

A CHEEGER INEQUALITY FOR THE GRAPH CONNECTION LAPLACIAN

AFONSO S. BANDEIRA*, AMIT SINGER†, AND DANIEL A. SPIELMAN‡

Abstract. The $O(d)$ Synchronization problem consists of estimating a set of n unknown orthogonal $d \times d$ matrices O_1, \dots, O_n from noisy measurements of a subset of the pairwise ratios $O_i O_j^{-1}$. We formulate and prove a Cheeger-type inequality that relates a measure of how well it is possible to solve the $O(d)$ synchronization problem with the spectra of an operator, the graph Connection Laplacian. We also show how this inequality provides a worst case performance guarantee for a spectral method to solve this problem.

Keywords: $O(d)$ Synchronization, Cheeger Inequality, Graph Connection Laplacian, Vector Diffusion Maps.

1. Introduction. While the graph Laplacian is used to encode similarities between connected vertices, the graph Connection Laplacian endows the edges with transformations that describe the nature of the similarity. For example, consider a collection of two-dimensional photos of a three-dimensional object that are taken from many angles as it spins in mid-air. We could form a graph with one vertex for each photo by connecting photos that are similar. By applying simple transformations, such as in-plane rotations, we will discover similarities between photos that are not apparent when we merely treat them as vectors of pixels. In this case, we may also wish to keep track of the transformation under which a pair of photos is similar. Singer and Wu [27] defined the Connection Laplacian to encode this additional information. The problem of using this information to assign a viewpoint to each picture is an instance of a synchronization problem on the graph.

More formally, the input to a synchronization problem over a group \mathcal{G} is an undirected graph $G = (V, E)$ and a group element $\rho_{ij} \in \mathcal{G}$ for each edge $(i, j) \in E$, such that $\rho_{ji} = \rho_{ij}^{-1}$. We say that an assignment of group elements to vertices, also called a *group potential*, $g : V \rightarrow \mathcal{G}$, satisfies an edge ρ_{ij} if $g_i = \rho_{ij} g_j$. The objective in a synchronization problem is to find a group potential that satisfies the edges as much as possible. When there is a g that satisfies all of the edges, it is easy to find: one can arbitrarily fix the value of g at one vertex and then iteratively set g at neighbors of vertices whose values have already been set. When there is no group potential that satisfies all of the edges, we must specify a measure of how well a group potential satisfies the edges. This is achieved by the notion of frustration which will be defined in Section 1.3.

We focus on the group $O(d)$ of $d \times d$ orthogonal matrices (rotations, reflections and compositions of both on \mathbb{R}^d). Please keep in mind that the “ O ” has nothing to do with asymptotic notation. This group is of particular interest in several applications. For example, when $d = 1$, i.e. $\mathcal{G} = O(1) \cong \mathbb{Z}/2\mathbb{Z}$, the solution to the synchronization problem can be used to determine whether a manifold is orientable [28]. The $\mathcal{G} = O(1)$ case is also a generalization of the MAX-CUT problem [13]: the group potential defines a partition of the vertices into two parts, and the group elements on edges specify whether the vertices they connect should be on the same or opposite sides of the partition. In fact, our inequality for *partial synchronization* can be understood as a generalization of Trevisan’s inequality relating eigenvalues of the graph Laplacian and the maximum cut [32].

Synchronization over $O(d)$ plays a major role in an algorithm for the sensor network localization problem [10]. The similar problem of synchronization over $SO(d)$, the group of rotations in \mathbb{R}^d , also has several applications. The problem over $SO(3)$ can be used for global alignment of 3-D scans [33], and the problem over $SO(2)$ plays a major role in new algorithms for the cryo-electron microscopy problem (see [26, 29]). Other applications of $SO(2)$ synchronization may be found in [25, 35, 17].

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

†PACM and Department of Mathematics, Princeton University, Princeton, NJ 08544, USA (amits@math.princeton.edu).

‡Department of Applied Mathematics and Department of Computer Science, Yale University, PO Box 208285, New Haven, CT 06520-8285, USA (spielman@cs.yale.edu).

Singer [25] proposed solving the $SO(2)$ synchronization problem by constructing a matrix, the Connection Laplacian, whose eigenvectors associated with the smallest eigenvalues would provide the group potential if it were possible to satisfy all the edges. He then showed that under a model of random noise that prevents such a solution, a good solution can still be obtained by rounding the smallest eigenvectors. A similar algorithm was proposed in [11, 26] for $SO(3)$. These algorithms are analogous to spectral graph partitioning—the use of the smallest eigenvector of the Laplacian matrix to partition a graph [15]. The analysis of the $SO(2)$ algorithm under random noise can be viewed as an analog of McSherry’s [21] analysis of spectral partitioning.

Our paper contains three theorems. We begin by considering the simpler problem of finding an assignment of unit vectors to each vertex that agrees with the transformations on the edges. Motivated by Trevisan [32], we first consider a variant of the problem in which we are allowed to find a *partial* assignment in which we assign the zero vector to some of the vertices. In Theorem 2.2, we prove a quadratic relationship between the smallest eigenvalue of the connection Laplacian and the minimum frustration of a partial assignment. To analyze the case in which we must assign a unit vector to every vertex, we observe that partial assignments are really only required when the underlying graph has poor connectivity. In Theorem 2.4, we prove a relation between the minimum frustration of a full assignment and the smallest eigenvalue of the Connection Laplacian *and* the second-smallest eigenvalue of the underlying graph Laplacian. We see that when the second-smallest eigenvalue of the underlying graph Laplacian is large (i.e., when it is a good expander), the minimum frustration is well approximated by the smallest eigenvalue of the Connection Laplacian. Our main result is Theorem 2.6, in which we relate the minimum frustration of a group potential to the sum of the smallest d eigenvalues of the Connection Laplacian *and* the second-smallest eigenvalue of the underlying graph Laplacian.

In the same way that the classical Cheeger’s inequality provides a worst-case performance guarantee for spectral clustering, the three Theorems described above provide worst-case performance guarantees for a spectral method to solve each of the described Synchronization problems. It is worth noting that, in a different setup, Trevisan [32] showed a particular case of this inequality, when the group is $O(1) \cong \mathbb{Z}/2\mathbb{Z}$ and all the offsets are -1 . In that case the problem is equivalent to the MAX-CUT problem. Therefore, this $O(1)$ inequality gives a performance guarantee for a spectral method to solve MAX-CUT [32].

It is worth mentioning that there are several other adaptations of Cheeger’s inequality. Recent progress in multi-way partitioning problem gives a Cheeger inequality for the partitioning problem where one wants to partition the graph into more than 2 subsets (see [12, 19]). There is also a generalization of the Cheeger inequality for simplicial complexes, instead of graphs (see, [23, 24, 30]).

The rest of this section includes both mathematical preliminaries that are needed in later sections and the formulation of the problem. Section 2 consists of our main contributions, algorithms to solve different forms of the Synchronization problem and Cheeger-type inequalities that provide guarantees for these methods, as well as a brief overview of the proofs. In Section 3 we provide the rigorous proofs for the core results described in Section 2. We discuss an ℓ_1 version of the Synchronization problem in Section 4 and show a few tightness results in Section 5. We end with some concluding remarks and a few open problems in Section 6.

1.1. Notation. Throughout the paper we use the notation $[n]$ to refer to $\{1, \dots, n\}$. Also, we make use of several matrix and vector notations. Given a matrix A we denote by $\|A\|_F$ its Frobenius norm. If A is symmetric we denote by $\lambda_1(A), \lambda_2(A), \dots$ its eigenvalues in increasing order. Assuming further that A is positive semi-definite we define the A -inner product of vectors x and y as $\langle x, y \rangle_A = x^T A y$ (and say that two vectors are A -orthogonal if this inner product is zero). Also, we define the A -norm of x as $\|x\|_A = \sqrt{\langle x, x \rangle_A}$ and we represent the ℓ_2 norm of x by $\|x\|$. Given $x \in (\mathbb{R}^d)^n$ we denote by x_i the i -th $d \times 1$ block of x (that will correspond to the value of x on the vertex i) and, for any $u > 0$ we define x_i^u as

$$x_i^u = \begin{cases} \frac{x_i}{\|x_i\|} & \text{if } \|x_i\|^2 \geq u, \\ 0 & \text{if } \|x_i\|^2 < u. \end{cases} \quad (1.1)$$

Furthermore, for $u = 0$ we denote x^0 by \tilde{x} , that is,

$$\tilde{x}_i = \begin{cases} \frac{x_i}{\|x_i\|} & \text{if } x_i \neq 0, \\ 0 & \text{if } x_i = 0. \end{cases} \quad (1.2)$$

Finally, \mathbb{S}^{d-1} denotes the unit sphere in \mathbb{R}^d .

1.2. Cheeger's Inequality and the Graph Laplacian. Before considering the synchronization problem we will briefly present the classical graph Cheeger's inequality in the context of spectral partitioning. The material presented in this Section is well known but it will help motivate the ideas that follow in the next Sections.

Let $G = (V, E)$ be an undirected weighted graph with n vertices. In this section we discuss the problem of partitioning the vertices in two similarly sized sets in a way that minimizes the cut: the volume of edges across the subsets (of the partition).

There are several ways to measure the performance of a particular partition of the graph, we will consider the one known as the Cheeger constant. Given a partition (S, S^c) of V let $h_S := \frac{\text{cut}(S)}{\min\{\text{vol}(S), \text{vol}(S^c)\}}$, where the value of the cut associated with S is $\text{cut}(S) = \sum_{i \in S} \sum_{j \in S^c} w_{ij}$, its volume is $\text{vol}(S) = \sum_{i \in S} d_i$, and $d_i = \sum_{j \in V} w_{ij}$ is the weighted degree of vertex i . We want to partition the graph so that h_S is minimized, and the minimum value is referred to as the Cheeger number of the graph, denoted $h_G = \min_{S \subset V} h_S$. Finding the optimal S is known to be NP-hard, as it seems to require searching over an exponential number of possible partitions.

