiring $SX = 0$ (a condition known as *complementary slackness*), and hence requiring $Sz = 0$.

For ease of exposition, we now assume (without loss of generality) that $z = \mathbb{1}$. Tentatively, let $S = nI + \sigma\,\mathrm{diag}(W\mathbb{1}) - C$—for the complex case, we will see how to obtain this candidate without guessing. Then, by construction, $S + C$ is diagonal and $S\mathbb{1} = 0$. It remains to determine under what conditions S is positive semidefinite.

Define the diagonal matrix $D_W = \mathrm{diag}(W\mathbb{1})$ and the (Laplacian-like) matrix $L_W = D_W - W$. The candidate dual certificate is a sum of two Laplacian-like matrices:

$$S = nI - \mathbb{1}\mathbb{1}^\top + \sigma L_W. \tag{2.1}$$

The first part is the Laplacian of the complete graph. It has eigenvalues $0, n, \ldots, n$, with the zero eigenvalue corresponding to the all-ones vector $\mathbb{1}$. Since $L_W \mathbb{1} = 0$ by construction, S is positive semidefinite as long as σL_W does not "destroy" any of the large eigenvalues n. This is guaranteed in particular if $\|\sigma L_W\|_{\mathrm{op}} < n$. It is well-known from concentration results about Wigner matrices that, with high probability for growing n, $\|L_W\|_{\mathrm{op}} \leq 2\sqrt{n \log n}$ [26, Thm. 1]. Thus, it is expected that (SDP$\mathbb{R}$) will yield exact recovery of the signal z (with high probability as n grows) as long as

$$\sigma < \frac{1}{2}\sqrt{\frac{n}{\log n}}. \tag{2.2}$$

This is indeed compatible with the empirical observation of Figure 1.

2.2 Back to synchronization over SO(2)

We now return to the complex case, which is the focus of this paper. As was mentioned earlier, in the presence of even the slightest noise, one cannot reasonably expect the true signal z to be an optimal solution of (QP) anymore (this can be formalized using Cramér-Rao bounds [16]). Nevertheless, we set out to show that (under some assumptions on the noise) solutions of (QP) are close to z and they can be computed via (SDP).

The proof follows that of the real case in spirit, but requires more sophisticated arguments because the solution of (QP) is not known explicitly anymore. One effect of this is that the candidate dual certificate S will itself depend on the unknown solution of (QP). With that in mind, the proof of the upcoming main lemma (Lemma 3.2) follows this reasoning:

1. For small enough noise levels σ, any optimal solution x of (QP) is close to the sought signal z (Lemmas 4.1 and 4.2).

2. Solutions x are, a fortiori, local optimizers of (QP), and hence satisfy first-order necessary optimality conditions. These take up the form $Sx = 0$, where $S = \Re\{\mathrm{ddiag}(Cxx^*)\} - C$ depends smoothly on x (see (4.5)). Note that S is a function of x, the MLE (which is not explicitly known).

3. Remarkably, this Hermitian matrix S can be used as a dual certificate for solutions of (SDP). Indeed, $X = xx^*$ is a solution of (SDP) if and only if $Sx = 0$ and S is positive semidefinite—these are the Karush-Kuhn-Tucker conditions (KKT). Furthermore, that solution is unique if $\mathrm{rank}(S) = n - 1$ (Lemmas 4.3 and 4.4). Thus, it only remains to study the eigenvalues of S.

4. In the absence of noise, S is a Laplacian for a complete graph with unit weights (up to a unitary transformation), so that its eigenvalues are 0 with multiplicity 1 and n with multiplicity $n - 1$. Then, $X = zz^*$ is always the unique solution of (SDP).

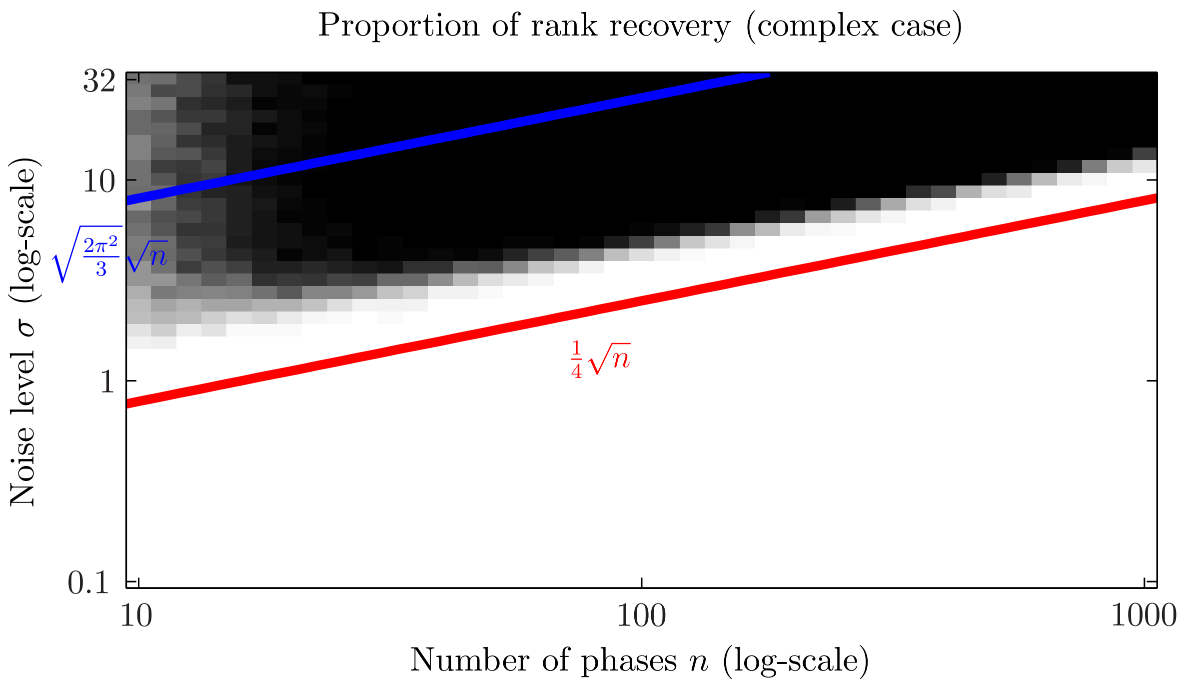

Figure 2: In the complex case (QP), *exact* recovery of the phases $z \in \{e^{i\theta} : \theta \in \mathbb{R}\}^n$ from the pairwise relative phase measurements $zz^* + \sigma W$ is hopeless as soon as $\sigma > 0$ [16]. Computing a maximum likelihood estimator (MLE) for z is still interesting. In particular, below the red line, with high probability, the MLE is closer to z than the estimator with maximum error (Lemma 4.1). Computing the MLE is hard in general, but solving the semidefinite relaxation (SDP) is easy. When (SDP) has a rank-one solution, that solution coincides with the MLE. This figure shows, empirically, how frequently the (SDP) admits a unique rank-one solution (same color code as Figure 1). For each pair (n, σ), 100 realizations of the noise W are generated independently and (SDP) is solved using a Riemannian optimization toolbox [15]. The (SDP) appears to be tight for remarkably large levels of noise. Theorem 2.2 (our main contribution) partly explains this phenomenon, by showing that σ can indeed grow unbounded while retaining rank recovery, although not at the rate witnessed here. We further note that, above the blue line, no unbiased estimator for z performs better than a random guess [16].

5. Adding small noise, because of the first point, the solution x will move only by a small amount, and hence so will S. Thus, the large eigenvalues should be controllable into remaining positive (Section 4.4).

6. The crucial fact follows: because of the way S is constructed (using first-order optimality conditions), the zero eigenvalue is "pinned down" (as long as x is a local optimum of (QP)). Indeed, both x and S change as a result of adding noise, but the property $Sx = 0$ remains valid. Thus, there is no risk that the zero eigenvalue from the noiseless scenario would become negative when noise is added.

Following this road map, most of the work in the proof below consists in bounding how far away x can be from z (as a function of the noise level and structure) and in using that to control the large eigenvalues of S.

Remark 2.1 (The role of smoothness)**.** *The third point in the road map, namely the special role of S, merits further comment. This way of identifying the dual certificate already appears explicitly in Journée et al. [34], who considered a different family of real, semidefinite programs which also admit a smooth geometry when the rank is constrained.*

In essence, KKT conditions capture the idea that, at a (local) optimizer, there is no escape direction that—up to first order—both preserves feasibility and improves the objective function. The KKT conditions for (SDP) *take up the classical form "if X is optimal, then there exists a dual certificate S which satisfies such and such conditions." For the purpose of certifying a solution analytically, this is impractical, because there is no explicit formula stating which S to verify. Fortunately,* (SDP) *is* nondegenerate *[6] (meaning, roughly, that its underlying geometry is smooth). This leads to uniqueness of dual certificates, and hence suggests there may be an analytical expression for the dual candidate.*

The convex problem (SDP) *is a relaxation of* (QP) *(up to the global phase). Hence, the KKT conditions for* (SDP) *at a rank-one solution xx^* are the KKT conditions for* (QP) *at x* plus *additional conditions: the latter ensure none of the new directions are improving directions either. Because* (QP) *is a smooth problem, the KKT conditions for* (QP) *are explicit: if x is a (local) optimizer of* (QP)*, then* $\operatorname{grad} g(x) = -2S(x)x = 0$ *(where $g(x) = x^*Cx$ and* $\operatorname{grad} g(x)$ *is its Riemannian gradient* (4.4)*).*

This gives an explicit expression for a candidate dual S that satisfies part of the KKT conditions of (SDP) *at xx^*. It then suffices to add the additional conditions of* (SDP) *(namely, that S be positive semidefinite) to obtain an explicit expression for the unique dual candidate.*

Our main theorem follows. In a nutshell, it guarantees that: under (complex) Wigner noise W, with high probability, solutions of (QP) are close to z, and, assuming the noise level σ is smaller than (on the order of) $n^{1/10}$, (SDP) admits a unique solution, it is of rank one and identifies the solution of (QP) (unique, up to a global phase shift).

