Theoretical and Algorithmic Aspects of Rigid Registration

A Thesis
Submitted for the Degree of
Master of Technology (Research)
in the Faculty of Engineering

by

Aditya Vikram Singh

Department of Electrical Engineering
INDIAN INSTITUTE OF SCIENCE
Bangalore – 560 012
January 2019

Abstract

In this thesis, we consider the rigid registration problem, which arises in applications such as sensor network localization, multiview registration, and protein structure determination. The abstract setup for this problem is as follows. We are given a collection of N labelled points in d-dimensional Euclidean space. There are M observers, each of whom observes a subset of points and assigns coordinates to them in their local frame of reference. For each observer, we know which points they observe, and the (possibly noisy) local coordinates assigned to these points. Based on this information, we wish to infer the global coordinates of the N points. We investigate the following questions in this context:

1. Uniqueness: Suppose that the local coordinates are noiseless. In this case, we know that the true global coordinates are a solution of the problem. But is this the only solution? We use results from graph rigidity theory to give a necessary and sufficient condition for the problem to have a unique solution. In two-dimensional space, this leads to a particularly efficient connectivity-based test for uniqueness.

2. Tightness of a convex relaxation: In general, when the local coordinates are noisy, we use least squares fitting to estimate the global coordinates. After a suitable reduction, this can be posed as a rank-constrained semidefinite program (REG-SDP). Dropping the rank-constraint yields a convex relaxation, which has been empirically observed to solve REG-SDP when the noise is below a certain threshold. Motivated by an analysis of Bandeira et al [1], we offer an explanation of this phenomenon by analyzing the Lagrange dual of the relaxed problem.

3. Convergence of an iterative solver: Instead of working with a convex relaxation, we can try directly solving REG-SDP by appropriately splitting the constraint set, and formally applying the alternating direction method of multipliers (ADMM). Empirically, this nonconvex ADMM algorithm has been demonstrated to perform well in the context of multiview registration. We analyze convergence of the iterates generated by this algorithm, and show how noise in the measurements affects convergence behavior.

Acknowledgements

First and foremost, I would like to thank my research advisor Prof. Kunal Chaudhury. I feel truly fortunate to be advised by someone who values theoretical research. He was always calling my attention to intriguing computational phenomena arising in applications of rigid registration that demanded rigorous mathematical justification, and allowed me complete freedom to investigate these; this has been instrumental in my growth as a researcher. I have also benefited immensely from his close reading of our manuscripts and his invaluable tips on formatting and presenting results.

The courses I have credited and audited at IISc have had a huge influence on me intellectually. I count myself especially fortunate to have been taught a course on linear algebra by Prof. Dilip Patil, who imparted to us not only the course-specific knowledge, but also an appreciation of what rigorous mathematical thinking entails.

I would like to thank the Ministry of Human Resource Development (Government of India) for their financial support in the form of a monthly stipend for the last 2.5 years.

Finally — and this goes without saying, but I will still say it because someday they may get curious and try to read this document and this is probably the only page that would make any sense to them — I will forever owe a debt of gratitude to my parents, who, in addition to manufacturing me, have always given me their unconditional love and support.

Publications

Journal:

1. A.V. Singh and K.N. Chaudhury, "On Uniquely Registrable Networks," *IEEE Transactions on Network Science and Engineering*, accepted for publication.
2. A.V. Singh and K.N. Chaudhury, "Convergence Analysis of Nonconvex ADMM for Rigid Registration," *Computational Optimization and Applications*, under review.

Conference:

1. R. Sanyal, A.V. Singh and K.N. Chaudhury, "An Iterative Eigensolver for Rank-Constrained Semidefinite Programming," *Proc. 2019 National Conference on Communications (NCC)*, pp. 1-6, 2019.
2. A.V. Singh and K.N. Chaudhury, "When Can a System of Subnetworks be Registered Uniquely?," *Proc. 2019 IEEE International Conference on Acoustic, Speech, and Signal Processing (ICASSP)*, pp. 4564-4568, 2019.

Contents

List of Figures

List of Tables

Notations and Abbreviations

Notations:

$[m : n]$	:	Set of integers $\{m, \ldots, n\}$.
$\mathbb{R}^{n\times n}$	:	Set of $n \times n$ real-valued matrices.
$\mathbb{S}^n$	:	Set of $n \times n$ symmetric matrices.
$\mathbb{S}^n_+$	:	Set of $n \times n$ positive semidefinite matrices.
$\lVert\cdot\rVert$	:	Norm of a vector (Euclidean), or of a matrix (Frobenius).
$\lVert\cdot\rVert_2$	:	Spectral norm of a matrix.
$\mathrm{Tr}(\cdot)$	:	Trace of a matrix.
$\Pi_\Xi(\cdot)$	:	Euclidean projection on the closed set Ξ.
$\langle\cdot,\cdot\rangle$	:	Inner-product between matrices.
$[\cdot]_{i,j}$	:	(i, j)-th $d \times d$ block of an $Md \times Md$ matrix.
$\lambda_i(\cdot)$	:	i-th eigenvalue (in non-decreasing order) of a matrix.

Abbreviations:

ADMM	:	Alternating Direction Method of Multipliers.

Introduction

Registration problems arise in situations where we wish to reconstruct an underlying global structure, given access to multiple local snapshots of that structure, such as in sensor network localization and multiview registration. Suppose there are N labelled points with unknown coordinates in a d-dimensional Euclidean space, and there are M observers each of whom observes a subset of points. Each observer assigns coordinates to the points they observe in their local frame of reference. The *registration problem* is to infer the global coordinates of the N points given the following: (i) the points observed by each observer, (ii) the local coordinates assigned to the observed points for every observer, (iii) the kind of transform that relates the local frame of reference of each observer to the global frame of reference.

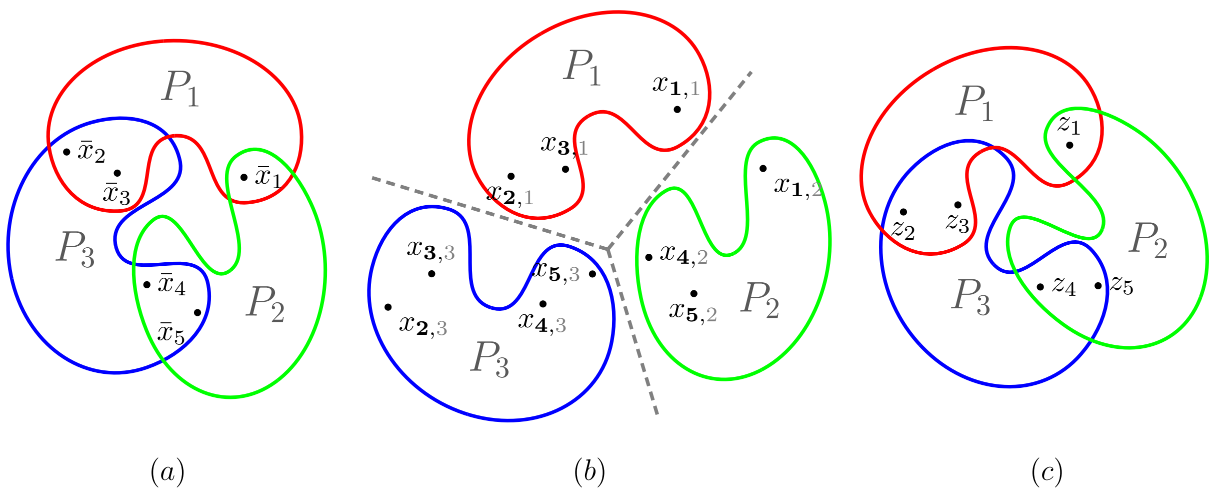

Figure 1.1: A rigid registration scenario. There are $N = 5$ labelled points, and $M = 3$ observers. The set of points observed by observer i is denoted by P_i; here $P_1 = \{1, 2, 3\}$, $P_2 = \{1, 4, 5\}$, $P_3 = \{2, 3, 4, 5\}$. (a) depicts the ground-truth, where $\bar{x}_i$ denotes the ground-truth global coordinates of point i. (b) depicts the three local coordinate systems corresponding to the observers; $x_{k,i}$ denotes the local coordinates of the k-th point assigned by the i-th observer (based on this information, we wish to recover the ground-truth global coordinates of every point). (c) A solution to the problem. Note that (a) and (c) are not exactly the same, but are related via a rigid transform, which is the best we can hope to do with the given information.

The *rigid* registration problem is the registration problem where the local frames of reference are related to the global frame of reference by unknown rigid (Euclidean) transforms (see Fig. 1.1). A rigid or Euclidean transform in $\mathbb{R}^d$ is a surjective distance-preserving transform. Any rigid transform can be seen as a composition of an orthogonal transform (rotation, reflection) and a translation, and is denoted by $(\mathbf{O}, \mathbf{t})$, where $\mathbf{O}$ denotes the orthogonal transform, and $\mathbf{t}$ denotes the translation. Specifically, rigid transform $(\mathbf{O}, \mathbf{t})$ maps a point $\mathbf{x}$ to the point $\mathbf{Ox} + \mathbf{t}$.

1.1 Application: Sensor Network Localization

As a concrete application of rigid registration, consider an adhoc wireless network consisting of geographically distributed sensor nodes with limited radio range. To make sense of the data collected from the sensors, one usually requires the positions of the individual sensors. It is often not feasible to equip each sensor with a GPS due to cost, power, and weight considerations. On the other hand, we can estimate (e.g. using time-of-arrival) the distances between sensor that are within the radio range of each other [2]. The problem of estimating the sensor locations from the available inter-sensor distances is referred to as *sensor network localization (SNL)* [2, 3]. Recently, scalable divide-and-conquer approaches for SNL were

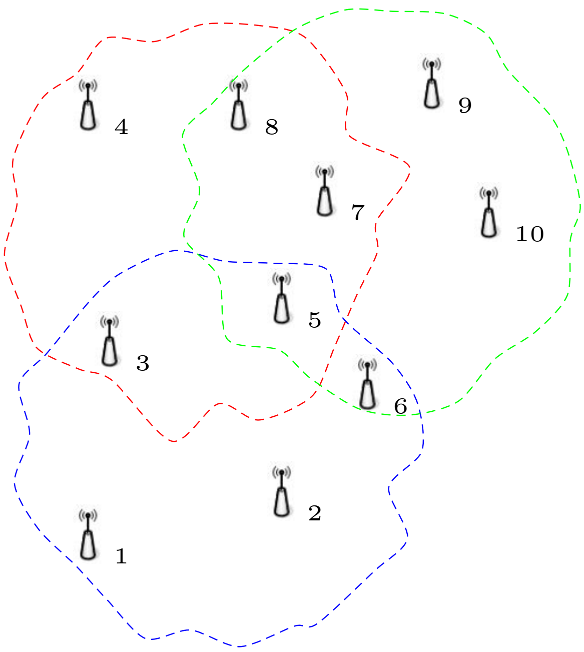

Figure 1.2: Sensor network localization scenario. The sensors within the radio range of each other are identified with a dashed closed curve (patch). For instance, sensor nodes 1, 2, 3, 5, 6 in the blue patch can communicate among themselves and thus all the inter-sensor distances among these sensors are known. Similarly, inter-sensor distances for sensors in the red patch, and inter-sensor distances for sensors in the green patch are known. Based on this information, we wish to infer the global structure of the sensor network. This, in particular, entails inferring distances between sensors that do not belong to the same patch (e.g., sensors 1 and 9).

proposed in [4, 5], where the network is first subdivided into smaller subnetworks that can be efficiently localized. The localized subnetworks are then *registered* in a global coordinate system to obtain the positions of the nodes in the original network.