There is another way to measure the performance of a partition (S, S^c) known as the normalized cut:

$$\text{Ncut}(S) = \text{cut}(S) \left(\frac{1}{\text{vol}(S)} + \frac{1}{\text{vol}(S^c)} \right).$$

As before, we want to find a subset with as small of an Ncut as possible. Note that the normalized cut and the Cheeger constant are closely related:

$$\frac{1}{2} \text{Ncut}(S) \leq h_S \leq \text{Ncut}(S).$$

Let us introduce a few important definitions. Let W_0 be the weighted adjacency matrix of G and D_0 the degree matrix, a diagonal matrix with elements d_i . If we consider a vector $f \in \mathbb{R}^n$ whose entries take only 2 possible values, one associated with vertices in S and another in S^c , then the quadratic form $Q_f = \frac{1}{2} \sum_{ij} w_{ij} (f_i - f_j)^2$ is of fundamental importance as a measure of the cut between the sets. The symmetric positive semi-definite matrix that corresponds to this quadratic form, L_0 , is known as the graph Laplacian of G . It is defined as $L_0 = D_0 - W_0$ and satisfies $v^T L_0 v = Q_v$ for any $v \in \mathbb{R}^n$. It is also useful to consider the normalized graph Laplacian $\mathcal{L}_0 = D_0^{-1/2} L_0 D_0^{-1/2} = I - D_0^{-1/2} W_0 D_0^{-1/2}$, which is also a symmetric positive semi-definite matrix.

Let us represent a partition (S, S^c) by a cut-function $f_S : V \rightarrow \mathbb{R}$ given by

$$f_S(i) = \begin{cases} \sqrt{\frac{\text{vol}(S^c)}{\text{vol}(S) \text{vol}(G)}} & \text{if } i \in S, \\ -\sqrt{\frac{\text{vol}(S^c)}{\text{vol}(S) \text{vol}(G)}} & \text{if } i \in S^c. \end{cases}$$

It is straightforward to show that $Q_{f_S} = f_S^T L_0 f_S = \text{Ncut}(S)$, $f_S^T D_0 f_S = 1$, and $f_S^T D_0 \mathbf{1} = 0$, where $\mathbf{1}$ is the all-ones vector in \mathbb{R}^n . This is the motivation for a spectral method to approximate the minimum normalized cut problem. If we drop the constraint that f needs to be a cut-function and simply enforce the properties established above then one would formulate the following relaxed problem

$$\min_{f: V \rightarrow \mathbb{R}, f^T D_0 f = 1, f^T D_0 \mathbf{1} = 0} f^T L_0 f. \quad (1.3)$$

Since $\mathbf{1}^T L_0 \mathbf{1} = 0$, we know by the Courant-Fisher formula that (1.3) corresponds to an eigenvector problem whose minimum is $\lambda_2(\mathcal{L}_0)$ and whose minimizer can be obtained by the corresponding eigenvector.

Since problem (1.3) is a relaxation of the minimum Ncut problem we automatically have $\frac{1}{2}\lambda_2(\mathcal{L}_0) \leq \frac{1}{2} \min_{S \subset V} \text{Ncut} \leq h_G$. Remarkably one can show that the relaxation is not far from the partitioning problem. In fact, one can round the solution of (1.3) so that it corresponds to a partition (S, S^c) of G , whose h_S we can control. This is made precise by the following classical result in spectral graph theory (several different proofs for this inequality can be found in [8]):

THEOREM 1.1 (Cheeger Inequality [2, 3]). *Let $G = (V, E)$ be a graph and \mathcal{L}_0 its normalized graph Laplacian. Then*

$$\frac{1}{2}\lambda_2(\mathcal{L}_0) \leq h_G \leq \sqrt{2\lambda_2(\mathcal{L}_0)},$$

where h_G is the Cheeger constant of G . Furthermore, the bound is constructive: using the solution of the eigenvector problem one can produce partition (S, S^c) that achieves the upper bound $\sqrt{2\lambda_2(\mathcal{L}_0)}$.

An alternative way to interpret Theorem 1.1 is through random walks on graphs. Note that the matrix $D_0^{-1}W_0$ is the transition probability matrix of a random walk in G , whose transition probabilities are proportional to the edge weights. It is known that the eigenvalues of \mathcal{L}_0 encode important information about the random walk. In fact, the second smallest eigenvalue¹ is a good measure of how well the random walk mixes. More specifically, the smaller $\lambda_2(\mathcal{L}_0)$, the slower the convergence to the limiting stationary distribution. It is clear that clusters will constitute obstacles to rapid mixing of the random walk, since the probability mass might be trapped inside such a set for a while. Cheeger's inequality (Theorem 1.1) shows that, in some sense, these sets are the only obstacles to rapid mixing.

1.3. Frustration, Vector-Valued Walks and the Connection Laplacian. If, in addition to a graph, we are given an orthogonal transformation $\rho_{ij} \in O(d)$ for each edge $(i, j) \in E$, we can consider a random walk that takes the transformations into account. One way of doing this is by defining a random walk that, instead of moving point masses, moves a vector from vertex to vertex and transforms it via the orthogonal transformation associated with the edge. One can similarly define a random walk that moves group elements on vertices. The Connection Laplacian was defined by Singer and Wu [27] to measure the convergence of such random walks. The construction requires that $\rho_{ji} = \rho_{ij}^{-1} = \rho_{ij}^T$. Define the symmetric matrix $W_1 \in \mathbb{R}^{dn \times dn}$ so that the (i, j) -th $d \times d$ block is given by $(W_1)_{ij} = w_{ij}\rho_{ij}$, where w_{ij} is the weight of the edge (i, j) . Also, let $D_1 \in \mathbb{R}^{dn \times dn}$, be the diagonal matrix such that $(D_1)_{ii} = d_i I_{d \times d}$. We assume $d_i > 0$, for every i . The graph Connection Laplacian L_1 is defined to be $L_1 = D_1 - W_1$, and the normalized graph Connection Laplacian is

$$\mathcal{L}_1 = I - D_1^{-1/2}W_1D_1^{-1/2}.$$

If $v : V \rightarrow \mathbb{S}^{d-1}$ assigns a unit vector in \mathbb{R}^d to each vertex, we may think of v as a vector in dn dimensions. In this case the quadratic form

$$v^T L_1 v = \sum_{(i,j) \in E} w_{ij} \|v_i - \rho_{ij} v_j\|^2 = \frac{1}{2} \sum_{i,j} w_{ij} \|v_i - \rho_{ij} v_j\|^2$$

is a measure of how well v satisfies the edges. This will be zero if $v_i = \rho_{ij} v_j$ for all edges (i, j) . As $w_{ij} = 0$ when $(i, j) \notin E$, we can sum over all pairs of vertices without loss of generality. An assignment satisfying all edges will correspond to a stationary distribution in the vector-valued random walk.

Following our analogy with Cheeger's inequality for the normalized graph Laplacian, we normalize this

¹We note that the smallest eigenvalue is always 0.

measure by defining the *frustration* of v as

$$\eta(v) = \frac{v^T L_1 v}{v^T D_1 v} = \frac{1}{2} \frac{\sum_{i,j} w_{ij} \|v_i - \rho_{ij} v_j\|^2}{\sum_i d_i \|v_i\|^2}. \quad (1.4)$$

We then define the \mathbb{S}^{d-1} frustration constant of G as

$$\eta_G = \min_{v: V \rightarrow \mathbb{S}^{d-1}} \eta(v). \quad (1.5)$$

The smallest eigenvalue of \mathcal{L}_1 provides a relaxation of η_G , as

$$\lambda_1(\mathcal{L}_1) = \min_{z \in \mathbb{R}^{dn}} \frac{z^T \mathcal{L}_1 z}{z^T z} = \min_{x \in \mathbb{R}^{dn}} \frac{(D_1^{\frac{1}{2}} x)^T \mathcal{L}_1 (D_1^{\frac{1}{2}} x)}{(D_1^{\frac{1}{2}} x)^T (D_1^{\frac{1}{2}} x)} = \min_{x \in \mathbb{R}^{dn}} \frac{x^T L_1 x}{x^T D_1 x} = \min_{x: V \rightarrow \mathbb{R}^d} \eta(x).$$

If there is a group potential $g : V \rightarrow O(d)$ that satisfies all the edges (which would again correspond to a stationary distribution for the $O(d)$ -valued random walk), then we can obtain d orthogonal vectors on which the quadratic form defined by \mathcal{L}_1 is zero. For each $1 \leq k \leq d$ we obtain one of these vectors by setting $v(i)$ to the k th column of $g(i)$ for all $i \in V$. In particular, this means that the columns of the matrices of the group potential that satisfies all of the edges lie in the nullspace of \mathcal{L}_1 . Since $g(i) \in O(d)$ these vectors are orthogonal. If G is connected, one can show that these are the only vectors in the nullspace of \mathcal{L}_1 . This observation is the motivation for the use of a spectral algorithm for synchronization.

We define the frustration of a group potential $g : V \rightarrow O(d)$ to be

$$\nu(g) = \frac{1}{2d} \frac{1}{\text{vol}(G)} \sum_{i,j} w_{ij} \|g_i - \rho_{ij} g_j\|_F^2. \quad (1.6)$$

We then define the $O(d)$ frustration constant of G to be

$$\nu_G = \min_{g: V \rightarrow O(d)} \nu(g).$$

In Theorem 2.6, we prove that this frustration constant is small if and only if the sum of the first d eigenvalues of \mathcal{L}_1 is small as well.

2. Cheeger's type inequalities for the synchronization problem. In this Section we present our main results. We present three spectral algorithms to solve three different formulations of synchronization problems and obtain for each a guarantee of performance in the form of a Cheeger's type inequality. We will briefly summarize both the results and the ideas to obtain them, leaving the rigorous proofs to Section 3.

We start by considering the \mathbb{S}^{d-1} synchronization problem. This corresponds to finding, for each vertex i of the graph, a vector $v_i \in \mathbb{S}^{d-1}$ in way that for each edge (i, j) the vectors agree with the edges, meaning $v_i = \rho_{ij} v_j$. Since this might not always be possible we look for a function $v : V \rightarrow \mathbb{S}^{d-1}$ for which the frustration $\eta(v)$ is minimum (see (1.4)). Motivated by an algorithm to solve MAX-CUT by Trevisan [32], we first consider a version of the problem for which we allows ourselves to synchronize only a subset of the vertices, corresponding to the partial synchronization in \mathbb{S}^{d-1} . We then move on to consider the full synchronization problem in \mathbb{S}^{d-1} .