Theorem 2.2. *Let $z \in \mathbb{C}^n$ be a vector with unit modulus entries, $W \in \mathbb{C}^{n\times n}$ a Hermitian Gaussian Wigner matrix and let $C = zz^* + \sigma W$. Let $x \in \mathbb{C}^n$ be a global optimizer of (QP). With probability at least $1 - 4n^{-\frac{1}{4}\log n + 2}$, the following is true. The (unidentifiable) global phase of x can be chosen such that x is close to z in the following two senses:*

$$\begin{aligned} \|x - z\|_\infty &\leq 2\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)\sigma n^{-1/4}, \text{ and} \\ \|x - z\|_2^2 &\leq 8\sigma n^{1/2} \min\left[1, 23\sigma\right]. \end{aligned}$$

Furthermore, there exists a universal constant $K > 0$ such that, if

$$\sigma + \sigma^{2/5} \leq K \frac{n^{1/10}}{\log(n)^{2/5}}, \tag{2.3}$$

then the semidefinite program (SDP), given by

$$\max_{X\in\mathbb{C}^{n\times n}} \operatorname{Trace}(CX), \quad \textit{subject to } \operatorname{diag}(X) = \mathbb{1} \textit{ and } X \succeq 0,$$

has, as its unique solution, the rank-one matrix $X = xx^*$.

Notice that the numerical experiments (Figure 2) suggest it should be possible to allow σ to grow at a rate of $\frac{n^{1/2}}{\operatorname{polylog}(n)}$ (as in the real case), but we were not able to establish that. Nevertheless, we do show that σ can grow unbounded with n. To the best of our knowledge, this is the first result of this kind. We hope it might inspire similar results in other problems where the same phenomenon has been observed [12].

Remark 2.3 (On the square-root rate)**.** *The relaxation (SDP) can be further relaxed by summarizing the constraints* $\operatorname{diag}(X) = \mathbb{1}$ *into the single constraint* $\operatorname{Trace}(X) = n$. *In so doing, the new relaxation always admits a rank-one solution $X = vv^*$ [40] such that v is a dominant eigenvector of C. The data C can be seen as a rank-one perturbation zz^* of a random matrix σW [45]. For i.i.d. Gaussian noise, as soon as the operator norm of the noise matrix is smaller than (twice) the operator norm of the signal (equal to n), dominant eigenvectors of C are expected to have better-than-random correlation with the signal z. Since the operator norm of such W's grows as $n^{1/2}$, this explains why, even when σ grows as $n^{1/2}$ itself, some signal is still present in the data.*

3 Main result

In this section we present our main technical result and show how it can be used to prove Theorem 2.2. Let us start by presenting a central definition in this paper. Intuitively, this definition characterizes non-adversarial noise matrices W.

Definition 3.1 (z-discordant matrix)**.** *Let $z \in \mathbb{C}^n$ be a vector with unit modulus entries. A matrix $W \in \mathbb{C}^{n\times n}$ is called* z-discordant *if it is Hermitian and satisfies all of the following:*

1. $\|W\|_{\mathrm{e},\infty} \leq \log(n)$,
2. $\|W\|_{\mathrm{op}} \leq 3n^{1/2}$,
3. $\|Wz\|_\infty \leq n^{3/4}$,
4. $|z^*Wz| \leq n^{3/2}$.

Recall that $\|\cdot\|_{\mathrm{e},\infty}$ is the entry-wise infinity norm. The next lemma is the main technical contribution of this paper. Note that it is a deterministic, non-asymptotic statement.

Lemma 3.2. *Let $z \in \mathbb{C}^n$ be a vector with unit modulus entries, let $W \in \mathbb{C}^{n\times n}$ be a Hermitian, z-discordant matrix (see Definition 3.1), and let $C = zz^* + \sigma W$. Let $x \in \mathbb{C}^n$ be a global optimizer of (QP). The (unidentifiable) global phase of x can be chosen such that x is close to z in the following two senses:*

$$\|x - z\|_\infty \leq 2\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)\sigma n^{-1/4}, \text{ and}$$
$$\|x - z\|_2^2 \leq 8\sigma n^{1/2}\min\left[1, 23\sigma\right].$$

Furthermore, there exists a universal constant $K > 0$ such that, if

$$\sigma + \sigma^{2/5} \leq K\frac{n^{1/10}}{\log(n)^{2/5}}, \tag{3.1}$$

then the semidefinite program (SDP), given by

$$\max_{X\in\mathbb{C}^{n\times n}} \ \mathrm{Trace}\,(CX), \quad \text{subject to } \mathrm{diag}(X) = \mathbb{1} \text{ and } X \succeq 0,$$

has, as its unique solution, the rank-one matrix $X = xx^*$.

We defer the proof of Lemma 3.2 to Section 4. The following proposition, whose proof we defer to Appendix A, shows how this lemma can be used to prove Theorem 2.2.

Proposition 3.3. *Let $z \in \mathbb{C}^n$ be a (deterministic) vector with unit modulus entries. Let $W \in \mathbb{C}^{n\times n}$ be a Hermitian, random matrix with i.i.d. off-diagonal entries following a complex normal distribution and zeros on the diagonal. Thus, $W_{ii} = 0$, $W_{ij} = \overline{W_{ji}}$, $\mathbb{E}W_{ij} = 0$ and $\mathbb{E}|W_{ij}|^2 = 1$ (for $i \neq j$). Then, W is z-discordant with probability at least $1 - 4n^{-\frac{1}{4}\log n + 2}$.*

The latter result is not surprising. Indeed, the definition of z-discordance requires two elements. Namely, (1) that W be not too large (properties 1 and 2), and (2) that W be not too aligned with z (properties 3 and 4). For W a Wigner matrix independent of z, those are indeed expected to hold.

The definition of z-discordance is not tightly adjusted to Wigner noise. As a result, it is expected that Lemma 3.2 will be applicable to show tightness of semidefinite relaxations for a larger span of noise models. Section 5 presents an application to the orthogonal Procrustes problem.

4 The proof

In this section, we prove Lemma 3.2. See Section 2.2 for an outline of the proof. To ease the algebra involved in the proofs, and without loss of generality, we consider throughout that $z = \mathbb{1}$. This corresponds to the deterministic change of variables $C \mapsto \operatorname{diag}(z)^* C \operatorname{diag}(z)$. Certainly, W is z-discordant if and only if $\operatorname{diag}(z)^* W \operatorname{diag}(z)$ is $\mathbb{1}$-discordant.

4.1 Global optimizers of (QP) are close to z

Let x be any global optimizer of (QP). Choosing the global phase of x such that $\mathbb{1}^* x \geq 0$, we decompose x as follows:

$$x = \mathbb{1} + \Delta, \tag{4.1}$$

where $\Delta \in \mathbb{C}^n$ should be thought of as an error term, as it represents how far a global optimizer of (QP) is from the planted signal $\mathbb{1}$. This subsection focuses on bounding Δ. We bound both its ℓ_2 and ℓ_∞ norms.

The following easy ℓ_2 bound is readily available:

$$\|\Delta\|_2^2 = \|x - \mathbb{1}\|_2^2 = x^*x + \mathbb{1}^*\mathbb{1} - 2\Re\{\mathbb{1}^*x\} = 2(n - \mathbb{1}^*x) \leq 2n. \tag{4.2}$$

The next lemma provides an improved bound when $\sigma \leq \frac{1}{4}n^{1/2}$.

Lemma 4.1. *If W is $\mathbb{1}$-discordant, then*

$$\|\Delta\|_2^2 \leq 8\sigma n^{1/2}.$$

Proof. If $\sigma \geq \frac{1}{4}n^{1/2}$, the bound is trivial since $\|\Delta\|_2^2 \leq 2n \leq 8\sigma n^{1/2}$. We now prove the bound under the complementary assumption that $\sigma \leq \frac{1}{4}n^{1/2}$.

Since x is a global maximizer of (QP) it must, in particular, satisfy $x^*Cx \geq \mathbb{1}^*C\mathbb{1}$. Hence,

$$x^*(\mathbb{1}\mathbb{1}^* + \sigma W)x \geq \mathbb{1}^*(\mathbb{1}\mathbb{1}^* + \sigma W)\mathbb{1},$$

or equivalently, $\sigma\left(x^*Wx - \mathbb{1}^*W\mathbb{1}\right) \geq \mathbb{1}^*\mathbb{1}\mathbb{1}^*\mathbb{1} - x^*\mathbb{1}\mathbb{1}^*x$. This readily implies that

$$\sigma\left(\|x\|_2^2\,\|W\|_{\mathrm{op}} + |\mathbb{1}^*W\mathbb{1}|\right) \geq n^2 - |\mathbb{1}^*x|^2.$$

Hence, using $\|x\|_2^2 = n$ and $\mathbb{1}$-discordance of W (more specifically, $\|W\|_{\mathrm{op}} \leq 3n^{1/2}$ and $|\mathbb{1}^*W\mathbb{1}| \leq n^{3/2}$) we have

$$|\mathbb{1}^*x|^2 \geq n^2 - 3\sigma n^{3/2} - \sigma n^{3/2} = n^2 - 4\sigma n^{3/2}.$$

Since $\mathbb{1}^*x \geq 0$, under the assumption that $\sigma \leq \frac{1}{4}n^{1/2}$, we actually have:

$$\mathbb{1}^*x \geq \sqrt{n^2 - 4\sigma n^{3/2}}.$$

Combine the latter with the fact that $\|\Delta\|_2^2 = \|x - \mathbb{1}\|_2^2 = 2(n - \mathbb{1}^* x)$ to obtain

$$\|\Delta\|_2^2 \leq 2\left(n - \sqrt{n^2 - 4\sigma n^{3/2}}\right) \leq 8\sigma n^{1/2}.$$

The last inequality follows from $a - \sqrt{b} = \left(a - \sqrt{b}\right)\frac{a+\sqrt{b}}{a+\sqrt{b}} = \frac{a^2-b}{a+\sqrt{b}} \leq a - b/a$ for all $a > 0$ and $b \geq 0$ such that $a^2 \geq b$. □

The next lemma establishes a bound on the largest individual error, $\|\Delta\|_\infty$. This is informative for values of n and σ such that the bound is smaller than 2. Interestingly, for a fixed value of σ, the bound shows that increasing n drives Δ to 0, uniformly.