More specifically (see Fig. 1.2), suppose the sensor network contains N nodes. This network is divided into M subnetworks, where each subnetwork consists of sensors that are within the radio range of each other (following [5], we call each subnetwork a *patch*). Thus, the distance between any two sensor nodes in a patch are known. Now, for each patch, a computationally efficient algorithm called classical multidimensional scaling (cMDS) [3, 6] is used to compute coordinates of nodes in that patch. However, since this computation is based solely on the inter-sensor distances, the coordinates returned by cMDS for a patch will in general be an arbitrarily rotated, flipped, and translated version of the ground-truth coordinates. The network is thus divided into M patches, where each patch can be regarded as constituting a local coordinate system which is related to the global coordinate system by an unknown rigid transform. We now want to assign coordinates to the nodes in a global coordinate system based on these patch-specific local coordinates. In other words, we want to solve the rigid registration problem. Since the local coordinate systems are related to the ground-truth coordinate system via unknown rigid transforms, solving this problem involves estimating these rigid transforms and "undoing" them to obtain global coordinates of the nodes (Fig. 1.1). Similar registration problems arise in other applications such as manifold learning, computer vision, and protein structure determination [7, 4, 5, 8, 9, 10, 11, 12, 13].

1.2 Problem Statement

To better facilitate discussion of our contribution in this dissertation, we formally describe the rigid registration problem. Suppose that we have N points in $\mathbb{R}^d$, which we label using[1] $\mathcal{S} = [1 : N]$. Let $P_1, \cdots, P_M$ be subsets of $\mathcal{S}$, where P_i is the subset of points observed by observer i. We refer to each P_i as a *patch* and let $\mathcal{P} = \{P_1, \cdots, P_M\}$ be the collection of patches. A natural way to represent the point-patch correspondence is using the bipartite graph $\Gamma_C = (\mathcal{S}, \mathcal{P}, \mathcal{E})$, where $(k, P_i) \in \mathcal{E}$ if $k \in P_i$; with a slight abuse of notation, we use (k, i) in place of (k, P_i). We refer to Γ_C as the *correspondence graph*.

We associate with each patch a local coordinate system: if $(k, i) \in \mathcal{E}$, let $\mathbf{x}_{k,i} \in \mathbb{R}^d$ be the local coordinates of point k in patch P_i. Suppose the local coordinate measurements are clean, i.e., we do not incur any noise in the measurement of the local coordinates. Let $\bar{\mathbf{x}}_1, \ldots, \bar{\mathbf{x}}_N \in \mathbb{R}^d$ be the (unknown) true global coordinates of the N points, and let $(\bar{\mathbf{O}}_i, \bar{\mathbf{t}}_i)$ denote the (unknown) true rigid transform that relates local coordinates of points in P_i to their corresponding global coordinates. That is, for each $k \in P_i$, we have

$$\bar{\mathbf{x}}_k = \bar{\mathbf{O}}_i \mathbf{x}_{k,i} + \bar{\mathbf{t}}_i. \tag{1.1}$$

[1] we use $[m : n]$ to denote the set of integers $\{m, \ldots, n\}$.

Thus, the rigid registration problem (when the local coordinate measurements are clean) becomes:

Given a correspondence graph $\Gamma_C = (\mathcal{S}, \mathcal{P}, \mathcal{E})$ *and local coordinates* $\{\mathbf{x}_{k,i} : (k,i) \in \mathcal{E}\}$, *find an* N*-tuple of global coordinates* $\mathbf{Z} = (\mathbf{z}_k)_{k=1}^N$, *and an* M*-tuple of patch transforms* $\mathfrak{R} = ((\mathbf{O}_i, \mathbf{t}_i))_{i=1}^M$, *such that*

$$\text{(REG)} \qquad \mathbf{z}_k = \mathbf{O}_i \mathbf{x}_{k,i} + \mathbf{t}_i, \qquad (k,i) \in \mathcal{E}.$$

In general, the local coordinate measurements, $\mathbf{x}_{k,i}$ for $(k,i) \in \mathcal{E}$ are noisy. In this case, we cannot expect the system of equations REG to hold exactly. Instead, we resort to least-squares fitting to solve the registration problem. More precisely, we consider the following least-squares minimization over $((\mathbf{O}_i, \mathbf{t}_i))_{i=1}^M$ and $(\mathbf{z}_k)_{k=1}^N$ [14] to estimate the rigid transforms and the global coordinates:

$$\text{(LS-REG)} \qquad \min_{(\mathbf{O}_i),(\mathbf{t}_i),(\mathbf{z}_k)} \sum_{i=1}^{M} \sum_{k \in P_i} \|\mathbf{z}_k - \big(\mathbf{O}_i \mathbf{x}_{k,i} + \mathbf{t}_i\big)\|^2.$$

1.3 Main Contributions

We now give an overview of the main results we prove in this dissertation. Each subsection here corresponds to a chapter in the dissertation, where the subsection title indicates the relevant chapter number.

1.3.1 Chapter 2: On Unique Registrability

Suppose the given local coordinate measurements of the rigid registration problem are clean. In this case, the true global coordinates $(\bar{\mathbf{x}}_k)_{k=1}^N$ and the patch transforms $(\bar{\mathbf{O}}_i, \bar{\mathbf{t}}_i)$ satisfy REG (see equation (1.1)). But is this solution unique? Can two different assignments of global coordinates (not related via a rigid transform) lead to the same observed local coordinates? This is a fundamental question one would be faced with when coming up with algorithmic solutions to the registration problem [14, 5]. By uniqueness, we mean uniqueness *up to congruence*, i.e., any two solutions that are related through a global rigid transform are considered identical. For instance, in Fig. 1.1, solutions (a) and (c) are considered identical since they are related by a global rotation (loosely speaking, they have the same *shape*).

In Chapter 2, using results from graph rigidity theory, we give a necessary and sufficient condition for REG to have a unique solution. Namely, given an instance of REG, we construct a graph (called *body graph*) corresponding to the given problem, and check if this graph is *globally rigid* to determine if REG has a unique solution. We now briefly describe the concept of global rigidity of a graph and what we mean by a body graph, before stating an informal version of our result.

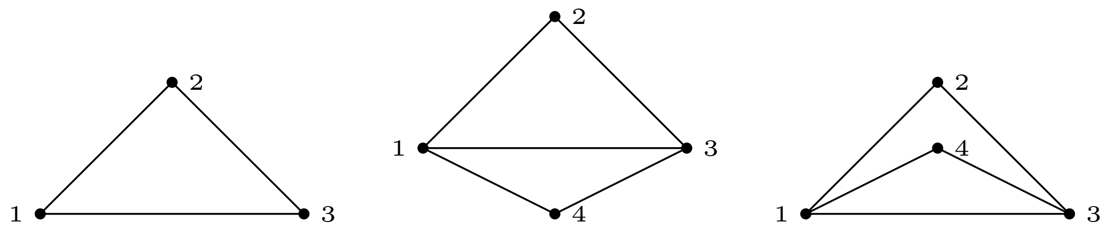

(a) (b) (c)

Figure 1.3: The triangular graph in (*a*) is globally rigid in $\mathbb{R}^2$ because any other edge-length-preserving embedding of this graph in $\mathbb{R}^2$ would be congruent to the embedding in (*a*). On the other hand, (*b*) and (*c*) are two embeddings of the same underlying graph. Note that the embeddings have the same edge lengths, but are not congruent; thus, this graph is not globally rigid in $\mathbb{R}^2$. We make these notions more precise in Chapter 2.

Consider a graph embedded in Euclidean space. The notion of global rigidity deals with the following question: Is there a non-congruent alternate embedding of the graph with the same edge lengths? Informally, is there an embedding that preserves the edge lengths but has a different shape? If the answer to this question is in the negative, the graph is said to be globally rigid. For instance, the triangular graph in Fig. 1.3a is globally rigid in $\mathbb{R}^2$. On the other hand, the graph in Fig. 1.3b is not globally rigid in $\mathbb{R}^2$ because Fig. 1.3c is an alternate embedding of the same graph which preserves all the edge lengths but has a different shape.

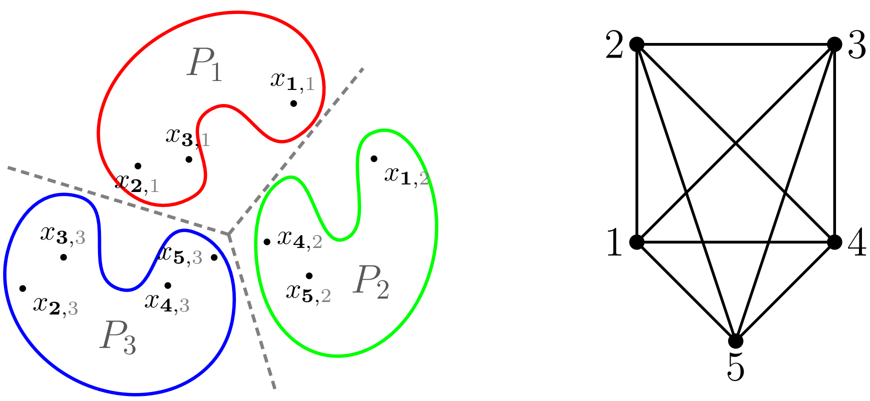

Figure 1.4: A problem instance of REG and its corresponding body graph. An edge between vertices *i* and *j* in the body graph corresponds to the fact that there is a patch that contains both point *i* and point *j*.

We now clarify what we mean by a *body graph*. Given a problem instance of REG, the corresponding body graph is a graph with the number of vertices equal to the number of points in the problem instance: each vertex corresponds to a point, and there is an edge

between vertices i and j of the body graph, if there exists a patch that contains both point i and point j (see Fig. 1.4).
Informally, our main results on the uniqueness of solution to REG are as follows.

Theorem (Informal)**.** *Under mild assumptions, REG has a unique solution if and only if the corresponding body graph is globally rigid.*

For rigid registration in 2-dimensional space, we obtain a particularly simple characterization for uniqueness of solution.

Corollary (Informal)**.** *Under mild assumptions, REG in* $\mathbb{R}^2$ *has a unique solution if and only if the corresponding body graph is* 3*-connected.*

Furthermore, using these results, we resolve a conjecture on uniqueness posed in [5], which was stated in terms of *quasi-connectedness* of the correspondence graph. To do this, we relate quasi-connectedness of the corresondence graph to the standard vertex-connectivity of the body graph, and show that the conjecture holds in 2-dimensional space, and that it fails to hold in spaces of dimension 3 and higher.