Finally we will present our main result, an algorithm for $O(d)$ synchronization and a Cheeger-like inequality that equips it with a worst-case guarantee. Recall that the $O(d)$ synchronization corresponds to finding an assignment of an element $g_i \in O(d)$ to each vertex i in a way that minimizes the discrepancy with the pairwise measurements $\rho_{ij} \sim g_i g_j^{-1}$ obtained for each edge. This corresponds to minimizing the $O(d)$ frustration, $\nu(g)$, (see (1.6)).

2.1. Partial synchronization in \mathbb{S}^{d-1} . The motivation for considering a spectral relaxation for the synchronization problem in \mathbb{S}^{d-1} is the observation that $\lambda_1(\mathcal{L}_1) = \min_{x:V \rightarrow \mathbb{R}^d} \eta(x)$. In order to understand how tight the relaxation is we need to relate $\lambda_1(\mathcal{L}_1)$ with $\eta_G = \min_{x:V \rightarrow \mathbb{S}^{d-1}} \eta(x)$.

Consider, however, the following example: a graph consisting of two disjoint components, one whose ρ_{ij} measurements are perfectly compatible and another one on which they are not. Its graph Connection Laplacian would have a non-zero vector in its null space, corresponding to synchronizations on the compatible component and zero on the incompatible part (thus $\lambda_1(\mathcal{L}_1) = 0$). On the other hand, the constraint that v has to take values on \mathbb{S}^{d-1} , will force it to try to synchronize the incompatible part thereby bounding η_G away from zero. This example motivates a different formulation of the \mathbb{S}^{d-1} synchronization problem where vertices are allowed not to be labeled (labeled with 0). We thus define the partial \mathbb{S}^{d-1} frustration constant of G , as the minimum possible frustration value for such an assignment,

$$\eta_G^* = \min_{v:V \rightarrow \mathbb{S}^{d-1} \cup \{0\}} \eta(v). \quad (2.1)$$

We propose the following algorithm to solve the partial \mathbb{S}^{d-1} synchronization problem.

ALGORITHM 2.1. *Given a graph $G = (V, E)$ and a function $\rho : E \rightarrow O(d)$, construct the normalized Connection Laplacian \mathcal{L}_1 and the degree matrix D_1 . Compute z , the eigenvector corresponding to the smallest eigenvalue of \mathcal{L}_1 . Let $x = D_1^{-\frac{1}{2}} z$. For each vertex index i , let $u_i = \|x_i\|$, and set $v^i : V \rightarrow \mathbb{S}^{d-1} \cup \{0\}$ as $v^i = x^{u_i}$, according to (1.1). Output v equal to the v^i that minimizes $\eta(v^i)$.*

Lemma 3.1 guarantees that the solution v given by Algorithm 2.1 satisfies $\eta(v) \leq \sqrt{10\eta(x)}$. Since x was computed so that $\eta(x) = \lambda_1(\mathcal{L}_1)$, Algorithm 2.1 is guaranteed to output a solution v such that

$$\eta(v) \leq \sqrt{10\lambda_1(\mathcal{L}_1)}.$$

Note that $\lambda_1(\mathcal{L}_1) \leq \eta_G^*$, which is the optimum value for the partial \mathbb{S}^{d-1} synchronization problem (see (2.1)). The proof for Lemma 3.1 will appear below. The idea to show that the rounding, from x , to the solution v done by Algorithm 2.1 produces a solution with $\eta(v) \leq \sqrt{10\eta(x)}$ is to use the probabilistic method. One considers a random rounding scheme by rounding x as in Algorithm 2.1 and (1.1) but thresholding at a random value u , drawn from a well-chosen distribution. One then shows that, in expectation, the frustration of the rounded vector is bounded by $\sqrt{10\eta(x)}$. This automatically ensures that there must exist a value u that produces a solution with frustration bounded by $\sqrt{10\eta(x)}$. The rounding described in Algorithm 2.1 runs through all possible such roundings and is thus guaranteed to produce a solution satisfying the bound. An $O(1)$ version of this algorithm and analysis appeared in [32], when ρ is the constant function equal to -1 , in the context of the MAX-CUT problem. In fact, if $d = 1$ the factor 10 can be substituted by 8 and the stronger inequality holds $\eta(v) \leq \sqrt{8\lambda_1(\mathcal{L}_1)}$.

The above performance guarantee for Algorithm 2.1 automatically implies the following Cheeger-like inequality.

THEOREM 2.2. *Let $G = (V, E)$ be a weighted graph. Given a function $\rho : E \rightarrow O(d)$, let η_G^* be the partial \mathbb{S}^{d-1} frustration constant of G and $\lambda_1(\mathcal{L}_1)$ the smallest eigenvalue of the normalized graph Connection Laplacian. Then*

$$\lambda_1(\mathcal{L}_1) \leq \eta_G^* \leq \sqrt{10\lambda_1(\mathcal{L}_1)}. \quad (2.2)$$

Furthermore, if $d = 1$, the stronger inequality holds, $\eta_G^ \leq \sqrt{8\lambda_1(\mathcal{L}_1)}$.*

We note that Trevisan [32], in the context of MAX-CUT, iteratively performs this partial synchronization procedure in the subgraph composed of the vertices left unlabeled by the previous iteration, in order to label the entire graph. We, however, consider only one iteration.

2.2. Full synchronization in \mathbb{S}^{d-1} . In this section we adapt Algorithm 2.1 to solve (full) synchronization in \mathbb{S}^{d-1} and show performance guarantees, under reasonable conditions, by obtaining bounds for η_G , the frustration constant for synchronization in \mathbb{S}^{d-1} . The intuition given to justify the relaxation to partial \mathbb{S}^{d-1} synchronization was based on the possibility of poor connectivity of the graph (small spectral gap). In this section we show that poor connectivity, as measured by a small spectral gap in the normalized graph Laplacian, is the only condition under which one can have large discrepancy between the frustration constants and the spectra of the graph Connection Laplacian. We will show that, as long as the spectral gap is bounded away from zero, one can in fact control the full frustration constants.

ALGORITHM 2.3. *Given a weighted graph $G = (V, E)$ and a function $\rho : E \rightarrow O(d)$, construct the normalized Connection Laplacian \mathcal{L}_1 and the degree matrix D_1 . Compute z , the eigenvector corresponding to the smallest eigenvalue of \mathcal{L}_1 . Let $x = D_1^{-\frac{1}{2}}z$. Output the solution $v : V \rightarrow \mathbb{S}^{d-1} \cup \{0\}$ where each v_i is defined as*

$$v_i = \frac{x_i}{\|x_i\|}.$$

If $x_i = 0$, have v_i be any vector in \mathbb{S}^{d-1} .

Similarly to Algorithm 2.1, Lemma 3.6 guarantees that the solution v given by Algorithm 2.3 satisfies $\eta(v) \leq 44 \frac{1}{\lambda_2(\mathcal{L}_0)} \eta(x)$. Again, since x was computed so that $\eta(x) = \lambda_1(\mathcal{L}_1)$, then Algorithm 2.1 is guaranteed to output a solution v such that

$$\eta(v) \leq 44 \frac{\lambda_1(\mathcal{L}_1)}{\lambda_2(\mathcal{L}_0)}.$$

Recall that, trivially, $\lambda_1(\mathcal{L}_1) \leq \eta_G$, which is the optimum value for the (full) \mathbb{S}^{d-1} synchronization problem (see (1.5)). The proof for Lemma 3.6 is also deferred until Section 3. The idea here is to look at the vector of the local norms of x : $n_x \in \mathbb{R}^n$ where $n_x(i) = \|x_i\|$. It is not hard to show that $\frac{n_x^T \mathcal{L}_0 n_x}{n_x^T D_0 n_x} \leq \eta(x)$, which means that, if $\eta(x)$ is small then n_x cannot vary much between two vertices that share an edge. Since $\lambda_2(\mathcal{L}_0)$ is large one can show that such a vector needs to be close to constant, which means that the norms of x across the vertices are similar. If the norms were all the same then the rounding $v_i = \frac{x_i}{\|v_i\|}$ would not affect the value of $\eta(\cdot)$, we take this slightly further by showing that if the norms are similar then we can control how much the rounding affects the penalty function.

The above performance guarantee for Algorithm 2.3 automatically implies another Cheeger-like inequality.

THEOREM 2.4. *Let $G = (V, E)$ be a graph. Given a function $\rho : E \rightarrow O(d)$, let η_G be the \mathbb{S}^{d-1} frustration constants of G , $\lambda_1(\mathcal{L}_1)$ the smallest eigenvalue of the normalized graph Connection Laplacian and $\lambda_2(\mathcal{L}_0)$ the second smallest eigenvalue of the normalized graph Laplacian. Then,*

$$\lambda_1(\mathcal{L}_1) \leq \eta_G \leq 44 \frac{\lambda_1(\mathcal{L}_1)}{\lambda_2(\mathcal{L}_0)}.$$

2.3. The $O(d)$ synchronization problem. We present now our main contribution, a spectral algorithm for $O(d)$ synchronization together with a Cheeger-type inequality that provides a worst-case performance guarantee for the algorithm.

Before presenting the Algorithm and the results let us note the differences between this problem and the \mathbb{S}^{d-1} synchronization problem, presented above. For the \mathbb{S}^{d-1} case, the main difficulty that we faced in trying to obtain candidate solutions from eigenvectors was the local unit norm constraint. This is due to the fact that the synchronization problem requires its solution to be a function from V to \mathbb{S}^{d-1} , corresponding to a vector in \mathbb{R}^{dn} whose vertex subvectors have unit norm, while the eigenvector, in general, does not satisfy such a constraint. Nevertheless, the results in the previous section show that, by simply rounding the eigenvector,

one does not lose more than a linear term, given that the graph Laplacian has a spectral gap bounded away from zero.