Lemma 4.2. *If W is $\mathbb{1}$-discordant, then*

$$\|\Delta\|_\infty \leq 2\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)\sigma n^{-1/4}.$$

Proof. We wish to upper bound, for all $i \in \{1, 2, \ldots, n\}$, the value of $|\Delta_i|$. Let $e_i \in \mathbb{R}^n$ denote the i^{th} vector of the canonical basis (its i^{th} entry is 1 whereas all other entries are zero). Consider $\hat{x} = x + (1 - x_i)e_i$, a feasible point of (QP) obtained from the optimal x by changing one of its entries to 1. Since x is optimal, it performs at least as well as $\hat{x}$ according to the cost function of (QP):

$$x^*Cx \geq \hat{x}^*C\hat{x} = x^*Cx + |1 - x_i|^2 C_{ii} + 2\Re\{(1 - \overline{x_i})e_i^*Cx\}.$$

Further develop the last term by isolating the diagonal term C_{ii}:

$$2\Re\{(1 - \overline{x_i})e_i^*Cx\} = 2C_{ii}\Re\{x_i - 1\} + 2\Re\Big\{(1 - \overline{x_i})\sum_{j\neq i} C_{ij}x_j\Big\}.$$

Since $|1 - x_i|^2 C_{ii} = -2C_{ii}\Re\{x_i - 1\}$, combining the two equations above yields the following inequality:

$$\Re\Big\{(\overline{x_i} - 1)\sum_{j\neq i}(1 + \sigma W_{ij})x_j\Big\} \geq 0.$$

Injecting $x = \mathbb{1} + \Delta$ we get:

$$\Re\Big\{\overline{\Delta_i}\sum_{j\neq i}(1 + \sigma W_{ij})(1 + \Delta_j)\Big\} \geq 0.$$

Expand the product, remembering that $W_{ii} = 0$ by definition, to obtain:

$$\begin{aligned}(n-1)\Re\{\overline{\Delta_i}\} &\geq -\Re\Big\{\overline{\Delta_i}\sum_{j\neq i}\left(\sigma W_{ij} + \Delta_j + \sigma W_{ij}\Delta_j\right)\Big\} \\ &= |\Delta_i|^2 - \Re\Big\{\overline{\Delta_i}\sum_{j}\left(\sigma W_{ij} + \Delta_j + \sigma W_{ij}\Delta_j\right)\Big\}.\end{aligned}$$

At this point, recall we want to bound $|\Delta_i|$. Since

$$|\Delta_i|^2 = |x_i - 1|^2 = 2(1 - \Re\{x_i\}) = -2\Re\{\Delta_i\} = -2\Re\{\overline{\Delta_i}\},$$

the above inequality is equivalent to:

$$\begin{aligned}
|\Delta_i|^2 &\leq \frac{2}{n+1} \Re\Big\{\overline{\Delta_i} \sum_j (\sigma W_{ij} + \Delta_j + \sigma W_{ij} \Delta_j)\Big\} \\
&\leq \frac{2}{n+1} |\Delta_i| \left| \sum_j (\sigma W_{ij} + \Delta_j + \sigma W_{ij} \Delta_j) \right| \\
&\leq \frac{2}{n+1} |\Delta_i| \left(\sigma \, |e_i^* W \mathbb{1}| + |\mathbb{1}^* \Delta| + \sigma |e_i^* W \Delta| \right) \\
&\leq \frac{2}{n+1} |\Delta_i| \left(\sigma \|W\mathbb{1}\|_\infty + \frac{1}{2} \|\Delta\|_2^2 + \sigma \, \|W\|_{\mathrm{op}} \, \|\Delta\|_2 \right),
\end{aligned}$$

where we used the triangular inequality multiple times and the simple identity $\mathbb{1}^*\Delta = -\|\Delta\|_2^2/2$. The above inequality holds for all $1 \leq i \leq n$. We now leverage the $\mathbb{1}$-discordance of W (more precisely, $\|W\mathbb{1}\|_\infty \leq n^{3/4}$ and $\|W\|_{\mathrm{op}} \leq 3n^{1/2}$) together with Lemma 4.1 to finally obtain:

$$\begin{aligned}
\|\Delta\|_\infty &\leq \frac{2}{n+1} \left(\sigma n^{3/4} + 4\sigma n^{1/2} + 6\sqrt{2}\, \sigma^{3/2} n^{3/4} \right) \\
&\leq 2 \left(1 + 4n^{-1/4} + 6\sqrt{2}\, \sigma^{1/2} \right) \sigma n^{-1/4} \\
&\leq 2 \left(5 + 6\sqrt{2}\, \sigma^{1/2} \right) \sigma n^{-1/4}.
\end{aligned}$$

This concludes the proof. $\square$

As a side note, notice that, using the bound on $\|\Delta\|_\infty$, one obtains another bound on $\|\Delta\|_2^2$ as follows:

$$\|\Delta\|_2^2 \leq n\|\Delta\|_\infty^2 \leq 8\sigma n^{1/2} \left[\frac{(5 + 6\sqrt{2}\sigma^{1/2})^2}{2} \sigma \right]. \tag{4.3}$$

The factor in brackets is an increasing function of σ that hits 1 for $\sigma \approx 0.0436$. Below that value, the above bound improves on Lemma 4.1 and the factor in brackets is bounded by 23σ, thus yielding the bound as stated in Lemma 3.2. Nevertheless, in the remainder of the paper, we use only Lemma 4.1 to bound $\|\Delta\|_2^2$. This is because we aim to allow σ to grow with n, and Lemma 4.1 is sharper in that regime. The interest of the above bound is to show that, for small noise, the root mean squared error $\|\Delta\|_2/\sqrt{n}$ is at most on the order of $\sigma/n^{1/4}$.

4.2 Optimality conditions for (SDP)

The global optimizers of the semidefinite program (SDP) can be characterized completely via the Karush-Kuhn-Tucker (KKT) conditions:

Lemma 4.3. *A Hermitian matrix $X \in \mathbb{C}^{n\times n}$ is a global optimizer of (SDP) if and only if there exists a Hermitian matrix $\hat{S} \in \mathbb{C}^{n\times n}$ such that all of the following hold:*

1. $\operatorname{diag}(X) = \mathbb{1}$;
2. $X \succeq 0$;
3. $\hat{S}X = 0$;
4. $\hat{S} + C$ *is (real) diagonal; and*
5. $\hat{S} \succeq 0$.

If, furthermore, $\operatorname{rank}(\hat{S}) = n-1$, *then X has rank one and is the unique global optimizer of (SDP).*

Proof. These are the KKT conditions of (SDP) [43, Example 3.36]. Conditions 1 and 2 are primal feasibility, condition 3 is complementary slackness and conditions 4 and 5 encode dual feasibility. Since the identity matrix I_n satisfies all equality constraints and is (strictly) positive definite, the so-called *Slater condition* is fulfilled. This ensures that the KKT conditions stated above are necessary and sufficient for global optimality [43, Theorem 3.34]. Slater's condition also holds for the dual. Indeed, let $\tilde{S} = \alpha I - C$, where $\alpha \in \mathbb{R}$ is such that $\tilde{S} \succ 0$ (such an α always exists); then $\tilde{S} + C$ is indeed diagonal and $\tilde{S}$ is strictly admissible for the dual. This allows to use results from [6]. Specifically, assuming $\operatorname{rank}(\hat{S}) = n - 1$, Theorem 9 in that reference implies that $\hat{S}$ is *dual nondegenerate*. Then, since $\hat{S}$ is also optimal for the dual (by complementary slackness), Theorem 10 in that reference guarantees that the primal solution X is unique. Since X is nonzero and $\hat{S}X = 0$, it must be that $\operatorname{rank}(X) = 1$. □

Certainly, if (SDP) admits a rank-one solution, it has to be of the form $X = xx^*$, with x an optimal solution of the original problem (QP). Based on this consideration, our proof of Lemma 3.2 goes as follows. We let x denote a global optimizer of (QP) and we consider $X = xx^*$ as a candidate solution for (SDP). Using the optimality of x and assumptions on the noise, we then construct and verify a dual certificate matrix S as required per Lemma 4.3. In such proofs, one of the nontrivial parts is to guess an analytical form for S given a candidate solution X. We achieve this by inspecting the first-order optimality conditions of (QP) (which x necessarily satisfies). The main difficulty is then to show the suitability of the candidate S, as it depends nonlinearly on the global optimum x, which itself is a complicated function of the noise W. Nevertheless, we show feasibility of S via a program of inequalities, relying heavily on the $\mathbb{1}$-discordance of the noise W (see Definition 3.1).

4.3 Construction of the dual certificate S

Every global optimizer of the combinatorial problem (QP) must, a fortiori, satisfy first-order necessary optimality conditions. We derive those now.