1.3.2 Chapter 3: Tightness of Convex Relaxation

Recall that when the local coordinate measurements in the registration problem are noisy, we use least-squares minimization LS-REG to solve the registration problem. As we shall see in Chapter 3, LS-REG essentially reduces to the following rank-contrained semidefinite program:

$$\text{(REG-SDP)} \qquad \begin{aligned} \min_{\mathbf{G}\in\mathbb{S}_+^{Md}} \quad & \mathrm{Tr}\,(\mathbf{CG}) \\ \text{subject to} \quad & [\mathbf{G}]_{ii} = \mathbf{I}_d, \quad i \in [1:M], \\ & \mathrm{rank}(\mathbf{G}) = d, \end{aligned}$$

where $\mathbb{S}_+^{Md}$ denotes the cone of symmetric $Md \times Md$ positive semidefinite matrices, $[\mathbf{G}]_{ii}$ denotes the i-th $d \times d$ diagonal block of the $Md \times Md$ matrix $\mathbf{G}$, and $\mathbf{I}_d$ denotes the $d \times d$ identity matrix. This is a standard semidefinite program (SDP), but with an additional rank constraint which makes it nonconvex, and computationally hard (for $d = 1$, this problem class includes MAXCUT [15]). However, by dropping the rank constraint, we obtain the following computationally tractable convex relaxation of REG-SDP:

$$\text{(C-SDP)} \qquad \begin{aligned} \min_{\mathbf{G}\in\mathbb{S}_+^{Md}} \quad & \mathrm{Tr}(\mathbf{CG}) \\ \text{subject to} \quad & [\mathbf{G}]_{ii} = \mathbf{I}_d, \quad i \in [1:M]. \end{aligned}$$

Suppose $\mathbf{G}^*$ is global optimum for the relaxed problem C-SDP. The rank of $\mathbf{G}^*$ is not guaranteed to be d, and so $\mathbf{G}^*$ might not even be feasible for REG-SDP. Thus, if $\mathrm{rank}(\mathbf{G}^*) > d$

(it can never be less than d due to the constraint $[\mathbf{G}^*]_{ii} = \mathbf{I}_d$), we have to "round" $\mathbf{G}^*$ to a rank-d matrix [14], which will generally be suboptimal for REG-SDP. However, if $\text{rank}(\mathbf{G}^*) = d$, then clearly $\mathbf{G}^*$ is global optimum for REG-SDP as well, meaning that we have solved the original nonconvex problem by solving the relaxed problem. In this case, we say that the relaxation is *tight* [1]. Empirically, we notice the well-known phenomena of phase transition for convex relaxations (e.g. see [16, 17]), where below a certain noise threshold, the relaxation remains tight (see Fig. 1.5).

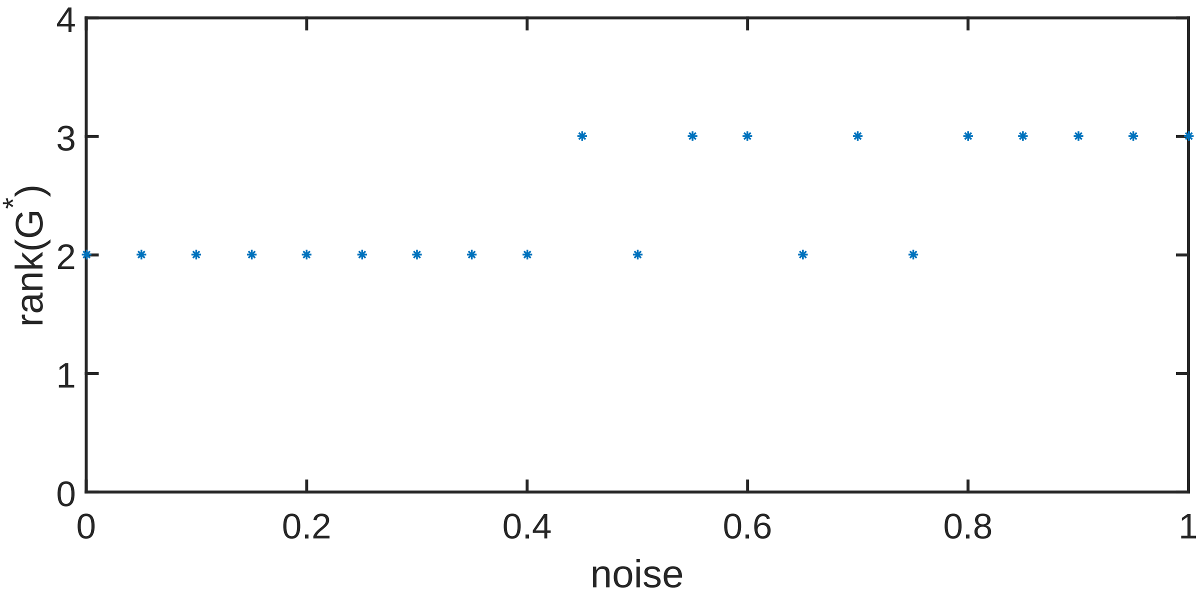

Figure 1.5: Phase transition in tightness of C-SDP ($d = 2$). The plot shows the rank of global optimum $\mathbf{G}^*$ of C-SDP as a function of noise level in the measurements. Below a certain noise threshold, the rank of $\mathbf{G}^*$ is exactly 2, i.e., the relaxation is tight. Above this threshold, the rank of $\mathbf{G}^*$ may exceed 2, making it infeasible for REG-SDP.

In Chapter 3, we give a theoretical justification of this phenomenon by analyzing the Lagrange dual of C-SDP. Our analysis is in the spirit of the analysis in [1], where the authors use Lagrange dual to study tightness of convex relaxation in the context of the phase synchronization problem. Our main result is the following.

Theorem (Informal)**.** *Let* $\mathbf{C}_0$ *be the data matrix when the local coordinate measurements are clean, and let* $\mathbf{C} = \mathbf{C}_0 + \mathbf{W}$, *where* $\mathbf{W}$ *is the noise matrix. Then there exists* $\eta > 0$ *such that* C-SDP *is tight if* $\|\mathbf{W}\| < \eta$.

1.3.3 Chapter 4: Convergence of Nonconvex ADMM

As discussed in the previous subsection, solving the convex relaxation C-SDP might produce a solution infeasible for the original problem REG-SDP, in which case we would need to resort to a suboptimal rounding step. This motivates the following question: Can we design an iterative procedure that directly attacks the nonconvex problem REG-SDP instead of working with the relaxation C-SDP? In [9], the authors proposed an iterative procedure based on the alternating direction method of multipliers (ADMM) [18] to solve REG-SDP, and demonstrated empirically that it performs remarkably well in the context of multiview registration. As we shall see in Chapter 4, in addition to obviating the need for any rounding, the algorithm in [9] (henceforth referred to as REG-ADMM) provides appreciable savings in computation over algorithms that solve the convex semidefinite program C-SDP.

ADMM algorithms applied to *convex* programs are well-studied, as exemplified by a much-cited paper by Boyd et al [18]. Recently, however, there have been a slew of results showing empirical success of ADMM algorithm applied to *nonconvex* programs [9, 19, 20, 21]. The theory for nonconvex ADMM, though, is yet to catch-up to these empirical results. A handful of theoretical analyses on nonconvex ADMM that do exist do not apply to REG-ADMM because they rely on certain regularity assumptions that are not satisfied by REG-SDP. In Chapter 4, we remedy this by bypassing these assumptions to prove convergence results for the REG-ADMM algorithm.

We briefly describe the kind of challenge involved in characterizing convergence behavior of REG-ADMM. Standard ADMM algorithm involves a fixed parameter ρ which affects the update of primal and dual variables by the algorithm. For ADMM applied to convex programs, the algorithm (under mild assumptions) can be shown to converge to global minimum for *any* ρ, and with *any* primal and dual initialization [18]. On the other hand (see Fig. 1.6), convergence of REG-ADMM depends not only on the parameter ρ, but also on the noise level in the data measurements. Even when the local coordinates measurements are clean, REG-ADMM may get stuck in a local minimum, depending on the value of parameter ρ. This dependence of the limit point of REG-ADMM iterates on ρ is observed both at low and high noise levels. Furthermore, when the noise in the data is relatively large (Fig. 1.6c), the iterates of the algorithm may oscillate without converging if the parameter ρ is set to a small value. What is perhaps surprising is that such non-attenuating oscillations for small values of ρ are not observed when the noise level in the data is low (Figure 1.6b).

Our analysis throws light on the convergence behaviour of REG-ADMM observed at different levels of noise in the data. Following are informal versions of the main results we prove in Chapter 4.

Theorem (Informal)**.** *If REG-ADMM converges, it does so to a stationary point of REG-SDP.*

Theorem (Informal)**.** *Suppose noise in the data is below a certain threshold. Then, the REG-ADMM iterates are attracted to the global minimum, if the initialization is in a close neighborhood of the global minimum.*

Theorem (Informal)**.** *Suppose the data is clean. Then given any primal initialization, we can compute ρ_0, such that REG-ADMM converges to global minimum for $\rho \leq \rho_0$.*

Further, using duality, we explain instability of REG-ADMM observed for high-noise data when the parameter ρ is small.

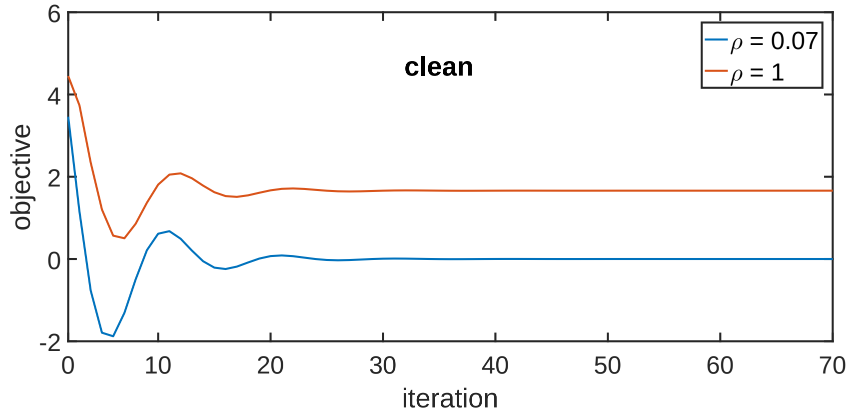

(a) Even with clean measurements, the iterates get stuck in a local minimum when $\rho = 1$ (the optimum value is zero in this case).

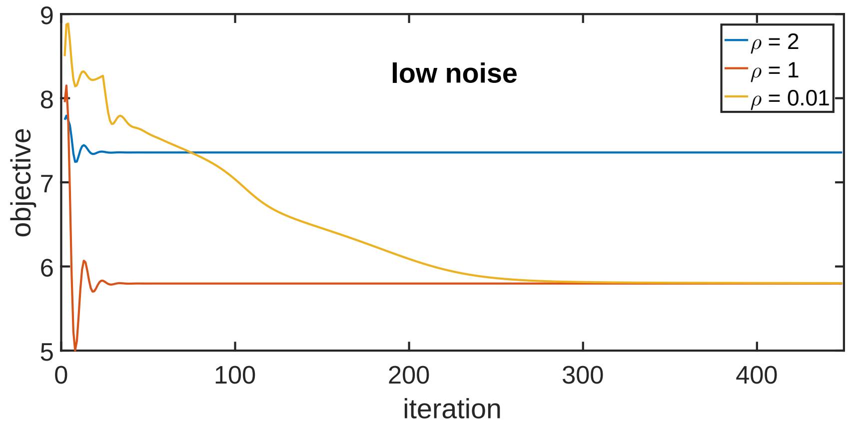

(b) There are no oscillations for small ρ when the noise is low.

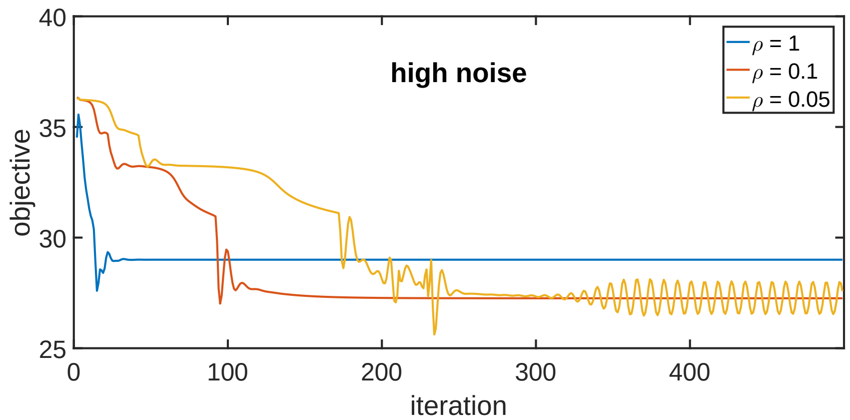

(c) The iterates oscillate for small ρ when the noise is high.