However, the $O(d)$ synchronization setting is more involved. The reason being that, besides the local normalization constraint, there is also a local orthogonality constraint (at each vertex, the d vectors have to be orthogonal so that they can be the columns of an orthogonal matrix). For \mathbb{S}^{d-1} we locally normalized the vectors, by choosing for each vertex the unit vector closest to x_i . For $O(d)$ synchronization we will pick, for each vertex, the orthogonal matrix closest (in the Frobenius norm) to the matrix $[x_i^1 \cdots x_i^d]$, where x_i^j corresponds to the d -dimensional vector assigned to vertex i by the j 'th eigenvector. This rounding can be achieved by the Polar decomposition. Given a $d \times d$ matrix X , the matrix $U(X)$, solution of $\min_{U \in O(d)} \|U(X) - X\|_F$, is one of the components of the Polar decomposition of X (see [16, 20] and references therein). We note that $U(X)$ can be computed efficiently through the SVD decomposition of X . In fact, given the SVD of X , $X = U\Sigma V^T$, the closest orthogonal matrix to X is given by $U(X) = UV^T$ (see [16]). This approach is made precise in the following spectral algorithm for $O(d)$ -synchronization.

ALGORITHM 2.5. *Given a weighted graph $G = (V, E)$ and a function $\rho : E \rightarrow O(d)$, construct the normalized Connection Laplacian \mathcal{L}_1 and the degree matrix D_1 . Compute z^1, \dots, z^d , the first d eigenvectors corresponding to the d smallest eigenvalues of \mathcal{L}_1 . Let $x^j = D_1^{-\frac{1}{2}} z^j$, for each $j = 1, \dots, d$. Output the solution $g : V \rightarrow O(d)$ where each g_i is defined as*

$$g_i = U(X_i),$$

where $X_i = [x_i^1 \cdots x_i^d]$ and $U(X_i)$ is the closest orthogonal matrix of X_i , which can be computed via the SVD of X_i , if $X_i = U_i \Sigma_i V_i^T$, then $U(X_i) = U_i V_i^T$. If X_i is singular² simply set $U(X_i)$ to be I_d .

Similarly to how the performance of the \mathbb{S}^{d-1} synchronization algorithms was obtained, Lemma 3.11 bounds the effect of the rounding step in Algorithm 2.5. Before rounding, the frustration of the solution $[x^1 \cdots x^d]$ is $\frac{1}{d} \sum_{i=1}^d \eta(x^i)$. Lemma 3.11 guarantees that the solution g obtained by the rounding in Algorithm 2.5 satisfies $\nu(g) \leq 1026d^3 \frac{1}{\lambda_2(\mathcal{L}_0)} \sum_{i=1}^d \eta(x^i)$. Because of how the vectors x^1, \dots, x^d were built, $\sum_{i=1}^d \eta(x^i) = \sum_{i=1}^d \lambda_i(\mathcal{L}_1)$, and this means that the solution g computed by Algorithm 2.5 satisfies

$$\nu(g) \leq 1026d^3 \frac{1}{\lambda_2(\mathcal{L}_0)} \sum_{i=1}^d \lambda_i(\mathcal{L}_1).$$

This performance guarantee automatically implies our main result, a Cheeger inequality for the Connection Laplacian.

THEOREM 2.6. *Let $\lambda_i(\mathcal{L}_1)$ and $\lambda_i(\mathcal{L}_0)$ denote the i -th smallest eigenvalues of the normalized Connection Laplacian \mathcal{L}_1 and the normalized graph Laplacian \mathcal{L}_0 respectively. Let ν_G denote the frustration constant for $O(d)$ Synchronization. Then,*

$$\frac{1}{d} \sum_{i=1}^d \lambda_i(\mathcal{L}_1) \leq \nu_G \leq 1026d^3 \frac{1}{\lambda_2(\mathcal{L}_0)} \sum_{i=1}^d \lambda_i(\mathcal{L}_1).$$

Note that, once again, the lower bound is trivially obtained by noting that the eigenvector problem is a relaxation of the original synchronization problem.

Although the rigorous statement and proof of Lemma 3.11 will be presented in Section 3, we give a brief intuitive explanation of how the result is obtained.

²In this case the uniqueness of $U(X_i)$ is not guaranteed and thus the map is not well-defined.

As discussed above, the performance guarantee for Algorithm 2.3 relies on a proper understanding of the effect of the rounding step. In particular we showed that if $\lambda_2(\mathcal{L}_0)$ is small, then locally normalizing the candidate solution (which corresponds to the rounding step) has an effect over the penalty function that we can control. The case of $O(d)$ Synchronization is dealt with similarly. Instead of local normalization, the rounding step for Algorithm 2.5 is based on the polar decomposition. We start by understanding when the polar decomposition is stable (in the sense of changing the penalty function on a given edge) and see that this is the case when the candidate solution $X_i \in \mathbb{R}^{d \times d}$ is not close to being singular. The idea then is to show that only a small portion (which will depend on $\sum_{i=1}^d \lambda_i(\mathcal{L}_1)$ and $\lambda_2(\mathcal{L}_0)$) of the graph can have candidate solutions X_i close to singular and use that to show that, overall, we can bound the harmful contribution potentially caused by the rounding procedure on the penalty function.

3. Proof of the main results. In this Section we prove the results described above.

3.1. Proofs for Synchronization in \mathbb{S}^{d-1} . We start with the main Lemma regarding partial \mathbb{S}^{d-1} Synchronization.

LEMMA 3.1. *Given $x \in \mathbb{R}^{dn}$ there exists $u > 0$ such that*

$$\eta(x^u) \leq \sqrt{10\eta(x)}.$$

Moreover, if $d = 1$ the right-hand side can be replaced by $\sqrt{8\eta(x)}$.

Proof. This Lemma immediately follows from Lemma 3.2 as

$$\eta(x^u) = \frac{1}{2} \frac{\sum_{ij} w_{ij} \|x_i^u - \rho_{ij} x_j^u\|^2}{\sum_i d_i \|x_i^u\|^2} \leq \left(\frac{1}{2} \max_{i,j} \|x_i^u - \rho_{ij} x_j^u\| \right) \frac{\sum_{ij} w_{ij} \|x_i^u - \rho_{ij} x_j^u\|}{\sum_i d_i \|x_i^u\|} \leq \frac{\sum_{ij} w_{ij} \|x_i^u - \rho_{ij} x_j^u\|}{\sum_i d_i \|x_i^u\|},$$

where the last inequality was obtained by noting that $\|x_i^u - \rho_{ij} x_j^u\| \leq \|x_i^u\| + \|x_j^u\| \leq 2\mathbb{1}$

LEMMA 3.2. *Given $x \in \mathbb{R}^{dn}$ there exists $u > 0$ such that*

$$\frac{\sum_{ij} w_{ij} \|x_i^u - \rho_{ij} x_j^u\|}{\sum_i d_i \|x_i^u\|} \leq \sqrt{10\eta(x)}.$$

Moreover, if $d = 1$ the right-hand side can be replaced by $\sqrt{8\eta(x)}$.

Proof. Let us suppose, without loss of generality, that x is normalized so that $\max_i \|x_i\| = 1$. We will use a probabilistic argument. Let us consider the random variable u drawn uniformly from $[0, 1]$ and recall that x^u is defined by $x_i^u = \frac{x_i}{\|x_i\|}$ if $\|x_i\|^2 > u$ or $x_i^u = 0$ if $\|x_i\|^2 \leq u$. We will show that $\frac{1}{2} \frac{\mathbb{E} \sum_{ij} w_{ij} \|x_i^u - \rho_{ij} x_j^u\|}{\mathbb{E} \sum_i d_i \|x_i^u\|} \leq \sqrt{\frac{5}{2}\eta(x)}$, which implies that at least one of the realizations of u must satisfy the inequality, and proves the Lemma.

We start by showing that, for each edge (i, j) ,

$$\mathbb{E} \|x_i^u - \rho_{ij} x_j^u\| \leq \frac{\sqrt{5}}{2} \|x_i - \rho_{ij} x_j\| (\|x_i\| + \|x_j\|). \quad (3.1)$$

Without loss of generality we can consider $\rho_{ij} = I$ and $\|x_j\| \leq \|x_i\|$ and get,

$$\mathbb{E} \|x_i^u - x_j^u\| = \|x_j\|^2 \left\| \frac{x_i}{\|x_i\|} - \frac{x_j}{\|x_j\|} \right\| + (\|x_i\|^2 - \|x_j\|^2).$$

Thus, it suffices to show

$$\|x_j\|^2 \left\| \frac{x_i}{\|x_i\|} - \frac{x_j}{\|x_j\|} \right\| + (\|x_i\|^2 - \|x_j\|^2) \leq \frac{\sqrt{5}}{2} \|x_i - x_j\| (\|x_i\| + \|x_j\|),$$

which is a consequence of Proposition A.1 for $y = \frac{x_j}{\|x_j\|}$, $z = \frac{x_i}{\|x_i\|}$ and $\alpha = \frac{\|x_i\|}{\|x_j\|}$. Now, using (3.1), the linearity of expectation, and the Cauchy-Schwartz inequality we have

$$\begin{aligned} \mathbb{E} \sum_{ij} w_{ij} \|x_i^u - \rho_{ij} x_j^u\| &\leq \frac{\sqrt{5}}{2} \sum_{ij} w_{ij} \|x_i - \rho_{ij} x_j\| (\|x_i\| + \|x_j\|) \\ &\leq \frac{\sqrt{5}}{2} \sqrt{\sum_{ij} w_{ij} \|x_i - \rho_{ij} x_j\|^2} \sqrt{\sum_{ij} w_{ij} (\|x_i\| + \|x_j\|)^2}. \end{aligned}$$

Since $\sum_{ij} w_{ij} \|x_i - \rho_{ij} x_j\|^2 = 2\eta(x) \sum_i d_i \|x_i\|^2$ and

$$\sum_{ij} w_{ij} (\|x_i\| + \|x_j\|)^2 \leq 2 \sum_{ij} w_{ij} (\|x_i\|^2 + \|x_j\|^2) = 4 \sum_i d_i \|x_i\|^2,$$

we have

$$\mathbb{E} \sum_{ij} w_{ij} \|x_i^u - \rho_{ij} x_j^u\| \leq \frac{\sqrt{5}}{2} \sqrt{8\eta(x)} \sum_i d_i \|x_i\|^2 = \frac{\sqrt{5}}{2} \sqrt{8\eta(x)} \mathbb{E} \sum_i d_i \|x_i^u\| = 2\sqrt{\frac{5}{2}\eta(x)} \mathbb{E} \sum_i d_i \|x_i^u\|,$$

which completes the proof. When $d = 1$ the sharper result can be obtained by noting that (3.1) holds even without the $\frac{\sqrt{5}}{2}$ factor. \square

In the context of full \mathbb{S}^{d-1} Synchronization, the intuition given to justify a requirement in connectivity of the graph is that it forces the solution of the relaxed problem to have balanced norm across the vertices of the graph. The following Lemma makes this thought precise.