We endow the complex plane $\mathbb{C}$ with the Euclidean metric

$$\langle y_1, y_2 \rangle = \Re\{y_1^* y_2\}.$$

This is equivalent to viewing $\mathbb{C}$ as $\mathbb{R}^2$ with the canonical inner product, using the real and imaginary parts of a complex number as its first and second coordinates. Denote the complex circle by

$$\mathcal{S} = \{y \in \mathbb{C} : y^* y = 1\}.$$

The circle can be seen as a submanifold of $\mathbb{C}$, with tangent space at each y given by (simply differentiating the constraint):

$$T_y\mathcal{S} = \{\dot{y} \in \mathbb{C} : \dot{y}^* y + y^* \dot{y} = 0\} = \{\dot{y} \in \mathbb{C} : \langle y, \dot{y} \rangle = 0\}.$$

Restricting the Euclidean inner product to each tangent space equips $\mathcal{S}$ with a Riemannian submanifold geometry. The search space of (QP) is exactly $\mathcal{S}^n$, itself a Riemannian submanifold of $\mathbb{C}^n$ with the product geometry. Thus, problem (QP) consists in maximizing a smooth function $g(x) = x^* C x$ over the smooth Riemannian manifold $\mathcal{S}^n$. Therefore, the first-order necessary optimality conditions for (QP) (i.e., the KKT conditions) can be stated simply as $\operatorname{grad} g(x) = 0$, where $\operatorname{grad} g(x)$ is the Riemannian gradient of g at $x \in \mathcal{S}^n$ [3]. This gradient is given by the orthogonal projection of the Euclidean (the classical) gradient of g onto the tangent space of $\mathcal{S}^n$ at x [3, eq. (3.37)]. The projector and the Euclidean gradient are given respectively by:

$$\begin{aligned}
&\operatorname{Proj}_x \colon \mathbb{C}^n \to T_x\mathcal{S}^n \colon \dot{x} \mapsto \operatorname{Proj}_x \dot{x} = \dot{x} - \Re\{\operatorname{ddiag}(\dot{x}x^*)\}x, \\
&\nabla g(x) = 2Cx,
\end{aligned}$$

where $\operatorname{ddiag} \colon \mathbb{C}^{n\times n} \to \mathbb{C}^{n\times n}$ sets all off-diagonal entries of a matrix to zero. For x a global optimizer of (QP), it holds that

$$0 = \operatorname{grad} g(x) = \operatorname{Proj}_x \nabla g(x) = 2(C - \Re\{\operatorname{ddiag}(Cxx^*)\})x. \tag{4.4}$$

This suggests the following definitions:

$$X = xx^*, \qquad\qquad S = \Re\{\operatorname{ddiag}(Cxx^*)\} - C. \tag{4.5}$$

Note that S is Hermitian and $Sx = 0$. Referring to the KKT conditions in Lemma 4.3, it follows immediately that X is feasible for (SDP) (conditions 1 and 2); that $SX = (Sx)x^* = 0$ (condition 3); and that $S+C$ is a diagonal matrix (condition 4). It thus only remains to show that S is also positive semidefinite and has rank $n-1$. If such is the case, then X is the unique global optimizer of (SDP), meaning that solving the latter solves (QP). Note the special role of the first-order necessary optimality conditions: they guarantee complementary slackness, without requiring further work.

The following lemma further shows that S is the "right" candidate dual certificate. More precisely, for x a critical point of (QP), it is necessary and sufficient for S to be positive semidefinite in order for $X = xx^*$ to be optimal for (SDP).

Lemma 4.4. *X (of any rank) is optimal for (SDP) if and only if it is feasible for (SDP) and $S = \Re\{\mathrm{ddiag}(CX)\} - C$ (4.5) is positive semidefinite and $SX = 0$. There exists no other dual certificate for X.*

Proof. The *if* part follows from Lemma 4.3. We show the *only if* part. Assume X is optimal. Then, by Lemma 4.3, there exists $\hat{S} \succeq 0$ which satisfies $\hat{S}X = 0$ and $\hat{S}+C = \hat{D}$, where $\hat{D}$ is diagonal. Thus, $CX = (\hat{D} - \hat{S})X = \hat{D}X$ and $\Re\{\mathrm{ddiag}(CX)\} = \hat{D}$. Consequently, $S = \hat{D} - C = \hat{S}$. $\square$

4.4 A sufficient condition for S to be positive semidefinite with rank $n-1$

First, observe that the diagonal matrix $\mathrm{diag}(x)$ is a unitary matrix. Thus,

$$Q = \mathrm{diag}(x)^* S \,\mathrm{diag}(x) = \Re\{\mathrm{ddiag}(Cxx^*)\} - \mathrm{diag}(x)^* C \,\mathrm{diag}(x) \tag{4.6}$$

is a Hermitian matrix whose spectrum is the same as that of S. In particular, S and Q share the same rank and they are simultaneously positive semidefinite, so that we now investigate Q.

Since $Sx = 0$, it follows that $Q\mathbb{1} = 0$ and Q is positive semidefinite with rank $n-1$ if and only if $u^*Qu > 0$ for all $u \in \mathbb{C}^n$ such that $u \neq 0$ and $\mathbb{1}^*u = 0$. We set out to find sufficient conditions for the latter.

To this end, separate u in its real and imaginary parts as $u = \alpha + i\beta$, with $\alpha, \beta \in \mathbb{R}^n$ satisfying $\mathbb{1}^\top\alpha = \mathbb{1}^\top\beta = 0$. The quadratic form expands as:

$$\begin{aligned} u^*Qu &= (\alpha^\top - i\beta^\top)Q(\alpha + i\beta) \\ &= \alpha^\top Q\alpha + \beta^\top Q\beta + i(\alpha^\top Q\beta - \beta^\top Q\alpha) \\ &= \alpha^\top \Re\{Q\}\alpha + \beta^\top \Re\{Q\}\beta - 2\alpha^\top \Im\{Q\}\beta. \end{aligned} \tag{4.7}$$

Let us inspect the last term more closely:

$$\begin{aligned} \Im\{Q\} &= -\Im\{\mathrm{diag}(x)^* C \,\mathrm{diag}(x)\} \\ &= -\Im\{\mathrm{diag}(x)^* \mathbb{1}\mathbb{1}^* \,\mathrm{diag}(x) + \sigma \,\mathrm{diag}(x)^* W \,\mathrm{diag}(x)\} \\ &= \Im\{xx^*\} - \sigma \Im\{\mathrm{diag}(x)^* W \,\mathrm{diag}(x)\}. \end{aligned}$$

At this step, we leverage the fact that, as per lemmas 4.1 and 4.2, if the noise level σ is small enough, then x is close to $z = \mathbb{1}$. We continue with the global phase convention $\mathbb{1}^*x \geq 0$ and the notation $x = \mathbb{1} + \Delta$ (4.1). Owing to α and β having zero mean components, it follows that

$$\begin{aligned} \alpha^\top \Im\{xx^*\}\beta &= \alpha^\top \Im\{\mathbb{1}\mathbb{1}^* + \mathbb{1}\Delta^* + \Delta\mathbb{1}^* + \Delta\Delta^*\}\beta \\ &= \alpha^\top \Im\{\Delta\Delta^*\}\beta. \end{aligned} \tag{4.8}$$

Keeping in mind the intuition that Δ is an error term, we expect (4.8) to be small. We will make this precise later, and now turn our attention to the real part of Q, which turns out to be a graph Laplacian. Indeed, this structure becomes apparent when the individual entries of the matrix are written out explicitly, starting from (4.6):

$$\Re\{Q\}_{ij} = \begin{cases} \Re\{(Cxx^*)_{ii} - C_{ii}\} = \sum_{\ell \neq i} \Re\{\bar{x}_i x_\ell C_{i\ell}\} & \text{if } i = j, \\ -\Re\{\bar{x}_i x_j C_{ij}\} & \text{if } i \neq j. \end{cases} \tag{4.9}$$

We recognize the Laplacian of a graph with n nodes and weight $\Re\{\bar{x}_i x_j C_{ij}\} = \langle x_i \bar{x}_j, C_{ij} \rangle$ on the edge (i, j). For a more compact notation, let

$$\mathcal{L} \colon \mathbb{C}^{n\times n} \to \mathbb{C}^{n\times n} \colon A \mapsto \mathcal{L}(A) = \operatorname{diag}(A\mathbb{1}) - A \tag{4.10}$$

be the linear map from a graph's adjacency matrix A to its Laplacian. Then,

$$\Re\{Q\} = \mathcal{L}(\Re\{\overline{C} \odot xx^*\}), \tag{4.11}$$

where $\odot$ denotes the entry-wise (or Hadamard) product of matrices. Certainly, the symmetric matrix $\Re\{Q\}$ admits a zero eigenvalue associated to the all-ones vector, since for any A, $\mathcal{L}(A)\mathbb{1} = 0$. We define the *spectral gap* of the Laplacian $\Re\{Q\}$ as its smallest eigenvalue associated to an eigenvector orthogonal to the all-ones vector:

$$\lambda_{\Re\{Q\}} = \min_{v \in \mathbb{R}^n, v \neq 0, \mathbb{1}^\top v = 0} \frac{v^\top \Re\{Q\} v}{v^\top v}. \tag{4.12}$$

Although this value could, in principle, be negative due to potential negative weights, the hope is that it will be positive and rather large (again, to be made precise later).