Figure 1.6: Plot of REG-ADMM iterates for different noise levels in the data measurements.

We conclude this dissertation with Chapter 5, where we review our results and suggest future directions to further strengthen these results.

Chapter 2

On Unique Registrability

2.1 Introduction

In this chapter, we investigate the notion of uniqueness of solution to a rigid registration problem. Since the sensor network localization problem discussed in Chapter 1 (Section 1.1) is a canonical example of the rigid registration problem, we shall use the term *network* to denote a collection of points in a Euclidean space, and the term *node* to refer to a point in the network. Consider a d-dimensional network consisting of N labelled nodes, i.e. the nodes are embedded in $\mathbb{R}^d$. There are M subsets of these nodes (*patches*). Each patch has associated with it a local coordinate system, i.e., nodes in a patch are assigned coordinates in a local coordinate system specific to that particular patch (*patch coordinate system*). The patch coordinate systems are related to each other by unknown rigid transforms. Equivalently, each patch coordinate system is related to the global coordinate system by an unknown rigid transform. Our problem is to assign coordinates to every node in a global coordinate system. In such problems, a question that naturally arises is that of uniqueness: Can we uniquely identify the global topology of the network that is consistent with the information in various local coordinate systems? (We call the network *uniquely registrable* if this is indeed the case.) Additionally, do we have computationally efficient tests to determine if the network is uniquely registrable? In this chapter, we investigate these questions using results from graph rigidity theory.

More precisely, we formulate a necessary and sufficient condition for a network to be uniquely registrable in terms of rigidity of the *body graph* of the network. This leads to a particularly simple characterization of unique registrability for planar networks. As an upshot, we resolve a conjecture posed recently [5] in the context of unique localizability of sensor networks. The conjecture was made in terms of *quasi-connectivity* of a certain bipartite graph; we translate this notion to that of vertex-connectivity of the body graph, and show that the conjecture holds for planar sensor networks, and that it fails to hold for three and higher dimensional networks.

2.1.1 Uniqueness in Rigid Registration

We begin by making precise the notion of uniqueness in the context of the rigid registration problem. We first review the formal problem statement of the rigid registration problem.

The node-patch correspondence (i.e. which nodes belong to which patch) is encoded by a bipartite correpondence graph $\Gamma_C = (\mathcal{S}, \mathcal{P}, \mathcal{E})$, where $\mathcal{S} = [1 : N]$, $\mathcal{P} = [1 : M]$, and $(k, i) \in \mathcal{E}$ if and only if node k belongs to patch P_i. We associate with each patch a local coordinate system: if $(k, i) \in \mathcal{E}$, let $x_{k,i} \in \mathbb{R}^d$ be the local coordinates assigned to node k. Suppose the local coordinate measurements are *exact*, i.e., the local coordinates given to us are noiseless. Then, the rigid registration problem is as follows:

Given the correspondence graph $\Gamma_C = (\mathcal{S}, \mathcal{P}, \mathcal{E})$ *and the local coordinates* $\{x_{k,i} : (k, i) \in \mathcal{E}\}$*, find an N-tuple of global coordinates* $\mathbf{Z} = (z_k)_{k=1}^N$*, and an M-tuple of patch transforms* $\mathfrak{R} = (\mathcal{R}_i)_{i=1}^M$*, such that*

$$\text{(REG)} \qquad z_k = \mathcal{R}_i(x_{k,i}), \qquad (k, i) \in \mathcal{E}.$$

Suppose the true coordinates of the N nodes in the network are $\bar{x}_1, \cdots, \bar{x}_N$, and the patch coordinate system P_i is related to the global coordinate system by the rigid transform $\bar{\mathcal{R}}_i$, $i \in [1 : M]$. Then, clearly, the true global coordinates $(\bar{x}_k)_{k=1}^N$ and the patch transforms $(\bar{\mathcal{R}}_i)_{i=1}^M$ satisfy REG. But is this solution unique? This is a fundamental question one would be faced with when coming up with algorithmic solutions to the registration problem [14, 5]. Of course, by uniqueness, we mean uniqueness up to *congruence*, i.e., any two solutions that are related through a Euclidean transform are considered identical. Note that a solution to REG has two components: the global coordinates, and the patch transforms. We will define uniqueness for each of these components. Suppose $(\mathbf{X}, \mathfrak{R})$ is a solution to (2). By *uniqueness of global coordinates*, we mean that given any other solution $(\mathbf{Y}, \mathfrak{T})$ to REG, there exists a Euclidean transform $\mathcal{Q}$ such that $y_k = \mathcal{Q}(x_k)$, $k \in \mathcal{S}$. Similarly, by *uniqueness of patch transforms*, we mean that there exists a Euclidean transform $\mathcal{U}$ such that $\mathcal{T}_i = \mathcal{U} \circ \mathcal{R}_i$, $i \in [1 : M]$, where $\circ$ denotes the composition of transforms. At this point, we make the following observation.

Observation 1. *It is clear that uniqueness of patch transforms implies uniqueness of global coordinates. That is, given two solutions* $(\mathbf{X}, \mathfrak{R})$ *and* $(\mathbf{Y}, \mathfrak{T})$ *to REG, if there exists a Euclidean transform* $\mathcal{U}$*, such that* $\mathcal{T}_i = \mathcal{U} \circ \mathcal{R}_i, i \in [1 : M]$*, then there exists a Euclidean transform* $\mathcal{Q}$*, such that* $y_k = \mathcal{Q}(x_k), k \in \mathcal{S}$ *(in particular, take* $\mathcal{Q} = \mathcal{U}$*). However, uniqueness of global coordinates does not imply uniqueness of patch transforms. That is, given two solutions* $(\mathbf{X}, \mathfrak{R})$ *and* $(\mathbf{Y}, \mathfrak{T})$ *to REG, there may not exist a Euclidean transform* $\mathcal{U}$*, such that* $\mathcal{T}_i = \mathcal{U} \circ \mathcal{R}_i, i \in [1 : M]$*, even if there exists a Euclidean transform* $\mathcal{Q}$*, such that* $y_k = \mathcal{Q}(x_k), k \in \mathcal{S}$*. (This is explained with an example in Fig. 2.1.)*

Notice that each patch has just two nodes in the example in Fig. 2.1. However, we know that a Euclidean transform in $\mathbb{R}^d$ is completely specified by its action on a set of $d + 1$ non-

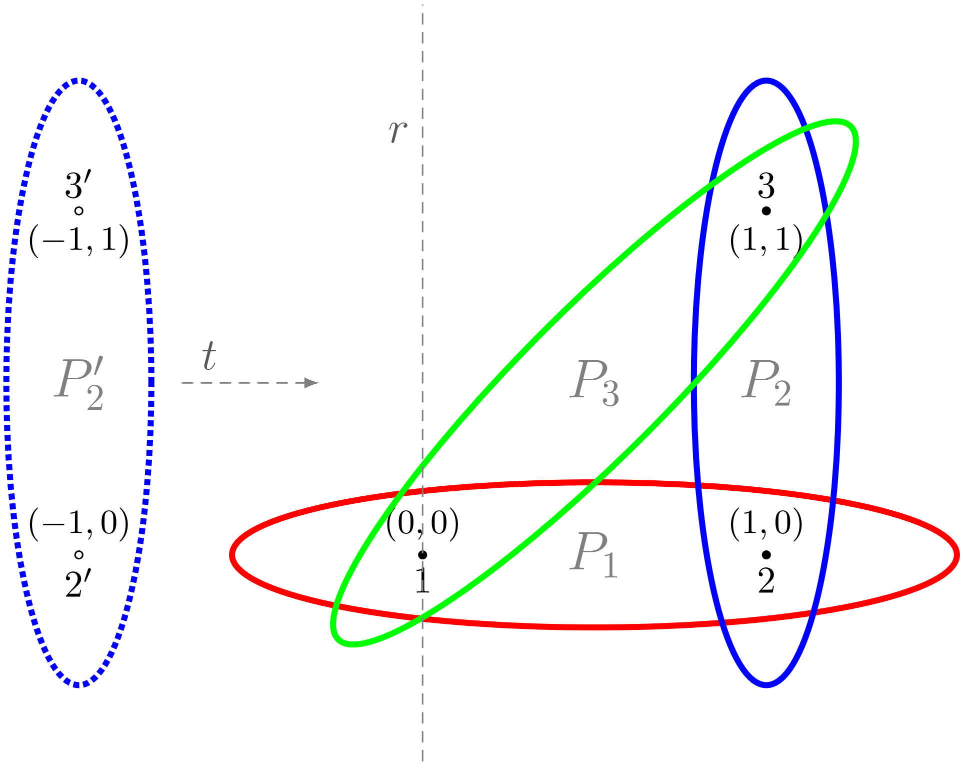

Figure 2.1: Consider the nodes $\mathcal{S} = \{1, 2, 3\}$, and the patches $\mathcal{P} = \{P_1, P_2, P_3\}$, where $P_1 = \{1, 2\}, P_2 = \{2, 3\}, P_3 = \{1, 3\}$. The true global coordinates are $\bar{\mathbf{X}} = ((0, 0), (1, 0), (1, 1))$, and the true patch transforms are $\bar{\mathfrak{R}} = (\mathcal{I}_d, \mathcal{I}_d, \mathcal{I}_d)$, where $\mathcal{I}_d$ is the identity transform (i.e., each patch coordinate system is same as the global coordinate system). Consider the Euclidean transform $\mathcal{T}$, which is a reflection along the dotted line marked r, followed by a translation of 2 units along the dotted ray marked t. Let $\mathfrak{R} = (\mathcal{I}_d, \mathcal{T}, \mathcal{I}_d)$. Notice that even though both $(\bar{\mathbf{X}}, \bar{\mathfrak{R}})$ and $(\bar{\mathbf{X}}, \mathfrak{R})$ are solutions to REG, $\mathfrak{R}$ is not congruent to $\bar{\mathfrak{R}}$.

degenerate nodes[1]. Equivalently, if $d + 1$ or more non-degenerate nodes are left fixed by a Euclidean transform, then the transform must be identity. This leads to the following proposition.

Proposition 2. *If every patch contains at least $d + 1$ non-degenerate nodes, then uniqueness of global coordinates is equivalent to uniqueness of patch transforms.*

Proof. Following Observation 1, we need only prove that uniqueness of global coordinates implies uniqueness of patch transforms. Suppose we have two solutions $(\mathbf{X}, \mathfrak{R})$ and $(\mathbf{Y}, \mathfrak{T})$. Following the uniqueness of global coordinates, there exists a Euclidean transform $\mathcal{Q}$, such that $y_k = \mathcal{Q}(x_k)$, $k \in \mathcal{S}$. Fix some $i \in [1 : M]$. Since $(\mathbf{Y}, \mathfrak{T})$ is a solution to REG, we have $y_k = \mathcal{T}_i(x_{k,i})$, $k \in P_i$. Thus, $\mathcal{Q}(x_k) = \mathcal{T}_i(x_{k,i})$, or $x_k = (\mathcal{Q}^{-1} \circ \mathcal{T}_i)(x_{k,i})$, $k \in P_i$. On the other hand, since $(\mathbf{X}, \mathfrak{R})$ is also a solution to REG, we have $x_k = \mathcal{R}_i(x_{k,i})$, $k \in P_i$. Combining the above, we get $(\mathcal{Q}^{-1} \circ \mathcal{T}_i)(x_{k,i}) = \mathcal{R}_i(x_{k,i})$, $k \in P_i$. Since $|P_i| \geq d + 1$, it follows that $\mathcal{Q}^{-1} \circ \mathcal{T}_i = \mathcal{R}_i$, or $\mathcal{T}_i = \mathcal{Q} \circ \mathcal{R}_i$. This holds for every $i \in [1 : M]$, which proves our claim. □

[1] A set of nodes in $\mathbb{R}^d$ is said to be *non-degenerate* if their affine span is $\mathbb{R}^d$.