LEMMA 3.3. *Given $x \in \mathbb{R}^{dn}$, there exists $\alpha_x \geq 0$, such that $r_x = x - \alpha_x \tilde{x}$ satisfies $\|r_x\|_{D_1}^2 \leq \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)} \|x\|_{D_1}^2$. *Proof.* Let us define $n_x \in \mathbb{R}^n$ by $(n_x)_i = \|x_i\|$ and recall \tilde{x} defined in (1.2). We now set $\alpha_x = \operatorname{argmin}_{\alpha} \|\alpha n_x - \mathbf{1}\|_{D_0}$. A simple calculation reveals that this gives*

$$\alpha_x = \frac{\mathbf{1}^T D_0 n_x}{\mathbf{1}^T D_0 \mathbf{1}}.$$

Since n_x is a non-negative vector, α_x is non-negative as well. Let us also define $u_x \in \mathbb{R}^n$ so that $(r_x)_i = (u_x)_i \tilde{x}_i$. This implies that $u_x = n_x - \left(\frac{\mathbf{1}^T D_0 n_x}{\mathbf{1}^T D_0 \mathbf{1}}\right) \mathbf{1}$. Thus,

$$u_x^T L_0 u_x = n_x^T L_0 n_x = \frac{1}{2} \sum_{ij} w_{ij} (\|x_i\| - \|x_j\|)^2 \leq \frac{1}{2} \sum_{ij} w_{ij} \|x_i - \rho_{ij} x_j\|^2 = \eta(x) \|x\|_{D_1}^2,$$

Since $u_x^T D_0 \mathbf{1} = 0$, we have $\frac{(u_x)^T L_0 u_x}{\|u_x\|_{D_0}^2} \geq \lambda_2(\mathcal{L}_0)$. This shows that $\|r_x\|_{D_1}^2 = \|u_x\|_{D_0}^2 \leq \frac{1}{\lambda_2(\mathcal{L}_0)} \eta(x) \|x\|_{D_1}^2$. \square

This allows one to bound the volume of the subset of vertices where the norm of x is not typical. Let us first define this set.

DEFINITION 3.4. *Given $x \in \mathbb{R}^{dn}$, normalized so that $\|x\|_{D_1}^2 = \operatorname{vol}(G)$, and a positive number δ , we define the ill-balanced vertex subset of the graph G as $\mathcal{I}b_x(\delta) = \{i \in V : \|\|x_i\| - 1\| \geq \delta\}$.*

The volume of $\mathcal{I}b_x(\delta)$ is controlled by the following Lemma.

LEMMA 3.5. *Let $x \in \mathbb{R}^{dn}$ satisfy $\|x\|_{D_1}^2 = \operatorname{vol}(G)$. Then,*

$$\frac{\operatorname{vol}(\mathcal{I}b_x(\delta))}{\operatorname{vol}(G)} \leq \frac{4}{\delta^2} \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}.$$

Proof. Lemma 3.3 guarantees the existence of $\alpha_x \in \mathbb{R}^+$ such that $r_x = x - \alpha_x \tilde{x}$ satisfies $\|r_x\|_{D_1}^2 \leq \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)} \|x\|_{D_1}^2$.

Let us start by bounding α_x ; by the triangle inequality,

$$(1 - \alpha_x)^2 \text{vol}(G) = (\|x\|_{D_1} - \alpha_x \|\tilde{x}\|_{D_1})^2 \leq \|r_x\|_{D_1}^2 \leq \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)} \text{vol}(G),$$

which implies $(1 - \alpha_x)^2 \leq \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}$.

If $i \in \mathcal{I}b_x(\delta)$ then $|\|x_i\| - 1| \geq \delta$, which implies $\|(r_x)_i\| = |\|x_i\| - \alpha_x| \geq |\|x_i\| - 1| - |1 - \alpha_x| \geq \delta - \sqrt{\frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}}$. Squaring both sides of the inequality and summing over all $i \in \mathcal{I}b_x(\delta)$ gives,

$$\frac{\eta(x)}{\lambda_2(\mathcal{L}_0)} \text{vol}(G) \geq \|r_x\|_{D_1}^2 \geq \sum_{i \in \mathcal{I}b_k} d_i \|(r_x)_i\|_2^2 \geq \text{vol}(\mathcal{I}b_k) \left(\delta - \sqrt{\frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}} \right)^2, \quad (3.2)$$

as long as $\delta > \sqrt{\frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}}$. Let us separate in two cases:

If $\frac{\delta}{2} > \sqrt{\frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}}$, then, using (3.2) we have,

$$\frac{\text{vol}(\mathcal{I}b_k)}{\text{vol}(G)} \leq \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)} \left(\delta - \sqrt{\frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}} \right)^{-2} \leq \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)} \left(\delta - \frac{\delta}{2} \right)^{-2} = \frac{4}{\delta^2} \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}.$$

If, on the other hand, $\frac{\delta}{2} \leq \sqrt{\frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}}$, then, since $\frac{\text{vol}(\mathcal{I}b_k)}{\text{vol}(G)} \leq 1$,

$$\frac{\text{vol}(\mathcal{I}b_k)}{\text{vol}(G)} \leq 1 \leq \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)} \left(\frac{\delta}{2} \right)^{-2} = \frac{4}{\delta^2} \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}.$$

□

By placing an upper bound on the number of Ill-balanced vertices, (Lemma 3.5) we can control how much $\eta(x)$ is affected when we locally normalize x . This is achieved in the following Lemma, which contains the central technical result regarding full \mathbb{S}^{d-1} Synchronization.

LEMMA 3.6. *For every $x \in \mathbb{R}^{dn}$, $\eta(\tilde{x}) \leq \frac{44}{\lambda_2(\mathcal{L}_0)} \eta(x)$.*

Proof. We want to bound $\eta(\tilde{x}) = \frac{1}{2 \text{vol}(G)} \sum_{ij} w_{ij} \|\tilde{x}_i - \rho_{ij} \tilde{x}_j\|^2$. Without loss of generality we can assume $\|x\|_{D_1}^2 = \text{vol}(G)$. Let $0 < \gamma < 1$, then

$$\begin{aligned} \eta(\tilde{x}) &\leq \frac{1}{2 \text{vol}(G)} \left(\sum_{i \in \mathcal{I}b_x(\gamma)} \sum_j w_{ij} \|\tilde{x}_i - \rho_{ij} \tilde{x}_j\|^2 + \sum_{j \in \mathcal{I}b_x(\gamma)} \sum_i w_{ij} \|\tilde{x}_i - \rho_{ij} \tilde{x}_j\|^2 + \sum_{i,j \notin \mathcal{I}b_x(\gamma)} w_{ij} \|\tilde{x}_i - \rho_{ij} \tilde{x}_j\|^2 \right) \\ &\leq 4 \frac{\text{vol}(\mathcal{I}b_x(\gamma))}{\text{vol}(G)} + \frac{1}{2 \text{vol}(G)} \sum_{i,j \notin \mathcal{I}b_x(\gamma)} w_{ij} \|\tilde{x}_i - \rho_{ij} \tilde{x}_j\|^2. \end{aligned}$$

By Lemma 3.5 we have $4 \frac{\text{vol}(\mathcal{I}b_x(\gamma))}{\text{vol}(G)} \leq \frac{16}{\gamma^2} \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)}$.

Note that, for any $y, z \in \mathbb{R}^d$, $\left\| \frac{y}{\|y\|} - \frac{z}{\|z\|} \right\| \leq \frac{\|y-z\|}{\min\{\|y\|, \|z\|\}}$. By setting $y = x_i$ and $z = \rho_{ij}x_j$ we get $\|\tilde{x}_i - \rho_{ij}\tilde{x}_j\| \leq \frac{\|x_i - \rho_{ij}x_j\|}{\min\{\|x_i\|, \|x_j\|\}}$. This implies that

$$\begin{aligned} \frac{1}{2 \operatorname{vol}(G)} \sum_{i,j \notin \mathcal{I}b_x(\gamma)} w_{ij} \|\tilde{x}_i - \rho_{ij}\tilde{x}_j\|^2 &\leq \frac{1}{2 \operatorname{vol}(G)} \sum_{i,j \notin \mathcal{I}b_x(\gamma)} w_{ij} \left(\frac{\|x_i - \rho_{ij}x_j\|}{\min\{\|x_i\|, \|x_j\|\}} \right)^2 \\ &\leq \frac{1}{2 \operatorname{vol}(G)} \frac{1}{(1-\gamma)^2} \sum_{i,j \notin \mathcal{I}b_x(\gamma)} w_{ij} \|x_i - \rho_{ij}x_j\|^2. \end{aligned}$$

This means that

$$\begin{aligned} \eta(\tilde{x}) &\leq \frac{16}{\gamma^2} \frac{\eta(x)}{\lambda_2(\mathcal{L}_0)} + \frac{1}{2 \operatorname{vol}(G)} \frac{1}{(1-\gamma)^2} \sum_{i,j \notin \mathcal{I}b_x(\gamma)} w_{ij} \|x_i - \rho_{ij}x_j\|^2 \\ &\leq \left(\frac{16}{\gamma^2} \frac{1}{\lambda_2(\mathcal{L}_0)} + \frac{1}{(1-\gamma)^2} \right) \eta(x). \end{aligned}$$

Since $\lambda_2(\mathcal{L}_0) \leq 1$ (see, e.g., [9]), it is possible to pick γ (e.g. 0.7) such that $\frac{16}{\gamma^2} \frac{1}{\lambda_2(\mathcal{L}_0)} + \frac{1}{(1-\gamma)^2} \leq \frac{44}{\lambda_2(\mathcal{L}_0)}$. \square

3.2. Proofs for Synchronization in $O(d)$. As described, the rounding procedure for $O(d)$ Synchronization is based on the polar decomposition. We need to understand how much the polar decomposition can potentially affect the penalty on each edge. With this purpose we use the following result of Li [20].