We now return to (4.7) and bound the expression:

$$\begin{aligned} u^*Qu &= \alpha^\top \Re\{Q\}\alpha + \beta^\top \Re\{Q\}\beta \\ &\qquad - 2\alpha^\top \Big(\Im\{\Delta\Delta^*\} - \sigma \Im\{\operatorname{diag}(x)^* W \operatorname{diag}(x)\} \Big) \beta \\ &\geq \left(\|\alpha\|_2^2 + \|\beta\|_2^2 \right) \lambda_{\Re\{Q\}} \\ &\qquad - 2\|\alpha\|_2 \|\beta\|_2 \Big(\|\Im\{\Delta\Delta^*\}\|_{\mathrm{op}} + \sigma \left\|\Im\{\operatorname{diag}(x)^* W \operatorname{diag}(x)\}\right\|_{\mathrm{op}} \Big). \end{aligned}$$

For this inequality to lead to a guarantee of positivity of u^*Qu, it is certainly necessary to require $\lambda_{\Re\{Q\}} > 0$. Using this and the facts that $\|\Im\{A\}\|_{\mathrm{op}} \leq \|A\|_{\mathrm{op}}$, that the operator norm is invariant under unitary transformations and the simple inequality

$$0 \leq (\|\alpha\|_2 - \|\beta\|_2)^2 = \|\alpha\|_2^2 + \|\beta\|_2^2 - 2\|\alpha\|_2\|\beta\|_2,$$

it follows that

$$u^*Qu \geq 2\|\alpha\|_2\|\beta\|_2 \left(\lambda_{\Re\{Q\}} - \|\Delta\|_2^2 - \sigma \|W\|_{\mathrm{op}} \right).$$

$$\|A\|_{\mathrm{op}}^2 = \max_{\substack{x \in \mathbb{C}^n \\ \|x\|=1}} \|Ax\|^2 \geq \max_{\substack{x \in \mathbb{R}^n \\ \|x\|=1}} \|Ax\|^2 = \max_{\substack{x \in \mathbb{R}^n \\ \|x\|=1}} \|\Re(A)x\|^2 + \|\Im(A)x\|^2 \geq \|\Im(A)\|_{\mathrm{op}}^2.$$

Hence, a sufficient condition for S to be positive semidefinite with rank $n-1$ is:

$$\lambda_{\Re\{Q\}} > \|\Delta\|_2^2 + \sigma \|W\|_{\mathrm{op}}. \tag{4.13}$$

Let us now pause to reflect on condition (4.13) and to describe why it should hold. A bound on the operator norm of W is readily available from $\mathbb{1}$-discordance of W, and $\|\Delta\|_2^2$ is bounded by Lemma 4.1. Perhaps less obvious is why one would expect $\lambda_{\Re\{Q\}}$ to be large. The intuition is that, for small enough noise, $x_i\bar{x}_j \approx z_i\bar{z}_j \approx C_{ij}$, so that $\Re\{Q\}$ is the Laplacian of a complete graph with large weights $\langle x_i\bar{x}_j, C_{ij}\rangle$. If this is the case, then it is known from graph theory that $\lambda_{\Re\{Q\}}$ is large, because the underlying graph is well connected. The bound derived below on the spectral gap will, together with (4.13), reveal how large we may allow the noise level σ to be.

4.5 Bounding the spectral gap of $\Re\{Q\}$

This section is dedicated to lower bounding the spectral gap term (4.12). The right-hand side of (4.13) is on the order of $\sigma n^{1/2}$, so that showing that the spectral gap is at least on the order of $n - \mathcal{O}(\sigma n^{1/2})$ would yield an acceptable noise level of $\mathcal{O}(n^{1/2})$ for σ, as the numerical experiment suggests (Figure 2) and similarly to the real case (Figure 1). Unfortunately, the bound we establish here is not as good, and thus constitutes the bottleneck in our analysis.

Lemma 4.5. *If W is $\mathbb{1}$-discordant, then*

$$\lambda_{\Re\{Q\}} \geq n - \left[8\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)^2 \sigma n^{-1/4} + \left(6 + 40\sigma + 68\sigma^{3/2}\right)\log(n)\right]\sigma n^{3/4}.$$

Proof. Working from equation (4.11), we find that

$$\begin{aligned}
\Re\{Q\} &= \mathcal{L}\left(\Re\{\overline{C} \odot xx^*\}\right) \\
&= \mathcal{L}\left(\Re\left\{xx^*\right\}\right) + \sigma\mathcal{L}\left(\Re\left\{\overline{W} \odot (\mathbb{1}\mathbb{1}^* + xx^* - \mathbb{1}\mathbb{1}^*)\right\}\right) \\
&= \mathcal{L}\left(\Re\left\{xx^*\right\}\right) + \sigma\mathcal{L}\left(\Re\left\{W\right\}\right) + \sigma\mathcal{L}\left(\Re\left\{\overline{W} \odot (xx^* - \mathbb{1}\mathbb{1}^*)\right\}\right).
\end{aligned}$$

Factor in the fact that for any $n \times n$ matrix A,

$$\begin{aligned}
\|\mathcal{L}(A)\|_{\mathrm{op}} &= \|\mathrm{diag}(A\mathbb{1}) - A\|_{\mathrm{op}} \leq \|A\mathbb{1}\|_\infty + \|A\|_{\mathrm{op}}, \text{ and} \\
\|A\|_{\mathrm{op}} &\leq \|A\|_{\mathrm{F}} \leq n\,\|A\|_{\mathrm{e},\infty}
\end{aligned}$$

Even assuming (incorrectly) that $\Delta = 0$, so that $x = z = \mathbb{1}$, we would only get a spectral gap of $n - \mathcal{O}(\sigma n^{3/4})$ (because of the bound on $\|W\mathbb{1}\|_\infty$), yielding a final acceptable rate of $\sigma = \mathcal{O}(n^{1/4})$, which still falls short of the target rate $\mathcal{O}(n^{1/2})$ (up to log factors).

to further obtain:

$$\begin{aligned}\lambda_{\Re\{Q\}} &\geq \lambda_{\mathcal{L}(\Re\{xx^*\})} - \sigma\left(\left\|\mathcal{L}\left(\Re\left\{W\right\}\right)\right\|_{\mathrm{op}} + \left\|\mathcal{L}\left(\Re\left\{\overline{W}\odot\left(xx^*-\mathbb{1}\mathbb{1}^*\right)\right\}\right)\right\|_{\mathrm{op}}\right)\\
&\geq \lambda_{\mathcal{L}(\Re\{xx^*\})} - \sigma\left(\left\|W\mathbb{1}\right\|_{\infty} + \left\|W\right\|_{\mathrm{op}}\right)\\
&\quad - \sigma\left(\left\|\left(\overline{W}\odot\left(xx^*-\mathbb{1}\mathbb{1}^*\right)\right)\mathbb{1}\right\|_{\infty} + n\left\|\overline{W}\odot\left(xx^*-\mathbb{1}\mathbb{1}^*\right)\right\|_{\mathrm{e},\infty}\right)\\
&\geq \lambda_{\mathcal{L}(\Re\{xx^*\})} - \sigma\left(\left\|W\mathbb{1}\right\|_{\infty} + \left\|W\right\|_{\mathrm{op}} + 2n\left\|W\right\|_{\mathrm{e},\infty}\left\|xx^*-\mathbb{1}\mathbb{1}^*\right\|_{\mathrm{e},\infty}\right).\end{aligned}$$

We now rely on Lemma 4.2 to bound $\left\|xx^*-\mathbb{1}\mathbb{1}^*\right\|_{\mathrm{e},\infty}$. For all $1 \leq i,j \leq n$,

$$|x_i\overline{x_j} - 1| = |x_i - x_j| = |(x_i - 1) - (x_j - 1)| \leq |\Delta_i| + |\Delta_j| \leq 2\|\Delta\|_\infty.$$

Thus,

$$\left\|xx^*-\mathbb{1}\mathbb{1}^*\right\|_{\mathrm{e},\infty} \leq 2\|\Delta\|_\infty \leq 4\left(5+6\sqrt{2}\,\sigma^{1/2}\right)\sigma n^{-1/4}. \tag{4.14}$$

Combining the last equations with $\mathbb{1}$-discordance of W and the fact that for $n \geq 2$ we have $4 \leq 6\log(n)$, gives:

$$\begin{aligned}\lambda_{\Re\{Q\}} &\geq \lambda_{\mathcal{L}(\Re\{xx^*\})} - \sigma\left(n^{3/4} + 3n^{1/2} + 8\left(5+6\sqrt{2}\,\sigma^{1/2}\right)\sigma n^{3/4}\log(n)\right)\\
&\geq \lambda_{\mathcal{L}(\Re\{xx^*\})} - \left(4 + 8\left(5+6\sqrt{2}\,\sigma^{1/2}\right)\sigma\log(n)\right)\sigma n^{3/4}\\
&\geq \lambda_{\mathcal{L}(\Re\{xx^*\})} - \left(6 + 8\left(5+6\sqrt{2}\,\sigma^{1/2}\right)\sigma\right)\sigma n^{3/4}\log(n)\\
&\geq \lambda_{\mathcal{L}(\Re\{xx^*\})} - \left(6 + 40\sigma + 68\sigma^{3/2}\right)\sigma n^{3/4}\log(n).\end{aligned} \tag{4.15}$$

It remains to bound the dominating part of the spectral gap. To this effect, we use the fact that $\Re\left\{x_i\overline{x_j}\right\}$ is nonnegative when the noise level is low enough, so that (restricting ourselves to that regime) $\lambda_{\mathcal{L}(\Re\{xx^*\})}$ is the spectral gap of a complete graph with all weights strictly positive. That spectral gap must be at least as large as the smallest weight multiplied by the spectral gap of the complete graph with unit weights, namely, $\lambda_{\mathcal{L}(\mathbb{1}\mathbb{1}^\top)} = n$. Formally, for all $v \in \mathbb{R}^n$ such that $\|v\|_2 = 1$ and $\mathbb{1}^\top v = 0$, by properties of Laplacian matrices it holds that

$$\begin{aligned}v^\top\mathcal{L}\left(\Re\left\{xx^*\right\}\right)v &= \sum_{i<j}\Re\{x_i\overline{x_j}\}(v_i - v_j)^2\\
&\geq \min_{i,j}\Re\{x_i\overline{x_j}\}\sum_{i<j}(v_i - v_j)^2\\
&= \min_{i,j}\Re\{x_i\overline{x_j}\}\; v^\top\mathcal{L}(\mathbb{1}\mathbb{1}^\top)v\\
&= n\min_{i,j}\Re\{x_i\overline{x_j}\}.\end{aligned}$$