In other words, if every patch contains at least $d + 1$ non-degenerate nodes, we need not distinguish between *uniqueness of global coordinates* and *uniqueness of patch transforms*, and we can generally talk about *unique registrability* without any ambiguity. Intuitively, it is clear that for REG to have a unique solution, there must be sufficient overlap among patches. In particular, Γ_C must be connected. In Section 2.3, we will see that the notion of unique registrability is essentially combinatorial in nature for *almost every* instance of the problem.

2.1.2 Related Work

The correspondence graph $\Gamma_C = (\mathcal{S}, \mathcal{P}, \mathcal{E})$ encodes the pattern of overlap among patches, which makes it desirable to relate the problem of unique registrability to the properties of Γ_C. In [14], the authors propose a lateration criterion which guarantees unique registrability. We recall that Γ_C is said to be *laterated* if there exists a reordering of the patch indices such that P_1 contains at least $d + 1$ non-degenerate nodes, and P_i and $P_1 \cup P_2 \cup \cdots \cup P_{i-1}$ have at least $d + 1$ non-degenerate nodes in common for $i \geq 2$. This criterion, however, has two major shortcomings. First, an efficient test for lateration is not known. Second, lateration is a rather strong condition. For instance, see Fig. 2.4, where Γ_C is not laterated, but, as we will see later, the network is uniquely registrable. More recently, the notion of *quasi connectedness* of Γ_C was introduced in [5], which was shown to be necessary for unique registrability, and conjectured to be sufficient.

In a related work [22], rigidity theory is used to deal with unique localizability of nodes in a general sensor network localization problem, where, given inter-node distances of a subset of node-pairs, a graph is constructed with the vertices corresponding to the nodes, and an edge between every node-pair whose inter-node distance is given; it is demonstrated that this graph has to be globally rigid for unique localizability of the sensor network. Tools from rigidity theory have also been used in network design problem [23], and in quantifying robustness of networks [24].

2.1.3 Organization

The rest of the chapter is organized as follows. In Section 2.2, we review relevant definitions and results from rigidity theory. In Section 2.3,we bring in the notion of *body graph*, introduced in [25] in the context of affine rigidity, and show that unique registrability is equivalent to global rigidity of the body graph. In Section 2.4, we address the conjecture posed in [5], namely that quasi $(d + 1)$-connectivity of Γ_C is necessary and sufficient for unique registrability in $\mathbb{R}^d$. We show that quasi connectivity of Γ_C is equivalent to vertex-connectivity of the body graph, and use this to establish the conjecture for $d = 2$. Next, we give counterexamples to show that the conjecture is false when $d \geq 3$. Detailed proofs of some of the technical results from Sections 2.3 and 2.4 are given in Section 2.6.

2.1.4 Graph Notations

We will work with undirected graphs in this chapter. If H is a subgraph of $G = (V, E)$, which we denote by $H \subseteq G$, then $V(H)$ denotes the set of vertices of H, and $E(H)$ denotes the set of edges of H. A complete graph (or clique) on n vertices is denoted by K_n. Given a graph $G = (V, E)$, and a set $V' \subseteq V$, the subgraph induced by V' is the graph $G' = (V', E')$, where $E' = \{(i, j) \in E : i, j \in V'\}$. The degree of a vertex v of a graph is the number of edges incident on v. A path in a graph $G = (V, E)$ is an ordered sequence of distinct vertices $v_1, \cdots, v_n \in V$ such that $(v_i, v_{i+1}) \in E, 1 \le i \le n - 1$. We denote a path by $v_1 - \cdots - v_n$; v_1 and v_n are called the end vertices of the path, and every other vertex of the path is an internal vertex. If $v_1 = a$ and $v_n = b$, we say that the path connects a and b, or that $v_1 - \cdots - v_n$ is a path between a and b. Given subgraphs A and B, an A-B path is a path $v_1 - \cdots - v_n$ where $v_1 \in V(A)$ and $v_n \in V(B)$. Given a subgraph A, a path $v_1 - \cdots - v_n$ is said to be within A, if $v_i \in V(A)$ for every $i \in [1 : N]$. Two paths are said to be disjoint if they do not have any vertex in common. Two paths are said to be independent if they do not have any internal vertex in common. A graph is said to be k-connected (or, k-vertex-connected) if it has more than k vertices and the subgraph obtained after removing fewer than k vertices remains connected; equivalently, by Menger's theorem [26], there exists k independent paths between every pair of vertices of the graph.

2.2 Rigidity Theory

Before moving on to our results, we recall some definitions and results from rigidity theory [27, 28, 29, 30, 31].

2.2.1 Basic Terminology

Given a graph $G = (V, E)$, a d-dimensional *configuration* is a map $\mathbf{p} : V \longrightarrow \mathbb{R}^d$. The pair $(G, \mathbf{p})$ is called a d-dimensional *framework*. Throughout this chapter, $\|\cdot\|$ denotes the Euclidean norm.

Definition 3 (Equivalent frameworks)**.** *Two frameworks* $(G, \mathbf{p})$ *and* $(G, \mathbf{q})$ *are said to be equivalent, denoted by* $(G, \mathbf{p}) \sim (G, \mathbf{q})$, *if* $\|\mathbf{p}(u) - \mathbf{p}(v)\| = \|\mathbf{q}(u) - \mathbf{q}(v)\|$, *for every* $(u, v) \in E$.

Definition 4 (Congruent frameworks)**.** *Two frameworks* $(G, \mathbf{p})$ *and* $(G, \mathbf{q})$ *are said to be congruent, denoted by* $(G, \mathbf{p}) \equiv (G, \mathbf{q})$, *if* $\|\mathbf{p}(u) - \mathbf{p}(v)\| = \|\mathbf{q}(u) - \mathbf{q}(v)\|$ *for every* $u, v \in V$.

Thus, for two frameworks to be congruent, we require the distance between *each* vertex-pair to be equal in the two frameworks (regardless of whether that vertex-pair corresponds to an edge in the underlying graph G). In other words, congruent frameworks are related through a rigid transform. Clearly, congruence implies equivalence, but the converse is generally not true (see Fig. 2.2). This leads to the concept of global rigidity.

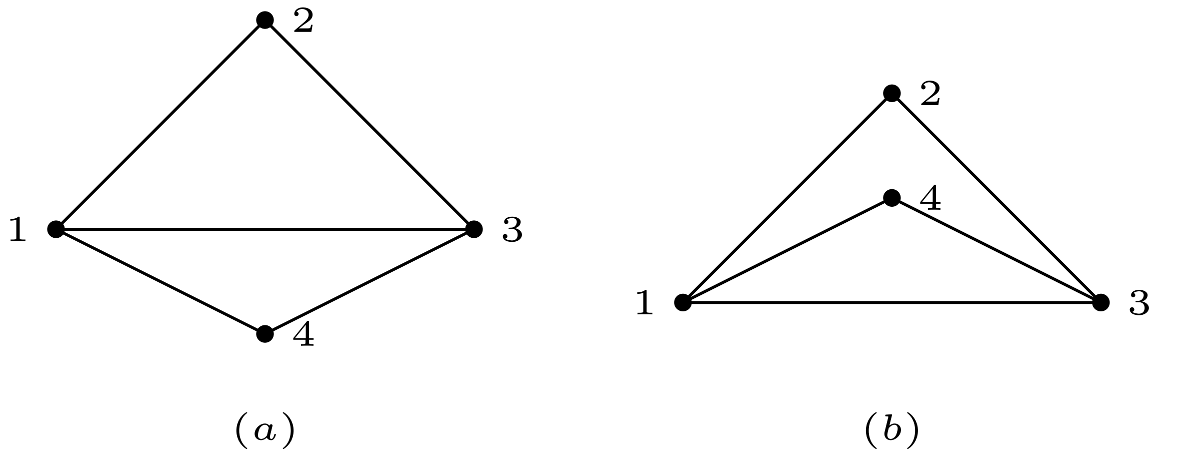

Figure 2.2: The frameworks in (a) and (b) are equivalent because the corresponding edge lengths are equal; however, they are not congruent because the distance between vertices 2 and 4 is not equal in the two frameworks. Thus, the framework in (a) is not globally rigid in $\mathbb{R}^2$. On the other hand, it can be shown that the framework is *locally* rigid in $\mathbb{R}^2$. Observe that there exists no continuous motion in $\mathbb{R}^2$ that takes (a) to (b). Also note that framework (a) is not locally rigid in $\mathbb{R}^3$ since the lower triangle 4-1-3 can be rotated in 3-dimensional space about the line 1-3 to get to framework (b), which is equivalent but non-congruent to framework (a).

Definition 5 (Globally rigid framework)**.** *A framework* $(G, \mathbf{p})$ *is said to be globally rigid if any framework equivalent to* $(G, \mathbf{p})$ *is also congruent to* $(G, \mathbf{p})$.

In other words, if $(G, \mathbf{q}) \sim (G, \mathbf{p})$, and $(G, \mathbf{p})$ is globally rigid, then $\mathbf{p}$ and $\mathbf{q}$ must be related via a rigid transform. We now define local rigidity of a framework, which formalizes the notion that the framework cannot be *continuously* deformed into an equivalent framework (see Fig. 2.2, where (a) and (b) are locally rigid in $\mathbb{R}^2$, since they cannot be continuously deformed while preserving edge lengths).

Definition 6 (Locally rigid framework)**.** *A framework* $(G, \mathbf{p})$ *is said to be locally rigid if there exists* $\epsilon > 0$ *such that any* $(G, \mathbf{q}) \sim (G, \mathbf{p})$ *satisfying* $\|\mathbf{p}(v) - \mathbf{q}(v)\| \leq \epsilon, v \in V$, *is congruent to* $(G, \mathbf{p})$.

2.2.2 Rigidity and Genericity

A fundamental problem in rigidity theory is the following: *Given a d-dimensional framework* $(G, \mathbf{p})$, *decide whether it is (locally or globally) rigid in* $\mathbb{R}^d$. The notions of local and global rigidity depend not only on the graph, but also on the configuration (see Fig. 2.3). This makes testing of rigidity computationally intractable [32, 33]. A standard way of getting around this is to make an additional assumption of *genericity*. A framework (or configuration) is said to be *generic* if there are no algebraic dependencies among the coordinates of the configuration, i.e., the coordinates of the configuration do not satisfy any non-trivial algebraic equation with rational coefficients. For a given graph, the set of non-generic configurations is a measure-zero set in the space of all possible configurations [34], and hence *almost every* configuration is generic.

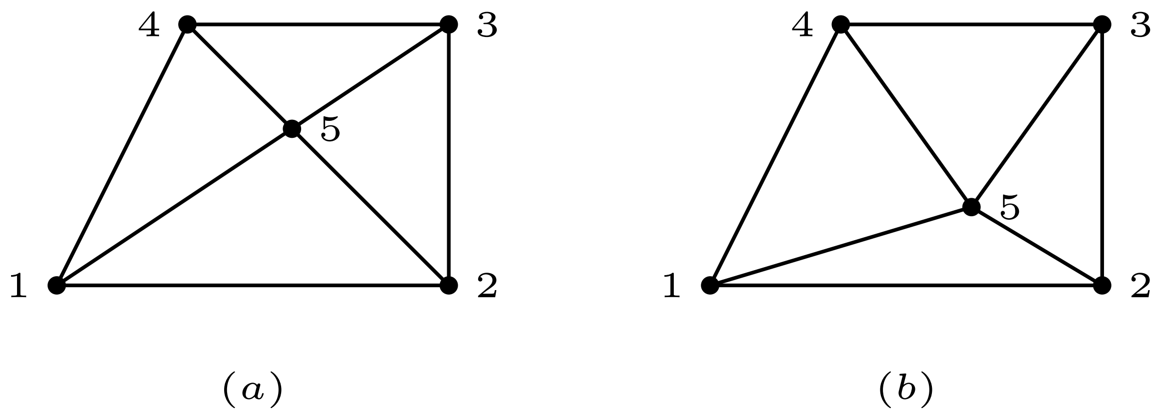

Figure 2.3: Frameworks (a) and (b) with the same underlying graph. Framework (a) is not globally rigid because vertex 4 can be reflected along the line 1-5-3, which results in an equivalent but non-congruent framework. Such an edge-length-preserving reflection is not possible in (b), which is globally rigid.