LEMMA 3.7 (Theorem 1 in [20]). *Let $A, B \in \mathbb{C}^{d \times d}$ be non-singular matrices with polar decompositions $A = U(A)P$ and $B = U(B)P'$. Then $\|U(A) - U(B)\|_F \leq \frac{2}{\sigma_{\min}(A) + \sigma_{\min}(B)} \|A - B\|_F$, where $\sigma_{\min}(A)$ is the smallest singular value of the matrix A .*

This means that, in order to bound possible instabilities of the polar decomposition, one needs to control the size of the smallest singular value of $X_i = [x_i^1 \cdots x_i^d]$. To achieve this, we will introduce a notion similar to $\mathcal{I}b_x(\delta)$ but designed to take into account local orthogonality instead of local normalization.

DEFINITION 3.8. *Given $x, y \in \mathbb{R}^{dn}$ two D_1 -orthogonal vectors, normalized so that $\|x\|_{D_1}^2 = \|y\|_{D_1}^2 = \operatorname{vol}(G)$, and a positive number δ , we define the following ill-balanced vertex subset of the graph G as*

$$\mathcal{I}b_{xy}(\delta) = \{i \in V : |\langle x_i, y_i \rangle| \geq \delta\}.$$

The next Lemma shows that, for an edge whose incident vertices are not ill-balanced (see Definitions 3.8 and 3.4), the polar decomposition only slightly affects the penalty function.

LEMMA 3.9. *Let $x^1, \dots, x^d \in \mathbb{R}^{dn}$ be D_1 -orthogonal vectors, normalized so that $\|x^k\|_{D_1}^2 = \operatorname{vol}(G)$. Let us define the “balanced” set \mathcal{B} as the complement of $\bigcup_{k \in [d]} \left(\mathcal{I}b_{x^k} \left(\frac{1}{8d} \right) \cup \bigcup_{m \in [d] \setminus \{k\}} \mathcal{I}b_{x^k x^m} \left(\frac{1}{2d} \right) \right)$. For all $i, j \in \mathcal{B}$, we have*

$$\|U(X_i) - \rho_{ij}U(X_j)\|_F \leq \sqrt{2} \|X_i - \rho_{ij}X_j\|_F.$$

Proof. For $i \in \mathcal{B}$, consider the gram matrix $X_i^T X_i$. Its k -th diagonal entry satisfies $\|x_i^k\|^2 \geq \left(1 - \frac{1}{8d}\right)^2 \geq 1 - \frac{1}{4d}$. On the other hand the non-diagonal entries are, in magnitude, smaller or equal to $\frac{1}{2d}$. By the Gershgorin circle theorem, the smallest eigenvalue of $X_i^T X_i$, which is equal to $\sigma_{\min}(X_i)^2$, satisfies $\sigma_{\min}(X_i)^2 \geq 1 - \frac{1}{4d} - (d-1)\frac{1}{2d}$. Hence, $\sigma_{\min}(X_i) \geq \frac{1}{\sqrt{2}}$. By observing that $U(\rho_{ij}X_j) = \rho_{ij}U(X_j)$, and

using Lemma 3.7, we get $\|U(X_i) - \rho_{ij}U(X_j)\|_F \leq \sqrt{2}\|X_i - \rho_{ij}X_j\|_F$. \square

The last step is to control the size of the ill-balanced sets.

LEMMA 3.10. *Let $x, y \in \mathbb{R}^{dn}$ be D_1 -orthogonal vectors such that $\|x\|_{D_1}^2 = \|y\|_{D_1}^2 = \text{vol}(G)$. Then,*

$$\frac{\text{vol}(\mathcal{I}b_{xy}(\frac{1}{2d}) \setminus (\mathcal{I}b_x(\frac{1}{8d}) \cup \mathcal{I}b_y(\frac{1}{8d})))}{\text{vol}(G)} \leq 4(8d)^2 \frac{\eta(x) + \eta(y)}{\lambda_2(\mathcal{L}_0)}.$$

Proof. Let us consider the vector $u = \frac{1}{\sqrt{2}}(x+y)$. It satisfies $\|u\|_{D_1}^2 = \text{vol}(G)$ and by the triangle inequality on the norm $\|\cdot\|_{L_1}$, $\eta(u) \leq \eta(x) + \eta(y)$. By Lemma 3.5 we get $\frac{\text{vol}(\mathcal{I}b_u(\frac{1}{8d}))}{\text{vol}(G)} \leq 4(8d)^2 \frac{\eta(x) + \eta(y)}{\lambda_2(\mathcal{L}_0)}$. We conclude the proof by noting that $\mathcal{I}b_{xy}(\frac{1}{2d}) \subset \mathcal{I}b_x(\frac{1}{8d}) \cup \mathcal{I}b_y(\frac{1}{8d}) \cup \mathcal{I}b_u(\frac{1}{8d})$. In fact, if $i \notin \mathcal{I}b_x(\frac{1}{8d}) \cup \mathcal{I}b_y(\frac{1}{8d}) \cup \mathcal{I}b_u(\frac{1}{8d})$ then $|\langle x_i, y_i \rangle| = \left| \|u_i\|^2 - \frac{\|x_i\|^2 + \|y_i\|^2}{2} \right| \leq \left(1 + \frac{1}{8d}\right)^2 - \left(1 - \frac{1}{8d}\right)^2 = \frac{1}{2d}$. \square

At this point we have build all the foundations needed for the proof of the central lemma regarding $O(d)$ Synchronization.

LEMMA 3.11. *Given $x^1, \dots, x^d \in \mathbb{R}^{dn}$ such that $\langle x^k, x^l \rangle_{D_1} = 0$ for all $k \neq l$, consider the potential $g : V \rightarrow O(d)$ given as $g_i = U(X_i)$ where $X_i = [x_i^1 \dots x_i^d]$ and $U(X)$ is the closest (in the Frobenius norm) orthogonal matrix of X . If X_i is singular $U(X_i)$ is simply set to be I_d . Then,*

$$\nu(g) \leq (2d^{-1} + 2^{10}d^3) \frac{1}{\lambda_2(\mathcal{L}_0)} \sum_{i=1}^d \eta(x^i).$$

Proof.

Let us consider \mathcal{B} as defined in Lemma 3.9, meaning $\mathcal{B}^c = \bigcup_{k \in [d]} \left(\mathcal{I}b_{x^k}(\frac{1}{8d}) \cup \bigcup_{m \in [d] \setminus \{k\}} \mathcal{I}b_{x^k x^m}(\frac{1}{2d}) \right)$. We want to bound $\nu(g) = \frac{1}{2d \text{vol}(G)} \sum_{ij} w_{ij} \|g_i - \rho_{ij}g_j\|_F^2$. Since $(\mathcal{B} \times \mathcal{B})^c \subset (\mathcal{B}^c \times V) \cup (V \times \mathcal{B}^c)$,

$$\begin{aligned} \sum_{ij} w_{ij} \|g_i - \rho_{ij}g_j\|_F^2 &\leq \sum_{(i,j) \in \mathcal{B} \times \mathcal{B}} w_{ij} \|U(X_i) - \rho_{ij}U(X_j)\|_F^2 + 2 \sum_{i \in \mathcal{B}^c} \sum_{j \in V} w_{ij} \|g_i - \rho_{ij}g_j\|_F^2 \\ &\leq 2 \sum_{ij} w_{ij} \|X_i - \rho_{ij}X_j\|_F^2 + 8d \text{vol}(\mathcal{B}^c), \end{aligned}$$

where the second inequality was obtained by using Lemma 3.9 and noting that O_i is an orthogonal matrix. To bound $\text{vol}(\mathcal{B}^c)$ we make use of Lemmas 3.5 and 3.10 and get that $\frac{\text{vol}(\mathcal{B}^c)}{\text{vol}(G)}$ is bounded above by

$$\sum_{k \in [d]} \frac{\text{vol}(\mathcal{I}b_{x^k}(\frac{1}{8d}))}{\text{vol}(G)} + \frac{1}{2} \sum_{k \in [d]} \sum_{m \in [d] \setminus \{k\}} \frac{\text{vol}(\mathcal{I}b_{x^k x^m}(\frac{1}{2d}) \setminus (\mathcal{I}b_x(\frac{1}{8d}) \cup \mathcal{I}b_y(\frac{1}{8d})))}{\text{vol}(G)} \leq 2^8 d^3 \frac{\sum_{k \in [d]} \eta(x^k)}{\lambda_2(\mathcal{L}_0)}.$$

Since $\sum_{ij} w_{ij} \|X_i - \rho_{ij}X_j\|_F^2 = 2 \text{vol}(G) \sum_{k=1}^d \eta(x^k)$, we get $\nu(g) \leq (2d^{-1} + 2^{10}d^3) \frac{1}{\lambda_2(\mathcal{L}_0)} \sum_{k=1}^d \eta(x^k)$. \square

4. An unsquared version of the frustration constant. We formulated the Synchronization problem as minimizing the square of the Frobenius norm of the incompatibilities (in some sense, an ℓ_2 penalty function). This penalty function is particularly nice because it is close in spirit to the Rayleigh quotient formulation of an eigenvector problem and thus more related to what the spectral method will try to minimize, as we indeed showed in the theorems above.

On the other hand, considering the sum of the Frobenius norms of the incompatibilities (in some sense, an ℓ_1 penalty function) will induce sparsity on the edge inconsistencies, meaning that it will favor candidate solutions for which some edges are perfectly correct even if there are some edges with large errors. This type of penalty function is often favorable under some noise models. In fact, the noise model analyzed in [25] consists of a few randomly chosen edges having a measurement that is randomly drawn with respect to the uniform distribution in the space of possible measurements (in our case $O(d)$, in [25] $SO(2)$), therefore the original rotation potential will perfectly agree with some edges and have a large error on others. This motivates us to look at ℓ_1 versions of frustration constants. Let us define

$$\nu_1(O) = \frac{1}{\sqrt{d} \text{vol}(G)} \sum_{ij} w_{ij} \|O_i - \rho_{ij} O_j\|_F,$$

and the $O(d)$ frustration ℓ_1 constant of G as $\nu_{G,1} = \min_{O:V \rightarrow O(d)} \vartheta(O)$. Similarly,

$$\eta_1(v) = \frac{\sum_{ij} w_{ij} \|v_i - \rho_{ij} v_j\|}{\sum_i d_i \|v_i\|}$$

and the ℓ_1 constant of G , as $\eta_{G,1} = \min_{v:V \rightarrow S^{d-1}} \zeta(v)$. We also define a partial version of it $\eta_{G,1}^* = \min_{v:V \rightarrow S^{d-1} \cup \{0\}} \zeta(v)$.