Let us investigate the smallest weight. Recall that $|x_i - x_j| \leq 2\|\Delta\|_\infty$, $\Re\{x_i\overline{x_j}\} = 1 - \frac{1}{2}|x_i - x_j|^2$, and that $\|\Delta\|_\infty \leq 2\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)\sigma n^{-1/4}$ to get

$$\Re\{x_i\overline{x_j}\} \geq 1 - 2\|\Delta\|_\infty^2 \geq 1 - 8\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)^2 \sigma^2 n^{-1/2}.$$

Hence,

$$\lambda_{\mathcal{L}(\Re\{xx^*\})} \geq n - 8\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)^2 \sigma^2 n^{1/2}. \tag{4.16}$$

Merging the bounds (4.15) and (4.16) gives:

$$\begin{aligned}
\lambda_{\Re\{Q\}} &\geq n - 8\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)^2 \sigma^2 n^{1/2} - \left(6 + 40\sigma + 68\sigma^{3/2}\right)\sigma n^{3/4}\log(n) \\
&\geq n - \left[8\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)^2 \sigma n^{-1/4} + \left(6 + 40\sigma + 68\sigma^{3/2}\right)\log(n)\right]\sigma n^{3/4}.
\end{aligned}$$

This establishes the lemma. □

4.6 Concluding the proof

Recall that (SDP) is tight, in particular, if (4.13) holds:

$$\lambda_{\Re\{Q\}} > \|\Delta\|_2^2 + \sigma\left\|W\right\|_{\mathrm{op}}. \tag{4.17}$$

Still assuming $\mathbb{1}$-discordance of W (as it gives $\left\|W\right\|_{\mathrm{op}} \leq 3n^{1/2}$) and collecting results from lemmas 4.1 and 4.5, we find that this condition is fulfilled in particular if

$$n - \left[8\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)^2 \sigma n^{-1/4} + \left(6 + 40\sigma + 68\sigma^{3/2}\right)\log(n)\right]\sigma n^{3/4} > 11\sigma n^{1/2}.$$

Reorder terms and divide through by $n^{3/4}$ to get the equivalent condition:

$$n^{1/4} > \left[\left(8\left(5 + 6\sqrt{2}\,\sigma^{1/2}\right)^2 \sigma + 11\right) n^{-1/4} + \left(6 + 40\sigma + 68\sigma^{3/2}\right)\log(n)\right]\sigma.$$

This can be written in the form (with some constant $c_1 > 0$):

$$n^{1/4} > f_1(\sigma) n^{-1/4} + f_2(\sigma)\log(n) + c_1\sigma^3 n^{-1/4},$$

where f_1 and f_2 are polynomials with nonnegative coefficients, lowest power σ and highest power $\sigma^{5/2}$. Thus, their sum is upper-bounded by $c_2(\sigma + \sigma^{5/2})$, for some constant $c_2 > 0$. Hence, there exists a constant $c_3 > 0$ such that, if

$$c_1\sigma^3 n^{-1/4} < 0.99 n^{1/4}, \tag{4.18}$$

then

$$n^{1/4} > c_3(\sigma + \sigma^{5/2}) \log(n)$$

is a sufficient condition. It is then easy to see that there exists a universal constant $K > 0$ such that

$$\sigma + \sigma^{2/5} \leq K \frac{n^{1/10}}{\log(n)^{2/5}}$$

is a sufficient condition for tightness of (SDP). Indeed, since $n^{-1/4} < 1.3 \log(n)$ for all $n \geq 2$, it is possible to make K small enough that, for all $n \geq 2$, the above condition also implies (4.18)—the latter essentially requiring $\sigma = \mathcal{O}(n^{1/8})$, a looser condition indeed. This concludes the proof of Lemma 3.2.

5 The orthogonal Procrustes problem

The definition of z-discordance (Definition 3.1) is not tightly fitted to the idealized scenario where W has i.i.d. Gaussian entries. In fact, for the observation model $C = zz^* + \tau W$, Lemma 3.2 gives a tightness result as soon as $\frac{\tau}{\sigma} W$ is z-discordant for σ satisfying (2.3). It is then clear that, for n large enough, there must exist $\tau > 0$ for which the main result applies. This also gives a procedure to find such τ, by simply controlling the bounds appearing in the definition of z-discordance.

One simple example would be to allow noise on different edges to be non-identically distributed. For this class of examples, the spectral norm of W can be easily estimated with the tail bounds in [13]. A more challenging example is to address scenarios where noise on different edges is not independent. This notably appears in the orthogonal Procrustes problem, which attracted considerable attention lately [39, 47]. We now address a two-dimensional version of it.

Let $a \in \mathbb{C}^m$, $a \neq 0$, represent a cloud of points to be estimated in the (complex) plane. We are given n noisy observations of this cloud, of the form

$$b_{(i)} = a\bar{z}_i + \delta\varepsilon_{(i)}, \tag{5.1}$$

where $z_i \in \mathcal{S} = \{e^{i\theta} : \theta \in \mathbb{R}\}$ is an unknown rotation applied to the reference cloud, $\delta \geq 0$ is the noise level and $\varepsilon_{(i)} \in \mathbb{C}^m$ is a random noise vector. Arranging the observations as the columns of a matrix $B \in \mathbb{C}^{m\times n}$ and arranging the noise vectors as the columns of a random matrix $N \in \mathbb{C}^{m\times n}$, the observation model can be stated compactly as

$$B = az^* + \delta N. \tag{5.2}$$

The task is to estimate $a \in \mathbb{C}^m$ and $z \in \mathcal{S}^n$ from B. Of course, for any angle θ, the mapping $(a, x) \mapsto (ae^{i\theta}, ze^{i\theta})$ leaves B unchanged, so that the estimation can only be performed up to a global rotation.

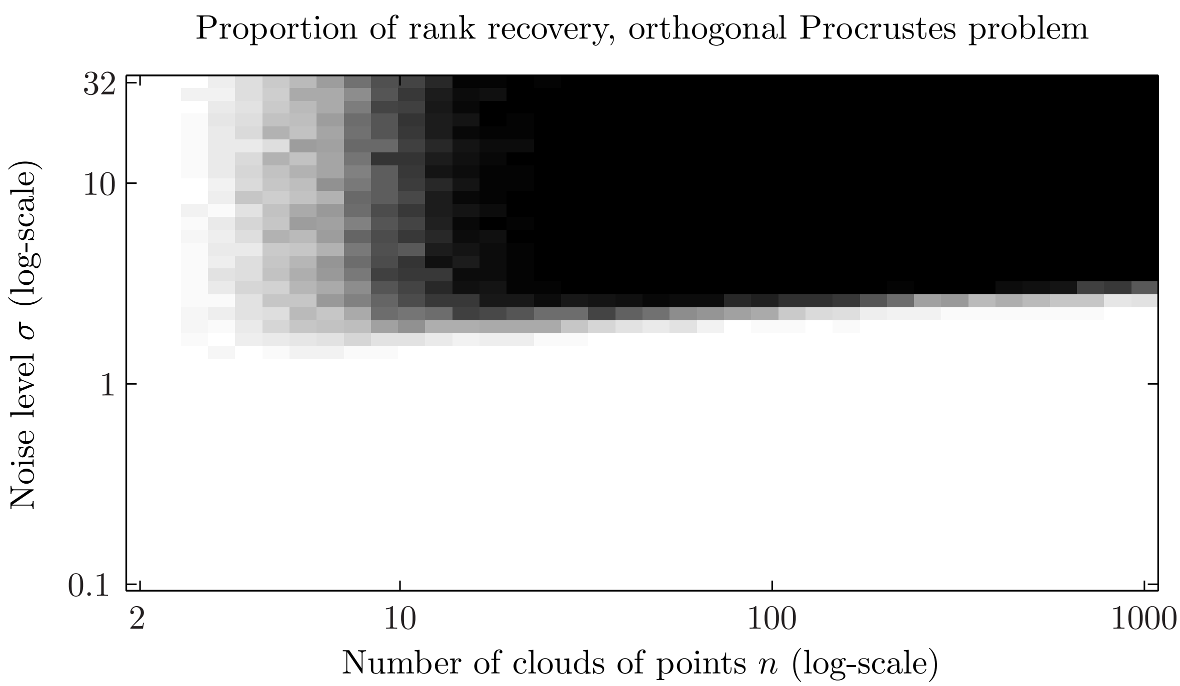

Figure 3: Tightness of the SDP relaxation for the orthogonal procrustes problem described in Section 5. The point cloud to estimate contains $m = 100$ points. For pairs of values (n, σ), where n is the number of observations of the cloud (each observation under an unknown, uniformly random in-plane rotation) and σ is the noise level, 100 instances of the problem are generated at random. Pixels code, on a gray scale, how frequently the semidefinite relaxation is tight (empirically).