In the next two subsections, we will discuss how the genericity assumption makes local and global rigidity a property of the graph, independent of any particular configuration. This opens up the possibility of coming up with combinatorial characterizations of rigidity. In particular, we will look at combinatorial characterizations of local and global rigidity in $\mathbb{R}^2$. In the final subsection, we will briefly discuss the algorithmic implications of genericity.

2.2.3 Generic Local Rigidity

We introduce the notion of *infinitesimal rigidity* of a framework $(G, \mathbf{p})$, which, in general, is a stronger notion than the local rigidity of $(G, \mathbf{p})$, but is equivalent to local rigidity if $\mathbf{p}$ is generic [28]. Loosely speaking, given a graph $G = (V, E)$ and a framework $(G, \mathbf{p})$, we look at motions that the vertices of G could make, starting with configuration $\mathbf{p}$, such that the edge lengths are preserved at the instant of starting out. If every such motion also *instantaneously* preserves the distance between vertex pairs that do not form an edge, we say that $(G, \mathbf{p})$ is infinitesimally rigid. More precisely, let $\mathbf{v}$ be the map that assigns the motion vector $\mathbf{v}(i)$ to vertex $i \in V$. For the edge lengths of the framework $(G, \mathbf{p})$ to be preserved instantaneously, the net motion of an edge should have no component along the edge itself, which can be expressed as

$$(\mathbf{p}(i) - \mathbf{p}(j))^\top (\mathbf{v}(i) - \mathbf{v}(j)) = 0, \quad (i, j) \in E. \tag{2.1}$$

Writing (2.1) in the matrix form, we get

$$R(G, \mathbf{p})\mathbf{v} = 0, \tag{2.2}$$

where $R(G, \mathbf{p})$ is called the *rigidity matrix* of $(G, \mathbf{p})$. The matrix $R(G, \mathbf{p})$ has $|E|$ rows and $d|V|$ columns, where each row corresponds to an edge of G and d consecutive columns corresponds to a vertex (columns indexed $(i-1)d : id$ correspond to vertex i). The matrix is populated as follows: if $e = (i, j) \in E$, then for the row corresponding to e, we fill the d columns corresponding to vertex i with the elements of $\mathbf{p}(i) - \mathbf{p}(j)$; similarly, we fill the d columns corresponding to vertex j with the elements of $\mathbf{p}(j) - \mathbf{p}(i)$, and the remaining columns in that row are set to 0. For instance, if the framework $(G, \mathbf{p})$ is as in Fig. 2.2(a), and if we let $\mathbf{p}(i) = (x_i, y_i) \in \mathbb{R}^2$, then $R(G, \mathbf{p})$ is given by

$$\begin{bmatrix} x_1 - x_2 & y_1 - y_2 & x_2 - x_1 & y_2 - y_1 & 0 & 0 & 0 & 0 \\ x_1 - x_3 & y_1 - y_3 & 0 & 0 & x_3 - x_1 & y_3 - y_1 & 0 & 0 \\ 0 & 0 & x_2 - x_3 & y_2 - y_3 & x_3 - x_2 & y_3 - y_2 & 0 & 0 \\ x_1 - x_4 & y_1 - y_4 & 0 & 0 & 0 & 0 & x_4 - x_1 & y_4 - y_1 \\ 0 & 0 & 0 & 0 & x_3 - x_4 & y_3 - y_4 & x_4 - x_3 & y_4 - y_3 \end{bmatrix}.$$

With a slight abuse of notation, $\mathbf{v}$ is interpreted as a $d|V| \times 1$ column vector formed by stacking all the motion vectors in a single column. From (2.2), it is clear that any motion $\mathbf{v}$ which instantaneously preserves all the edge lengths of $(G, \mathbf{p})$ belongs to the nullspace of $R(G, \mathbf{p})$. Thus, $\mathbf{v}$ corresponding to any Euclidean motion, i.e. a motion that moves the entire framework $(G, \mathbf{p})$ without deformation, would belong to the nullspace of $R(G, \mathbf{p})$. Assume that $\mathbf{p}$ is non-degenerate. Then, the dimension of the subspace formed by Euclidean motions equals the dimension of the d-dimensional Euclidean group, which is $d(d+1)/2$. This implies that the nullity of $R(G, \mathbf{p})$ is at least $d(d+1)/2$. We say that the framework $(G, \mathbf{p})$ is *infinitesimally rigid* if the nullity is exactly $d(d+1)/2$, or equivalently, if $\mathrm{rank}(R(G, \mathbf{p})) = d|V| - d(d+1)/2$.

Infinitesimal rigidity of $(G, \mathbf{p})$ turns out to be both necessary and sufficient for local rigidity of $(G, \mathbf{p})$ if $\mathbf{p}$ is generic [28]. This implies that local rigidity of a generic framework $(G, \mathbf{p})$ has a complete characterization in terms of $\mathrm{rank}(R(G, \mathbf{p}))$. Moreover, given a graph G, the rank of the rigidity matrix is the same for any generic framework [27, 28]. As a result, we have the following useful fact which illustrates the computational tractability afforded by the genericity assumption.

Proposition 7 ([27, 28]). *Local rigidity is a generic property, i.e., either all or none of the generic configurations of a graph form a locally rigid framework.*

Thus, we can talk of a *graph* being *generically locally rigid*, by which we mean that every generic configuration of the graph results in a locally rigid framework. This suggests the existence of combinatorial characterizations for generic local rigidity. In fact, any combinatorial characterization of generic local rigidity of a graph G is based on the following observations.

Observation 8.

(i) Local rigidity of a generic framework is entirely characterized in terms of the rank of the rigidity matrix.

(ii) Rank of the rigidity matrix is the same for any generic framework of G.

Observation 8.(i) highlights the importance of the linear dependence relation among the rows of the rigidity matrix, and motivates the following definition.

Definition 9 (Independence)**.** *Given a graph G, a subgraph $H \subseteq G$ is said to be independent in $\mathbb{R}^d$ if the rows of the rigidity matrix $R(G, \mathbf{p})$ indexed by the edge set $E(H)$ are linearly independent for a d-dimensional generic framework $(G, \mathbf{p})$.*

Because of Observation 8.(ii) above, independence is well-defined. That is, if the rows of the rigidity matrix indexed by $E(H)$ are linearly independent for a generic framework, then they are linearly independent for every generic framework. This can be seen by considering $(H, \mathbf{p}|_H)$ as a generic framework of the subgraph H, and noting that the rows of $R(H, \mathbf{p}|_H)$ are linearly independent if and only if the rows of $R(G, \mathbf{p})$ indexed by $E(H)$ are linearly independent.

From the rank condition for infinitesimal rigidity, it is clear that a graph G is generically locally rigid if and only if it has an independent subgraph with $d|V| - d(d+1)/2$ edges. Thus, combinatorial characterization of generic local rigidity boils down to characterizing independence of a graph solely in terms of the graph properties. We have such a characterization in $\mathbb{R}^2$ due to Laman [35].

Theorem 10 ([35])**.** *A graph $G = (V, E)$ is independent in $\mathbb{R}^2$ if and only if $|E(H)| \leq 2|V(H)| - 3$ for every subgraph $H \subseteq G$ with $|V(H)| \geq 2$.*

The next corollary follows immediately.

Corollary 11. *If $G = (V, E)$ is independent, and $|E| = 2|V| - 3$, then G is generically locally rigid in $\mathbb{R}^2$.*

Finding such a characterization in $\mathbb{R}^3$ is a long-standing open problem in rigidity theory. We present a few additional combinatorial characterizations for rigidity in $\mathbb{R}^2$ which we shall need later. Specifically, we need Theorem 16 which characterizes *redundant rigidity* in $\mathbb{R}^2$. A graph is said to be *redundantly rigid* in $\mathbb{R}^d$ if the graph is generically locally rigid in $\mathbb{R}^d$, and remains generically locally rigid in $\mathbb{R}^d$ after removal of any edge. Redundant rigidity plays an important role in the context of global rigidity. We first define the notion of an *M-circuit*.

Definition 12 (M-circuit)**.** *Given a graph G, a subgraph $H \subseteq G$ is said to be an M-circuit if H is not independent, but every proper subgraph of H is independent.*

In other words, an M-circuit H is a minimally dependent subgraph of G. The following theorem from [31] characterizes an M-circuit in $\mathbb{R}^2$.

Theorem 13 ([31])**.** *Consider a graph $G = (V, E)$. Then the following are equivalent in $\mathbb{R}^2$:*

(i) *G is an M-circuit.*

(ii) *$|E| = 2|V| - 2$ and $G - e$ is minimally rigid for every $e \in E$ (a graph is said to be minimally rigid if the graph is generically locally rigid, but is not generically locally rigid after removal of any edge).*

(iii) *$|E| = 2|V| - 2$ and $|E(H)| \leq 2|V(H)| - 3$ for every subgraph $H \subseteq G$ with $2 \leq |V(H)| \leq |V| - 1$.*

Corollary 14. *Complete graph K_4 is an M-circuit in $\mathbb{R}^2$.*

We now define *M-connectivity*, and then close our discussion of generic local rigidity with a theorem due to [36, 31], which gives us the characterization of redundant rigidity that we are after.

Definition 15 (M-connectivity)**.** *A graph G is M-connected if every pair of edges in G belongs to an M-circuit of G.*

Theorem 16 ([36, 31])**.** *The following are true in $\mathbb{R}^2$:*

(i) *If a graph G is M-connected, then G is redundantly rigid.*

(ii) *If a graph G is 3-connected and each edge of G belongs to an M-circuit, then G is M-connected.*

2.2.4 Generic Global Rigidity

We have seen that local rigidity is a generic property. The question of whether global rigidity is a generic property was open until recently, when it was finally resolved in the affirmative in [30]. Just as we had the rigidity matrix with which we could characterize local rigidity of a generic framework, we have the notion of an *equilibrium stress matrix* using which we can characterize global rigidity of a generic framework.

Given a graph $G = (V, E)$ and a framework $(G, \mathbf{p})$, assign a scalar ω_{ij} to every $(i, j) \in E$, such that

$$\sum_{j:(i,j)\in E} \omega_{ij}(\mathbf{p}(j) - \mathbf{p}(i)) = 0, \quad i \in V. \tag{2.3}$$

Then $\omega = (\cdots, \omega_{ij}, \cdots)$ is known as an *equilibrium stress vector* of the framework $(G, \mathbf{p})$. Equilibrium stress vector for a framework is not unique in general; this can be seen by observing from (2.3) that any vector in the nullspace of $R(G, \mathbf{p})^\top$ is an equilibrium stress vector of $(G, \mathbf{p})$ [29].

To ascertain global rigidity of $(G, \mathbf{p})$, we are interested in the existence of an equilibrium stress vector that satisfies a particular property. To state what the property is, we regard an equilibrium stress vector ω as inducing an $n \times n$ symmetric matrix Ω, called an *equilibrium stress matrix*, with the entries made in the following manner: $\Omega_{i,j} = -\omega_{ij}$ for $(i, j) \in E$, $\Omega_{i,j} = 0$ for $(i, j) \notin E, i \neq j$, and each diagonal element chosen to make the corresponding row and column sums zero.