From our results in the above section it is easy to obtain the following Cheeger type inequalities for these frustration constants

THEOREM 4.1. *Let $\lambda_i(\mathcal{L}_1)$ and $\lambda_i(\mathcal{L}_0)$ denote the i -th smallest eigenvalue of, respectively, the normalized Connection Laplacian \mathcal{L}_1 and the normalized graph Laplacian \mathcal{L}_0 . Let $\eta_{G,1}^*$, $\eta_{G,1}$, and $\nu_{G,1}$, denote the ℓ_1 frustration constants defined above. Then,*

$$\lambda_1(\mathcal{L}_1) \leq \eta_{G,1}^* \leq \sqrt{10\lambda_1(\mathcal{L}_1)}, \quad (4.1)$$

$$\lambda_1(\mathcal{L}_1) \leq \eta_{G,1} \leq 2\sqrt{\frac{22}{\lambda_2(\mathcal{L}_0)} \lambda_1(\mathcal{L}_1)}, \quad (4.2)$$

and,

$$\frac{1}{d} \sum_{i=1}^d \lambda_i(\mathcal{L}_1) \leq \nu_{G,1} \leq 6d \sqrt{\frac{57d}{\lambda_2(\mathcal{L}_0)} \sum_{i=1}^d \lambda_i(\mathcal{L}_1)}. \quad (4.3)$$

Proof.

For any $v : V \rightarrow S^{d-1} \cup \{0\}$ we have

$$\eta(v) = \frac{1}{2} \frac{\sum_{ij} w_{ij} \|v_i - \rho_{ij} v_j\|^2}{\sum_i d_i \|v_i\|^2} \leq \frac{\left(\sum_{ij} w_{ij} \|v_i - \rho_{ij} v_j\|\right) \frac{\max_{ij} \|v_i - \rho_{ij} v_j\|}{2}}{(\sum_i d_i \|v_i\|) \min \|v_i\|} \leq \frac{\sum_{ij} w_{ij} \|v_i - \rho_{ij} v_j\|}{\sum_i d_i \|v_i\|} = \eta_1(v),$$

which, together with Theorems 2.2 and 2.4, gives the lower bound on both (4.1) and (4.2).

Note that Lemma 3.2 actually guarantees that there exists $v : V \rightarrow S^{d-1} \cup \{0\}$ such that $\eta_1(v) \leq \sqrt{10\lambda_1(\mathcal{L}_1)}$ which concludes the proof of (4.1).

Let $v : V \rightarrow S^{d-1} \cup \{0\}$ be a solution that satisfies $\eta(v) \leq 44 \frac{\lambda_1(\mathcal{L}_1)}{\lambda_2(\mathcal{L}_0)}$, guaranteed to exist by Theorem 2.4.

We then have,

$$\begin{aligned} \eta_1(v) &= \frac{\sum_{ij} w_{ij} \|v_i - \rho_{ij} v_j\|}{\sum_i d_i \|v_i\|} = \frac{1}{\text{vol}(G)} \sum_{ij} \left(w_{ij}^{\frac{1}{2}} \|v_i - \rho_{ij} v_j\| \right) w_{ij}^{\frac{1}{2}} \\ &\leq \frac{1}{\text{vol}(G)} \left(\sum_{ij} w_{ij} \|v_i - \rho_{ij} v_j\|^2 \right)^{\frac{1}{2}} \left(\sum_{ij} w_{ij} \right)^{\frac{1}{2}} = \sqrt{\frac{\sum_{ij} w_{ij} \|v_i - \rho_{ij} v_j\|^2}{\sum_i d_i \|v_i\|^2}} = \sqrt{2\eta(v)}, \end{aligned} \quad (4.4)$$

where the inequality is obtained using Cauchy-Schwarz. This completes the proof of (4.2).

Inequality (4.3) is shown in the same way as (4.2): since $g_i, g_j \in O(d)$, we have $\|g_i - \rho_{ij} g_j\|_F \leq 2\sqrt{d}$ which gives $\nu_1(g) \geq \nu(g)$. On the other hand, using Cauchy Schwarz in the same way as in (4.4) gives $\nu_1(g) \leq \sqrt{2\nu(g)}$, which implies (4.3) and concludes the proof of the Theorem.

□

In fact Wang and Singer [34] recently showed that under the random outlier's noise model (described above), a semidefinite relaxation of the $O(d)$ Synchronization problem formulated with the ℓ_1 penalty function was able to recover the ground truth solution with high probability, provided the underlying graph is drawn from the Erdős-Rényi random graph model and the ratio of outliers is below a certain threshold.

5. Tightness of results. Let us consider the ring graph on n vertices $G_n = (V_n, E_n)$ with $V_n = [n]$ and $E = \{(i, (i+1) \bmod n), i \in [n]\}$ with the edge weights all equal to 1 and $\rho: V \rightarrow O(d)$ as $\rho_{(n,1)} = -I$ and $\rho = I$ for all other edges. Define $x \in \mathbb{R}^{dn}$ by $x_k = [2\frac{k}{n} - 1, 0, \dots, 0]^T$. It is easy to check that $\eta(x) = \mathcal{O}(n^{-2})$ and that, for any $u > 0$, if $x^u \neq 0$, there will have to be at least one edge that is not compatible with x^u , implying $\eta(x^u) \geq \frac{1}{2n}$. This shows that the $1/2$ exponent in Lemma 3.1 is needed. In fact, by adding a few more edges to the graph G_n one can also show the tightness of Theorem 2.2: Consider the ‘‘rainbow’’ graph H_n that is constructed by adding to G_n , for each non-negative integer k smaller than $n/2$, an edge between vertex k and vertex $n - k$ with $\rho_{(k, n-k)} = -I$. The vector x still satisfies $\eta(x) = \mathcal{O}(n^{-2})$, however, for any non-zero vector $v: V \rightarrow \mathbb{S}^{d-1} \cup \{0\}$, it is not hard to show that $\eta(v)$ has to be of order at least n^{-1} , meaning that η_G^* is $\Omega(\sqrt{\lambda_1(\mathcal{L}_1)})$. This also means that, even if considering η_G^* , one could not get a linear bound (as provided by Lemma 3.6) without the control on $\lambda_2(\mathcal{L}_0)$.

Theorem 2.6 provides a non-trivial bound only if $\lambda_2(\mathcal{L}_0)$ is sufficiently large. It is clear that if one wants to bound full frustration constants, a dependency on $\lambda_2(\mathcal{L}_0)$ is needed. It is, nevertheless, non-obvious that this dependency is still needed if we consider partial versions of $O(d)$ frustration constants, ϑ_G^* or ν_G^* . This can, however, be illustrated by a simple example in $O(2)$; consider a disconnected graph G with two sufficiently large complete components, $G^1 = (V^1, E^1)$ and $G^2 = (V^2, E^2)$. For each edge let $\rho_{i,j} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}$. It is clear that the vectors x^1 and x^2 defined such that $x_i^1 = [0, 1_{V^1}(i)]^T$ and $x_i^2 = [0, 1_{V^2}(i)]^T$ are orthogonal to each other and lie in the null space of the graph Connection Laplacian of G . This implies that $\lambda_2(\mathcal{L}_1) = 0$. On the other hand, it is straightforward to check that ν_G^* is not zero because it is impossible to perfectly synchronize the graph (or any of the components, for that matter).

6. Concluding Remarks. Synchronization is a challenging problem. Recent discoveries suggest that spectral relaxations are promising as feasible methods to solve this problem. In fact, in [25], probability guarantees of performance are given for the performance of a spectral method to solve the $SO(2)$ synchronization problem under a certain random noise model. Nevertheless, to the best of our knowledge, Algorithm 2.5 is the first method for $O(d)$ synchronization having a (deterministic) worst case performance guarantee. As one would expect, the worst case performance is significantly weaker than the kind of probabilistic guarantees, given a specific noise model, e.g. as the one given in [25]. In fact, the guarantees in [25] and [34] are given in terms of distance between the candidate solution and the ground truth, while the one we provide is given in terms of the compatibility error. Recently, in the context of the Phase Retrieval problem, the guarantees on this paper were used to obtain guarantees on the distance between the candidate solution and the ground truth (see [1]).

In special applications one knows, a priori, that every element in the potential has positive determinant. This corresponds to a synchronization problem in $SO(d)$. Although this can be viewed as a special case of the $O(d)$ problem it is expected that the additional structure can be leveraged to improve the algorithm (and the analysis). In particular, the first $d - 1$ columns of a matrix in $SO(d)$ completely determine the matrix. This suggests that the $SO(d)$ synchronization problem is solvable by just the first $d - 1$ eigenvectors of the Connection Laplacian, instead of the d first ones. In fact, the $SO(2)$ Synchronization problem is equivalent to the \mathbb{S}^1 localization one, and the guarantees for the \mathbb{S}^1 localization problem were given solely in terms of the first eigenvalue of the graph Connection Laplacian. We leave the improved $SO(d)$ analysis for future work.

The performance guarantee for this algorithm relies on the fact that the vectors obtained are D_1 -orthogonal. However, in practice, due to possible errors in the calculations this condition might be perturbed and the D_1 inner products of these vectors, although small, may no longer be exactly zero. It is easy to adapt the analysis to this setting and show that it is in fact robust to such perturbations.