Under the assumption that the entries of N are i.i.d. (complex) normal random variables, a maximum likelihood estimator for this estimation problem is a solution to the optimization problem

$$\min_{\hat{a}\in\mathbb{C}^m, \hat{z}\in\mathcal{S}^n} L(\hat{a}, \hat{z}) = \|\hat{a}\hat{z}^* - B\|_F^2. \tag{5.3}$$

(If there were no constraints on $\hat{z}$, the solution to this program would be given by the SVD of B. That is the PCA approximation to the orthogonal Procrustes problem.) There are no constraints on $\hat{a}$. Thus, any solution to (5.3) satisfies the first-order necessary optimality constraints,

$$\frac{\partial}{\partial \hat{a}} L(\hat{a}, \hat{z}) = 2(\hat{a}\hat{z}^* - B)\hat{z} = 0, \tag{5.4}$$

or, equivalently, $\hat{a} = \frac{1}{n}B\hat{z}$. Inject the latter into (5.3) to obtain an equivalent optimization problem:

$$\min_{\hat{z}\in\mathcal{S}^n} \left\| \frac{1}{n} B\hat{z}\hat{z}^* - B \right\|_F^2. \tag{5.5}$$

Expand the Frobenius norm and drop constant terms to find yet another equivalent formulation:

$$\max_{\hat{z}\in\mathcal{S}^n} \hat{z}^* B^* B \hat{z}. \tag{5.6}$$

Problem (QP) is recognizable, with cost matrix B^*B. We note in passing that, since B^*B is positive semidefinite, it is known that solving (SDP) followed by an appropriate projection (in case the solution is not of rank one) provides a guaranteed approximation to the solution of (QP), with ratio 0.785 [37]. In order to match the planted signal model, define:

$$C = \frac{1}{\|a\|^2} B^*B = zz^* + \frac{\delta}{\|a\|}\left(z\frac{a^*}{\|a\|}N + N^*\frac{a}{\|a\|}z^* \right) + \frac{\delta^2}{\|a\|^2}N^*N. \tag{5.7}$$

The rank-one signal component zz^* is duly apparent, corrupted by noise terms to first- and second-order with respect to a signal to noise ratio (SNR) $\|a\|/\delta$. The norm of a only appears through this ratio. Hence, without loss of generality, we may scale the problem such that $\|a\| = 1$. Then, solving the Procrustes problem is equivalent to solving (QP) with

$$C = zz^* + \tau W, \qquad \tau W = \delta\left(za^*N + N^*az^* \right) + \delta^2 N^*N. \tag{5.8}$$

Note how the noise matrix τW now depends on the signal z. Also, the noise originates at the nodes (the individual clouds) and creates perturbations on the edges, whereas for the synchronization problem we assumed the noise originated on the edges. Thus, noise on different edges here is neither identically distributed nor independent.

Forcing δ to be small enough is sufficient to guarantee (w.h.p for large n) that τW is z-discordant. Omitting details, we find that for $\delta = \mathcal{O}(n^{-1/2})$, the semidefinite relaxation approach to the Procrustes problem as described is tight, w.h.p. Even though this rate appears to be highly suboptimal when compared to experiments (Figure 3), to the best of our knowledge, this is the first result stating that the true MLE for a class of Procrustes problems can be computed in polynomial time, with a strictly positive noise level. Establishing such tightness for larger noise levels remains an open problem.

6 Conclusions

Computing maximum likelihood estimators is important for statistical problems, in particular when exact signal recovery is out of reach. In this paper, we have made partial progress toward understanding when this computation can be executed in polynomial time for the angular synchronization problem. It still remains to improve the proposed result to understand how this MLE can be computed for noise levels as high as those witnessed empirically.

For many other problems of interest, it has likewise been observed that semidefinite relaxations tend to be tight. Notably, this is the case for synchronization of rotations, a problem similar to angular synchronization, but with rotations in $\mathbb{R}^d$ rather than in $\mathbb{R}^2$. Analyzing the latter might prove more difficult, because rotations in $\mathbb{R}^d$ do not commute for $d > 2$. Tightness has also been observed for the Procrustes problem. The side result we proposed for the latter is much weaker than experiments suggest. Lastly, tightness of semidefinite relaxations has also been observed for the multi-reference alignment (MRA) problem [8]. The situation there differs markedly from the present setting, as the signal to recover is discrete and the noise originates at the vertices of a graph, not on the edges.

Notwithstanding the differences between the aforementioned problems and angular synchronization, we hope that the present work will help future attempts to explain the tightness phenomenon in various settings.

A Wigner matrices are discordant

This appendix is a proof for Proposition 3.3, namely, that for arbitrary $z \in \mathbb{C}^n$ such that $|z_1| = \cdots = |z_n| = 1$, complex Wigner matrices are z-discordant (Definition 3.1) with high probability.

A matrix W is z-discordant if and only if $\mathrm{diag}(z)^* W \,\mathrm{diag}(z)$ is $\mathbb{1}$-discordant. Since $\mathrm{diag}(z)^* W \,\mathrm{diag}(z)$ has the same distribution as W (owing to complex normal random variables having uniformly random phase), we may without loss of generality assume $z = \mathbb{1}$ in the remainder of the proof.

Before analyzing each of the properties, we start by noting that, by a suboptimal union bound argument, the maximum absolute value among k standard complex Gaussian random variables (not necessarily independent) is larger than t with probability at most $2ke^{-t^2/4}$. This will be useful multiple times.

1. $\Pr\left\{\|W\|_{\mathrm{e},\infty} > \log(n)\right\} \le n^{-\frac{1}{4}\log n+2}$.

   As there are $\frac{1}{2}n(n-1) \le \frac{1}{2}n^2$ nonzero independent entries in W, we can use the fact above and get

$$\Pr\left\{\|W\|_{\mathrm{e},\infty} > \log n\right\} \le n^2 e^{-\frac{\log^2 n}{4}} = n^{-\frac{1}{4}\log n+2}.$$

2. $\Pr\left\{\|W\|_{\mathrm{op}} > 3n^{1/2}\right\} \le e^{-n/2}$.

   Although tail bounds for the real version of this are well-known (see for example [52, 13]) and they mostly hold verbatim in the complex case, for the sake of completeness we include a classical argument, based on Slepian's comparison theorem and Gaussian concentration, for a tail bound in the complex valued case.

   We will bound the largest eigenvalue of W. It is clear that a simple union bound argument will allow us to bound also the smallest, and thus bound the largest in magnitude. Let $\lambda_+ = \max_{v\in\mathbb{C}^n:\|v\|=1} v^*Wv$ denote the largest eigenvalue of W. For any unit-norm $u, v \in \mathbb{C}^n$, the real valued Gaussian process $X_v = v^*Wv$ satisfies:

$$\begin{aligned}\mathbb{E}\left(X_v - X_u\right)^2 &= \mathbb{E}\left(\sum_{i<j} W_{ij}\left(\overline{v_i}v_j - \overline{u_i}u_j\right) + W_{ji}\left(\overline{v_j}v_i - \overline{u_j}u_i\right)\right)^2 \\ &= \sum_{i<j}\mathbb{E}\left[W_{ij}\left(\overline{v_i}v_j - \overline{u_i}u_j\right) + W_{ji}\left(\overline{v_j}v_i - \overline{u_j}u_i\right)\right]^2.\end{aligned}$$

   The variable W_{ij} has uniformly random phase, hence so does W_{ij}^2, so that $\mathbb{E}W_{ij}^2 = 0$. As a result,

$$\begin{aligned}\mathbb{E}\left(X_v - X_u\right)^2 &= \sum_{i<j} 2\mathbb{E}\left|W_{ij}\right|^2\left|\overline{v_i}v_j - \overline{u_i}u_j\right|^2 \\ &= 2\sum_{i<j}\left|\overline{v_i}v_j - \overline{u_i}u_j\right|^2 \\ &\le \sum_{i,j}\left|\overline{v_i}v_j - \overline{u_i}u_j\right|^2.\end{aligned}$$

Note that, since $\|u\| = \|v\| = 1$,

$$\begin{aligned}
\sum_{i,j} |\overline{v_i}v_j - \overline{u_i}u_j|^2 &= \sum_{ij} \left[|v_i|^2|v_j|^2 + |u_i|^2|u_j|^2 - \overline{v_i}v_j u_i \overline{u_j} - v_i \overline{v_j}\,\overline{u_i} u_j\right] \\
&= 2 - 2\,|v^* u|^2 \\
&\le 2\,(2 - 2\,|v^* u|) \\
&\le 4\,(1 - \Re\,[v^* u]) \\
&= 2\|v - u\|^2.
\end{aligned}$$

This means that we can use Slepian's comparison theorem (see for example [35, Cor. 3.12]) to get

$$\mathbb{E}\lambda_+ \le \sqrt{2}\mathbb{E} \max_{\tilde{v}\in\mathbb{R}^{2n}:\|\tilde{v}\|=1} \tilde{v}^T g \le 2\sqrt{n}, \tag{A.1}$$

where g is a standard Gaussian vector in $\mathbb{R}^{2n}$ and $\tilde{v}$ is a vector in $\mathbb{R}^{2n}$ obtained from $v \in \mathbb{C}^n$ by stacking its real and imaginary parts.

Since $|\|W_1\| - \|W_2\|| \le \|W_1 - W_2\| \le \|W_1 - W_2\|_{\mathrm{F}}$, Gaussian concentration [35] gives

$$\Pr\{\lambda_+ - \mathbb{E}\lambda_+ \ge t\} \le e^{-t^2/2}. \tag{A.2}$$

Using (A.1) and (A.2) gives

$$\Pr\left\{\|W\|_{\mathrm{op}} > 3n^{1/2}\right\} \le e^{-n/2}.$$

3. $\Pr\left\{\|W\mathbb{1}\|_\infty > n^{3/4}\right\} \le 2e^{-\frac{1}{4}n^{1/2}+\log n}$.

   The random vector given by $\frac{1}{(n-1)^{1/2}}W\mathbb{1}$ is jointly Gaussian where the marginal of each entry is a standard complex Gaussian. Hence,

$$\begin{aligned}
\Pr\left\{\|W\mathbb{1}\|_\infty > n^{3/4}\right\} &\le \Pr\left\{\left\|\frac{1}{(n-1)^{1/2}}W\mathbb{1}\right\|_\infty > n^{1/4}\right\} \\
&\le 2ne^{-\frac{1}{4}n^{1/2}} = 2e^{-\frac{1}{4}n^{1/2}+\log n}.
\end{aligned}$$

4. $\Pr\left\{|\mathbb{1}^* W\mathbb{1}| > n^{3/2}\right\} \le e^{-n/2}$.