Assume that the configuration $\mathbf{p}$ is non-degenerate. From (2.3) and the fact that the vector $\mathbf{1}$ of all ones lies in the nullspace of Ω, it can be shown that the nullity of Ω is at least $d + 1$ for any equilibrium stress matrix Ω of $(G, \mathbf{p})$ [30]. A generic framework $(G, \mathbf{p})$ with at least $d + 2$ vertices is globally rigid if [29] and only if [30], there exists an equilibrium stress matrix whose nullity is exactly $d + 1$. Moreover, if a generic framework of G has an equilibrium stress matrix with nullity $d + 1$, then any generic framework of G also has an equilibrium stress matrix with nullity $d + 1$ [29, 30]. In the case of a generic framework $(G, \mathbf{p})$ with $d + 1$ or fewer vertices, it is shown in [27] that $(G, \mathbf{p})$ is globally rigid if and only if G is a complete graph. Thus, we have the following proposition.

Proposition 17 ([30])**.** *Global rigidity is a generic property, i.e. either all or none of the generic configurations of a graph form a globally rigid framework.*

In other words, under the assumption of genericity, global rigidity becomes a property of the graph, and we can talk of a *graph* being *generically globally rigid.* This also tells us that there ought to be characterizations of generic global rigidity solely in terms of the graph properties.

Hendrickson [37] gave the following combinatorial conditions necessary for a graph to be generically globally rigid in $\mathbb{R}^d$.

Theorem 18 ([37])**.** *If a graph G with at least $d + 2$ vertices is generically globally rigid in $\mathbb{R}^d$, then*

(i) G is $(d + 1)$-connected,

(ii) G is redundantly rigid in $\mathbb{R}^d$.

Later, Jackson and Jordan [31] showed that the conditions in Theorem 18 are also sufficient for generic global rigidity in $\mathbb{R}^2$. Thus, we have the following combinatorial characterization of generic global rigidity in $\mathbb{R}^2$.

Theorem 19 ([31])**.** *A graph G is generically globally rigid in $\mathbb{R}^2$ if and only if either G is a triangle, or* *(i) G is* 3*-connected, and*

(ii) G is redundantly rigid in $\mathbb{R}^2$.

Conditions in Theorem 18 are not sufficient for generic global rigidity in $\mathbb{R}^d$ for $d \geq 3$. We will see instances of such graphs in section 2.4.

2.2.5 Testing Generic Rigidity

The combinatorial characterization of graph independence given by Theorem 10 can be used to obtain a deterministic polynomial-time algorithm for testing generic local rigidity in $\mathbb{R}^2$ [38]. Moreover, since redundant rigidity reduces to checking generic local rigidity after removing an edge (and doing this for every edge), and there are polynomial time algorithms for testing graph connectivity [39], Theorem 19 implies that there are deterministic polynomial time algorithms for testing generic global rigidity in $\mathbb{R}^2$. This state of affairs does not carry over to higher dimensions since we do not have a complete combinatorial characterization for generic local and global rigidity in $\mathbb{R}^d$ for $d \geq 3$. However, the fact that local and global rigidity are generic properties, and that a randomly chosen configuration of a graph is generic with high probability, has led to efficient randomized algorithms for testing rigidity. It is demonstrated in [30] that testing for generic local and global rigidity has complexity RP, which means that there is a polynomial-time randomized algorithm that never outputs a false positive, and outputs a false negative less than half of the time.

We briefly discuss the general algorithmic structure for randomized testing of rigidity as detailed in [30]. Given a graph G, the algorithm picks a framework $(G, \mathbf{p})$ with coordinates chosen at random. In order to make the linear algebraic computations efficient, it is desirable to choose framework $(G, \mathbf{p})$ with integer coordinates chosen from $[1 : N]$ at random. The probability of algorithm returning a false negative depends inversely on N [30], which necessitates that N be sufficiently large. For testing generic local rigidity, the algorithm outputs decision based on the rank of the rigidity matrix $R(G, \mathbf{p})$. For testing generic global rigidity, there is an additional randomized procedure which involves picking an equilibrium stress vector from the nullspace of $R(G, \mathbf{p})^\top$. The algorithm then computes the rank of the associated equilibrium stress matrix Ω, and outputs a decision based on this. It is shown in [30] that an equilibrium stress vector chosen in a 'suitably random' manner has an equilibrium stress matrix Ω of minimal nullity, which justifies the algorithm basing its decision on the rank of such an Ω.

2.3 Unique Registrability

We now formulate a necessary and sufficient condition for unique registrability under the following assumptions:

(A1) Each patch contains at least $d + 1$ non-degenerate nodes.

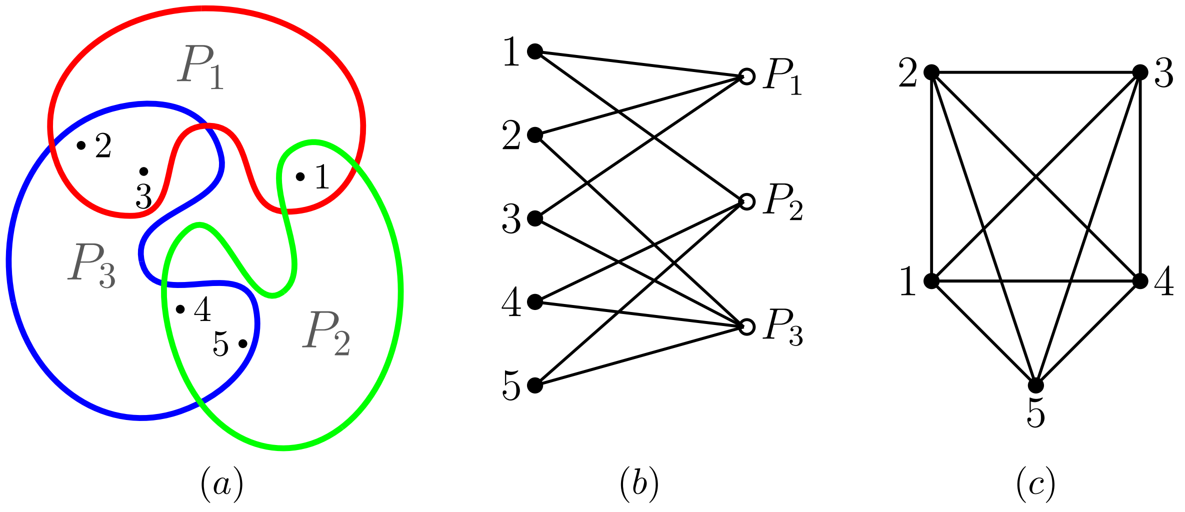

Figure 2.4: In this example, $\mathcal{S} = [1:5]$ and $\mathcal{P} = \{P_1, P_2, P_3\}$, where $P_1 = \{1,2,3\}$, $P_2 = \{1,4,5\}$ and $P_3 = \{2,3,4,5\}$. (a) Visualization of the node-patch correspondence, (b) Correspondence graph $\Gamma_C = (\mathcal{S}, \mathcal{P}, \mathcal{E})$, (c) Body graph Γ_B.

(A2) The nodes are in generic positions.

We briefly recall the rationale behind the assumptions. Under assumption (A1), uniqueness of the global coordinates and uniqueness of the patch transforms become equivalent, making unique registrability a well-defined notion (see Proposition 2). Moreover, we can easily force this assumption for divide-and-conquer algorithms [8, 4, 5]. assumption (A2) allows us to formulate conditions for unique registrability for *almost every* problem instance based solely on the combinatorial structure of the problem (see Proposition 7).

We now introduce the notion of a body graph, which will help us tie unique registrability to rigidity theory. For a network with correspondence graph $\Gamma_C = (\mathcal{S}, \mathcal{P}, \mathcal{E})$, consider a graph $\Gamma_B = (V, E)$, where $V = \mathcal{S}$, and $E = \{(k_1, k_2) : k_1, k_2 \in P_i \text{ for some } i \in [1:M]\}$. In other words, vertices of Γ_B correspond to the nodes in the network, and we connect two vertices by an edge if and only if the corresponding nodes belong to a common patch (see Fig. 2.4). We will call Γ_B the body graph of Γ_C. We derive the term body graph from [25], where a similar notion was introduced in the context of affine rigidity. Using the body graph, we now state our main result, whose proof is deferred to Section 2.6.

Theorem 20. *Under assumptions (A1) and (A2), the ground-truth solution* $(\bar{\mathbf{X}}, \bar{\mathfrak{R}})$ *is a unique solution of REG if and only if the body graph* Γ_B *is generically globally rigid.*

The import of Theorem 20 lies in the fact that generic global rigidity in an arbitrary dimension can be tested using a randomized polynomial-time algorithm [30]. Moreover, combining Theorem 20 with the combinatorial characterization of generic global rigidity in Theorem 19, and using results from rigidity theory, we get the following characterization of unique registrability in two dimensions; the proof is deferred to Section 2.6.

Corollary 21. *Under assumptions (A1) and (A2), a network is uniquely registrable in* $\mathbb{R}^2$ *if and only if the body graph* Γ_B *is* 3*-connected.*

The implication of Corollary 21 is that (assuming each patch has at least 3 nodes) we need only test for 3-connectivity to establish generic global rigidity of the body graph in $\mathbb{R}^2$. We need not perform an additional check for redundant rigidity, as required by Theorem 19. As is well-known, 3-connectivity can be tested efficiently using linear-time algorithms [39].

2.4 Quasi Connectivity

We now address the conjecture posed in [5] which asserts that, under assumption (A1) and the assumption that every set of $d + 1$ nodes is non-degenerate, quasi $(d + 1)$-connectivity of the correspondence graph Γ_C is sufficient for unique registrability in $\mathbb{R}^d$. We prove that, under assumptions (A1) and (A2), the conjecture holds for $d = 2$, but fails to hold for $d \geq 3$. We recall the definition of quasi connectivity [5].

Definition 22 (Quasi k-connectivity)**.** *The correspondence graph* $\Gamma_C = (\mathcal{S}, \mathcal{P}, \mathcal{E})$ *is said to be quasi* k*-connected if any two vertices in* $\mathcal{P}$ *have* k *or more* $\mathcal{S}$*-disjoint paths between them. (A set of paths is* $\mathcal{S}$*-disjoint if no two paths have a vertex from* $\mathcal{S}$ *in common.)*

Observation 23. *If a correspondence graph* Γ_C *is quasi* k*-connected, we can infer the following by dint of Definition 22:*

(a) There are at least k *participating nodes in every patch. (By a participating node, we mean a node that belongs to at least two patches.)*

(b) Let Γ_B *be the body graph of* Γ_C*. Let* H_i *be the clique of* Γ_B *induced by patch* P_i *where* $i \in [1 : M]$*. Then there are at least* k *disjoint* H_i*-*H_j *paths in the body graph, for every* $1 \leq i < j \leq M$ *(cf. Fig. 2.5).*

We relate quasi connectivity of the correspondence graph Γ_C to connectivity of the associated body graph Γ_B in the following theorem, whose proof we defer to Section 2.6.

Theorem 24 (Connectivity of Γ_C and Γ_B)**.**

(i) If the correspondence graph Γ_C *is quasi* k*-connected, then the body graph* Γ_B *is* k*-connected.*

(ii) If each patch has at least k *nodes and the body graph* Γ_B *is* k*-connected, then the correspondence graph* Γ_C *is quasi* k*-connected.*

We note some corollaries of Theorem 24. Corollary 25 was already proved in [5]; we give an alternative proof that utilizes the body graph. Corollary 26 establishes the conjecture posed in [5] for $d = 2$.