The results in this paper also suggest an alternative to Algorithm 2.5 which corresponds to, instead of solving the eigenvector problem, determining the d vectors sequentially and, at each step, constraining on the vector being locally orthogonal to the previous ones (this can still be done efficiently, see [14]). After the d vectors are obtained one can simply locally normalize each one and output that as the Synchronization solution candidate. The issue with this method is that its iterative nature ³ makes its analysis more difficult, as it is hard to guarantee that small errors in the first few vectors would not greatly affect the remaining ones. Also, numerical simulations suggest that the performance, in practice, of both methods is roughly the same. Independently of which method is used, the solution, although guaranteed to have a certain performance, is not guaranteed to be a local optimum. Recently, Boumal et al. [5] suggest that an iterative smooth optimization (in manifolds) method, when started in the solution given by the rounding procedure, can be used to locally search for a better solution.

One might argue that, in some applications, the weights on the edges of the graph do not have a clear meaning. The reason being that we may be given a few relative measurements ρ_{ij} and it is unclear how to give weights to such measurements. In such cases, since we have the freedom of choosing the weights of the edges and Theorems 2.4 and 2.6 suggest that our method will work better with a large $\lambda_2(\mathcal{L}_0)$, one could compute the weights of the edges in such a way that $\lambda_2(\mathcal{L}_0)$ is maximized. This problem is solved in [31]. The caveat is that, the new weights will affect the way the compatibility error is measured, as well as the eigenvalues of the Connection Laplacian. It is thus still unclear if such an approach would improve the method. Another interesting possible outcome of a procedure of this nature is a possible ranking of the edges, large weights would likely tend to be given to edges that are more important to ensure the connectivity of the graph.

The classical Cheeger inequality has an analogous result on smooth manifolds (actually, the first to be shown [7]). One interesting question is whether the theorems in this paper have an analogous smooth version. One difficulty is to understand what would correspond to the frustration constant in the smooth case. Recent work on Vector Diffusion Maps [27] suggests that an analogous result on smooth manifolds would be related to the parallel transport and its incompatibility and some results in Differential Geometry [4] suggest that the Holonomy could be a geometric property that corresponds to the frustration constant. These suggestions are “coherent” because Holonomy can, in some sense, be viewed as the incompatibility of the Parallel transport (due to the curvature of the manifold).

In some cases the incompatibilities have some structure; the continuous setting described above may be such an example. Although, in this paper, we make no attempt to understand or take advantage of such structure, we believe this would be an interesting direction for future work and we direct the reader to papers on which topological tools are used to understand (and leverage) the structure of inconsistencies in synchronization-like problems [6, 18, 22].

³The fact that the calculation of one of the vectors greatly depends on the ones already computed.

Acknowledgments. The authors thank Hau-Tieng Wu for interesting discussions on the topic of this paper and Leor Klainerman for reading a preliminary version of this manuscript. The authors also acknowledge the comments made by the referees and the editor, that helped to significantly improve this paper. Afonso S. Bandeira was supported by Award Number DMS-0914892 from the NSF. Amit Singer was partially supported by Award Numbers FA9550-09-1-0551 and FA9550-12-1-0317 from AFOSR, Award Number R01GM090200 from the National Institute of General Medical Sciences, the Alfred P. Sloan Foundation, and Award Number LTR DTD 06-05-2012 from the Simons Foundation. Daniel A. Spielman was supported by the NSF under Grant No. 091548.

REFERENCES

- [1] B. Alexeev, A. S. Bandeira, M. Fickus, and D. G. Mixon. Phase retrieval with polarization. *available online*, 2012.
- [2] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6:83–96, 1986.
- [3] N. Alon and V. Milman. Isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory*, 38:73–88, 1985.
- [4] W. Ballmann, J. Brüning, and G. Carron. Eigenvalues and holonomy. *Int. Math. Res. Not.*, (12):657–665, 2003.
- [5] N. Boumal, A. Singer, P.-A. Absil, and V. D. Blondel. Cramér-rao bounds for synchronization of rotations. *CoRR*, abs/1211.1621, 2012.
- [6] O. Candogan, I. Menache, A. Ozdaglar, and P. A. Parrilo. Flows and decompositions of games: Harmonic and potential games. *arXiv:1005.2405v2*, 2010.
- [7] J. Cheeger. A lower bound for the smallest eigenvalue of the laplacian. *Problems in analysis (Papers dedicated to Salomon Bochner, 1969)*, pp. 195–199. Princeton Univ. Press, 1970.
- [8] F. Chung. Four proofs for the cheeger inequality and graph partition algorithms. *Fourth International Congress of Chinese Mathematicians*, pp. 331–349, 2010.
- [9] F. Chung and L. Lu. *Complex graphs and networks*. Number 107. Amer Mathematical Society, 2006.
- [10] M. Cucuringu, Y. Lipman, and A. Singer. Sensor network localization by eigenvector synchronization over the euclidean group. *ACM Transactions on Sensor Networks*, In press.
- [11] 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.
- [12] S. O. Gharan and L. Trevisan. A higher-order cheeger’s inequality. 2011.
- [13] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the Association for Computing Machinery*, 42:1115–1145, 1995.
- [14] G. H. Golub. Some modified matrix eigenvalue problems. *SIAM Rev.*, 15(2), 1973.
- [15] Hagen and Kahng. New spectral methods for ratio cut partitioning and clustering. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 11:1074–1085, 1992.
- [16] N. J. Higham. Computing the polar decomposition with applications. *SIAM J. Sci. Stat. Comput.*, 7:1160–1174, October 1986.
- [17] S. D. Howard, D. Cochran, W. Moran, and F. R. Cohen. Estimation and registration on graphs. Arxiv preprint, 2010. Available at <http://arxiv.org/abs/1010.2983>.
- [18] X. Jiang, L.-H. LIM, Y. Yao, and Y. Ye. Statistical ranking and combinatorial hodge theory. *arXiv:0811.1067v2*, 2009.
- [19] J.R. Lee, S.O. Gharan, and L. Trevisan. Multi-way spectral partitioning and higher-order cheeger inequalities. 2011.
- [20] R.-C. Li. New perturbation bounds for the unitary polar factor. *SIAM J. Matrix Anal. Appl.*, 16(1):327–332, January 1995.
- [21] F. McSherry. Spectral partitioning of random graphs. In *FOCS ’01: Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*, page 529, 2001.
- [22] P. Molavi and A. Jadbabaie. A topological view of estimation from noisy relative measurements. In *American Control Conference (ACC), 2011*, pages 3615–3620, 2011.
- [23] O. Parzanchevski and R. Rosenthal. Simplicial complexes: spectrum, homology and random walks. *arXiv:1211.6775v1*, 2012.
- [24] O. Parzanchevski, R. Rosenthal, and R. J. Tessler. Isoperimetric inequalities in simplicial complexes. *arXiv:1207.0638v2*, 2012.
- [25] A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmon. Anal.*, 30(1):20 – 36, 2011.
- [26] 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.
- [27] A. Singer and H.-T. Wu. Vector diffusion maps and the connection laplacian. *Comm. Pure Appl. Math.*, in press.
- [28] A. Singer and H.-T. Wu. Orientability and diffusion maps. *Appl. Comput. Harmon. Anal.*, 31(1):44–58, 2011.
- [29] A. Singer, Z. Zhao, , Y. Shkolnisky, and R. Hadani. Viewing angle classification of cryo-electron microscopy images using eigenvectors. *SIAM Journal on Imaging Sciences*, 4:723–759, 2011.
- [30] J. Steenbergen, C. Klivans, and S. Mukherjee. A cheeger-type inequality on simplicial complexes. *arXiv:1209.5091v3*, 2012.
- [31] J. Sun, S. Boyd, L. Xiao, and P. Diaconis. The fastest mixing markov process on a graph and a connection to a maximum

- variance unfolding problem. *SIAM Review*, 48(4):681–699, 2006.
- [32] L. Trevisan. Max cut and the smallest eigenvalue. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC '09, pages 263–272, New York, NY, USA, 2009. ACM.
- [33] T. Tzeneva. Global alignment of multiple 3-d scans using eigenvector synchronization. *Senior Thesis, Princeton University (supervised by S. Rusinkiewicz and A. Singer)*, 2011.
- [34] L. Wang and A. Singer. Exact and stable recovery of rotations for robust synchronization. *arXiv:1211.2441v1*, 2012.
- [35] S. X. Yu. Angular embedding: A robust quadratic criterion. *IEEE Trans. Pattern Anal. Mach. Intell.*, 34(1):158–173, 2012.

Appendix A. Some Technical Steps.

PROPOSITION A.1. *For any y and z unit vectors in \mathbb{R}^d , the following holds for any $\alpha \geq 1$,*

$$\|y - z\| + \alpha^2 - 1 \leq \frac{\sqrt{5}}{2} \|y - \alpha z\| (1 + \alpha).$$

Proof. Let $t = \|y - z\|$, which implies that $0 \leq t \leq 2$. Since y and z are unit vectors it is straightforward to check that $\|y - \alpha z\| = \sqrt{1 + \alpha^2 - 2\alpha(1 - \frac{1}{2}t^2)}$. Thus, it suffices to show

$$t + \alpha^2 - 1 \leq \frac{\sqrt{5}}{2} \sqrt{1 + \alpha^2 - 2\alpha\left(1 - \frac{1}{2}t^2\right)} (1 + \alpha), \quad (\text{A.1})$$

for all $0 \leq t \leq 2$ and $\alpha \geq 1$. Since both sides of (A.1) are positive, it is enough to show the inequality with both sides squared. Squaring and rearranging yields,

$$t^2 + 2t\alpha^2 - 2t + (\alpha^2 - 1)^2 \leq \frac{5}{4} ((\alpha^2 - 1)^2 + \alpha t^2 (1 + \alpha)^2).$$

This is equivalent to the non-negativity, in the interval $[0, 2]$, of a certain quadratic function of t :

$$\left(\frac{5}{4}\alpha(1 + \alpha)^2 - 1\right) t^2 - (2\alpha^2 - 2)t + \frac{1}{4}(\alpha^2 - 1)^2 \geq 0.$$

Since, for $\alpha \geq 1$,

$$(2\alpha^2 - 2)^2 - 4\left(\frac{5}{4}\alpha(1 + \alpha)^2 - 1\right)\frac{1}{4}(\alpha^2 - 1)^2 = (\alpha^2 - 1)^2 \left(4 - \frac{5}{4}\alpha(1 + \alpha)^2 + 1\right) \leq 0,$$

the quadratic is always non-negative and thus non-negative in $[0, 2]$. \square