   It is easy to see that $\mathbb{1}^* W\mathbb{1}$ is a real Gaussian random variable with zero mean and variance $2\frac{n(n-1)}{2} = n(n-1)$. This implies that:

$$\begin{aligned}
\Pr\left\{|\mathbb{1}^* W\mathbb{1}| > n^{3/2}\right\} &\le \Pr\left\{\frac{1}{(n(n-1))^{1/2}}|\mathbb{1}^* W\mathbb{1}| > n^{1/2}\right\} \\
&\le \frac{1}{\sqrt{2\pi}}\frac{1}{n^{1/2}}e^{-\frac{n}{2}} \\
&\le e^{-n/2}.
\end{aligned}$$

Acknowledgments

A. S. Bandeira was supported by AFOSR Grant No. FA9550-12-1-0317. N. Boumal was supported by a Belgian F.R.S.-FNRS fellowship while working at the Université catholique de Louvain (Belgium) and by a Research in Paris grant at INRIA. A. Singer was partially supported by Award Number R01GM090200 from the NIGMS, by Award Numbers FA9550-12-1-0317 and FA9550-13-1-0076 from AFOSR, and by Award Number LTR DTD 06-05-2012 from the Simons Foundation.

References

[1] E. Abbe, A. S. Bandeira, A. Bracher, and A. Singer. Decoding binary node labels from censored edge measurements: Phase transition and efficient recovery. *Transactions on Network Science and Engineering, to appear. Available online at arXiv:1404.4749 [cs.IT]*, 2014.

[2] E. Abbe, A.S. Bandeira, and G. Hall. Exact recovery in the stochastic block model. *arXiv preprint arXiv:1405.3267*, 2014.

[3] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.

[4] A. Agrawal, R. Raskar, and R. Chellappa. What is the range of surface reconstructions from a gradient field? In A. Leonardis, H. Bischof, and A. Pinz, editors, *Computer Vision – ECCV 2006*, volume 3951 of *Lecture Notes in Computer Science*, pages 578–591. Springer Berlin Heidelberg, 2006.

[5] B. Alexeev, A. S. Bandeira, M. Fickus, and D. G. Mixon. Phase retrieval with polarization. *SIAM J. on Imaging Sci.*, 7(1):35–66, 2013.

[6] F. Alizadeh, J.-P. Haeberly, and M.L. Overton. Complementarity and nondegeneracy in semidefinite programming. *Mathematical Programming*, 77(1):111–128, 1997.

[7] D. Amelunxen, M. Lotz, M. B. McCoy, and J.A. Tropp. Living on the edge: A geometric theory of phase transitions in convex optimization. *arXiv:1303.6672 [cs.IT]*, 2013.

[8] A. S. Bandeira, M. Charikar, A. Singer, and A. Zhu. Multireference alignment using semidefinite programming. *available online*, 2013.

[9] A. S. Bandeira, Y. Chen, and D. G. Mixon. Phase retrieval from power spectra of masked signals. *Information and Inference: a Journal of the IMA*, 3:83–102, 2014.

[10] A. S. Bandeira, C. Kennedy, and A. Singer. Approximating the little grothendieck problem over the orthogonal group. *Available online at arXiv:1308.5207 [cs.DS]*, 2013.

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

[12] A.S. Bandeira, Y. Khoo, and A. Singer. Open problem: Tightness of maximum likelihood semidefinite relaxations. In Maria Florina Balcan, Vitaly Feldman, and Csaba Szepesvri, editors, *Proceedings of the 27th Conference on Learning Theory*, volume 35 of *JMLR W&CP*, pages 1265–1267, 2014.

[13] A.S. Bandeira and R. van Handel. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *arXiv preprint arXiv:1408.6185*, 2014.

[14] A.I. Barvinok. Problems of distance geometry and convex properties of quadratic maps. *Discrete & Computational Geometry*, 13(1):189–202, 1995.

[15] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014.

[16] N. Boumal, A. Singer, P.-A. Absil, and V.D. Blondel. Cramér-Rao bounds for synchronization of rotations. *Information and Inference*, 3:1–39, 2014.

[17] J. Briët, F.M. de Oliveira Filho, and F. Vallentin. Grothendieck inequalities for semidefinite programs with rank constraint. *Theory of Computing*, 10(4):77–105, 2014.

[18] E.J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

[19] E.J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52:489–509, 2006.

[20] E.J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59:1207–1223, 2006.

[21] E.J. Candès, T. Strohmer, and V. Voroninski. Phaselift: exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 2011.

[22] V. Chandrasekaran, P.A. Parrilo, and A.S. Willsky. Latent variable graphical model selection via convex optimization. *The Annals of Statistics*, 40(4):1935–1967, 08 2012.

[23] V. Chandrasekaran, B. Recht, P.A. Parrilo, and A.S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.

[24] L. Demanet and P. Hand. Stable optimizationless recovery from phaseless linear measurements. *arXiv:1208.1803*, 2012.

[25] L. Demanet and V. Jugnon. Convex recovery from interferometric measurements. *arXiv preprint arXiv:1307.6864*, 2013.

[26] X. Ding and T. Jiang. Spectral distribution of adjacency and laplacian matrices of random graphs. *The annals of applied probability*, 2010.

[27] D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52:1289–1306, 2006.

[28] 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.

[29] M.X. Goemans and D.P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.

[30] M.X. Goemans and D.P. Williamson. Approximation algorithms for Max-3-Cut and other problems via complex semidefinite programming. *Journal of Computer and System Sciences*, 68(2):442–470, 2004.

[31] A. Grothendieck. Resume de la theorie metrique des produits tensoriels topologiques (french). *Reprint of Bol. Soc. Mat. Sao Paulo*, page 179, 1996.

[32] R. Hartley, J. Trumpf, Y. Dai, and H. Li. Rotation averaging. *International Journal of Computer Vision*, 103(3):267–305, 2013.

[33] Q.X. Huang and L. Guibas. Consistent shape maps via semidefinite programming. In *Computer Graphics Forum*, volume 32, pages 177–186. Wiley Online Library, 2013.

[34] M. Journée, F. Bach, P.-A. Absil, and R. Sepulchre. Low-rank optimization on the cone of positive semidefinite matrices. *SIAM Journal on Optimization*, 20(5):2327–2351, 2010.

[35] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: isoperimetry and processes*, volume 23 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge / A Series of Modern Surveys in Mathematics*. 1991.

[36] S.H. Low. Convex relaxation of optimal power flow: a tutorial. In *Bulk Power System Dynamics and Control-IX Optimization, Security and Control of the Emerging Power Grid (IREP), 2013 IREP Symposium*, pages 1–15. IEEE, 2013.

[37] Z. Luo, W. Ma, AMC So, Y. Ye, and S. Zhang. Semidefinite relaxation of quadratic optimization problems. *Signal Processing Magazine, IEEE*, 27(3):20–34, 2010.

[38] D. Martinec and T. Pajdla. Robust rotation and translation estimation in multiview reconstruction. In *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, pages 1–8, June 2007.

[39] A. Nemirovski. Sums of random symmetric matrices and quadratic optimization under orthogonality constraints. *Math. Program.*, 109(2-3):283–317, 2007.

[40] G. Pataki. On the rank of extreme matrices in semidefinite programs and the multiplicity of optimal eigenvalues. *Mathematics of operations research*, 23(2):339–358, 1998.

[41] G. Pisier. Grothendieck's theorem, past and present. *Bull. Amer. Math. Soc.*, 49:237323, 2011.

[42] J. Rubinstein and G. Wolansky. Reconstruction of optical surfaces from ray data. *Optical Review*, 8(4):281–283, 2001.

[43] A.P. Ruszczyński. *Nonlinear optimization*. Princeton University Press, 2006.

[44] G. Sagnol. A class of semidefinite programs with rank-one solutions. *Linear Algebra and its Applications*, 435(6):1446–1463, 2011.

[45] A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1):20–36, 2011.

[46] A. Singer and Y. Shkolnisky. Three-dimensional structure determination from common lines in Cryo-EM by eigenvectors and semidefinite programming. *SIAM J. Imaging Sciences*, 4(2):543–572, 2011.

[47] A. M.-C. So. Moment inequalities for sums of random matrices and their applications in optimization. *Mathematical Programming*, 130(1):125–151, 2011.

[48] A. M.-C. So, J. Zhang, and Y. Ye. On approximating complex quadratic optimization problems via semidefinite programming relaxations. *Mathematical Programming*, 110(1):93–110, 2007.

[49] S. Sojoudi and J. Lavaei. Exactness of semidefinite relaxations for nonlinear optimization problems with underlying graph structure. *SIAM Journal on Optimization*, 24(4):1746–1778, 2014.

[50] J.A. Tropp. Just relax: convex programming methods for identifying sparse signals in noise. *Information Theory, IEEE Transactions on*, 52(3):1030–1051, March 2006.

[51] L. Vanderberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38:49–95, 1996.

[52] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. *Chapter 5 of: Compressed Sensing, Theory and Applications. Edited by Y. Eldar and G. Kutyniok. Cambridge University Press*, 2012.

[53] L. Wang and A. Singer. Exact and stable recovery of rotations for robust synchronization. *Information and Inference*, 2(2):145–193, 2013.

[54] D.P. Williamson and D.B. Shmoys. *The design of approximation algorithms*. Cambridge University Press, 2011.

[55] S. Zhang and Y. Huang. Complex quadratic optimization and semidefinite programming. *SIAM Journal on Optimization*, 16(3):871–890, 2006.