Corollary 25. *Under assumptions (A1) and (A2), quasi* $(d+1)$*-connectivity of* Γ_C *is a necessary condition for unique registrability in* $\mathbb{R}^d$.

Proof. From Theorem 20, unique registrability is equivalent to global rigidity of Γ_B. From Theorem 18, $(d+1)$-connectivity of Γ_B is a necessary condition for generic global rigidity of Γ_B in $\mathbb{R}^d$. The result now follows from Theorem 24. □

Corollary 26. *Under assumptions (A1) and (A2), quasi* 3*-connectivity of the correspondence graph* Γ_C *is sufficient for unique registrability in* $\mathbb{R}^2$.

Proof. Follows from Theorem 24 and Corollary 21. □

Corollary 26, in effect, says that the constraints imposed by quasi 3-connectivity of Γ_C ensure that Γ_B is redundantly rigid in addition to being 3-connected, and hence generically globally rigid in $\mathbb{R}^2$. But this trend does not carry over to $d \geq 3$. We demonstrate it with two examples for $d = 3$ (which appear in [40]), and then give a prescription for generating such counterexamples in higher dimensions.

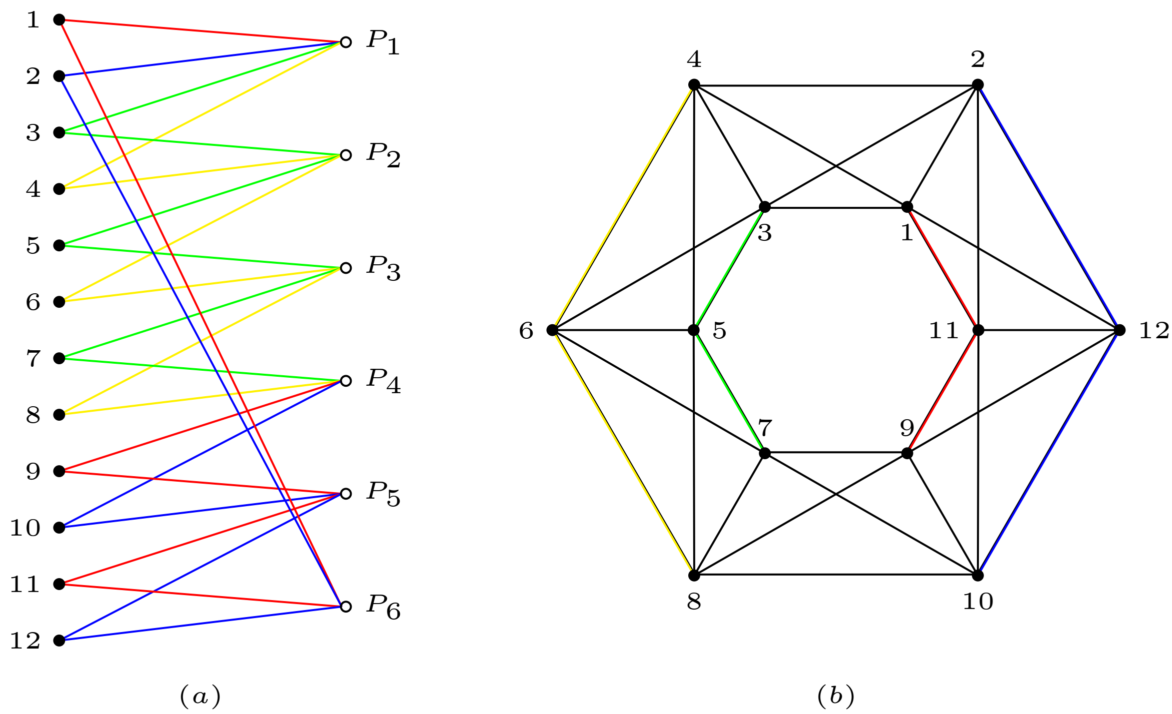

Figure 2.5: The figure shows a counterexample (Example 1) to the sufficiency of quasi 4-connectivity of the correspondence graph for unique registrability in $\mathbb{R}^3$. (a) Correspondence graph Γ_{C1}, (b) Body graph Γ_{B1}. The colored paths in (a) show the four $\mathcal{S}$-disjoint paths between P_1 and P_4. The corresponding disjoint H_1-H_4 paths in the body graph Γ_{B1} are colored in (b), where H_1 and H_4 are cliques induced by patches P_1 and P_4 (see text for details).

Example 1

Let $\mathcal{S} = [1 : 12]$, and $\mathcal{P} = \{P_1, \cdots, P_6\}$. That is, we have 12 nodes and 6 patches. Consider the following node-patch correspondence:

$$\begin{aligned} &P_1 = \{1, 2, 3, 4\}, P_2 = \{3, 4, 5, 6\}, \cdots, \\ &P_5 = \{9, 10, 11, 12\}, P_6 = \{11, 12, 1, 2\}. \end{aligned} \tag{2.4}$$

The correspondence graph Γ_{C1} and the associated body graph Γ_{B1} are shown in Fig. 2.5. It is easy to verify that Γ_{C1} is quasi 4-connected, or equivalently (Theorem 24), that Γ_{B1} is 4-connected. But, it can be shown [40] that the body graph Γ_{B1} is minimally rigid in $\mathbb{R}^3$, i.e. Γ_{B1} is generically locally rigid, but removing any edge destroys generic local rigidity. Hence Γ_{B1} is not redundantly rigid in $\mathbb{R}^3$. This implies (Theorem 18) that Γ_{B1} is not generically globally rigid, and thus (Theorem 20), the network is not uniquely registrable in $\mathbb{R}^3$.

Example 2

In this example, we will see that quasi $(d + 1)$-connectivity of the correspondence graph is not sufficient for generic global rigidity of the body graph, even if we ensure that the body graph is redundantly rigid. Let $\mathcal{S} = [1 : 18]$, and $\mathcal{P} = \{P_1, \cdots, P_6\}$, where

$$\begin{aligned} &P_1 = \{1, 2, 3, 4, 13\}, P_2 = \{3, 4, 5, 6, 14\}, \cdots, \\ &P_5 = \{9, 10, 11, 12, 17\}, P_6 = \{11, 12, 1, 2, 18\}. \end{aligned} \tag{2.5}$$

That is, we have added a non-participating node in each patch of Example 1. The correspondence graph Γ_{C2}, and the associated body graph Γ_{B2} are shown in Fig. 2.6. It is easy to verify that Γ_{C2} is quasi 4-connected, or equivalently, that Γ_{B2} is 4-connected. Moreover, Γ_{B2} is redundantly rigid [40]. But, from the fact that Γ_{B1} in Example 1 is not generically globally rigid in $\mathbb{R}^3$, it can be deduced (Proposition 34) that Γ_{B2} is also not generically globally rigid in $\mathbb{R}^3$. Thus, the network is not uniquely registrable in $\mathbb{R}^3$.

Graphs such as Γ_{B2} in Example 2 above, which satisfy both conditions in Theorem 18, but are not generically globally rigid in $\mathbb{R}^d$, are known as *H-graphs*. By an operation called *coning*, which takes a graph G and adds a new vertex adjacent to every vertex of G, a d-dimensional H-graph can be turned into a $(d + 1)$-dimensional H-graph [41, 42]. In terms of node-patch correspondence, this equates to adding a new node that belongs to every patch. Thus, by applying $d - 3$ coning operations to Γ_{B2}, we can generate a node-patch correspondence with a quasi $(d+1)$-connected correspondence graph, thus obtaining a network which is not uniquely registrable in $\mathbb{R}^d$ for $d > 3$.

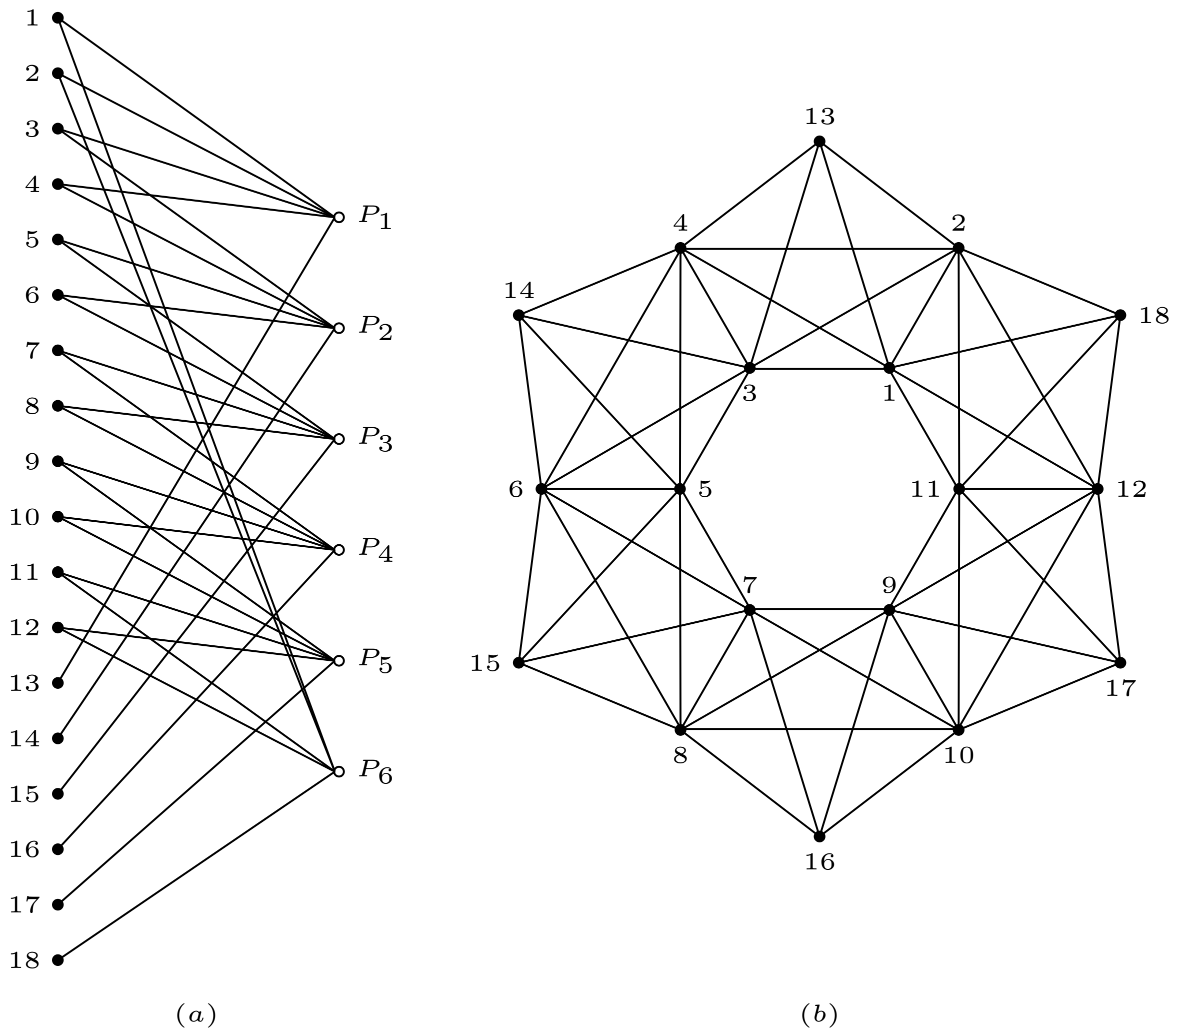

Figure 2.6: The figure shows a counterexample (Example 2) to sufficiency of quasi 4-connectivity of the correspondence graph for unique registrability in $\mathbb{R}^3$ even when the body graph is redundantly rigid. (*a*) Correspondence graph Γ_{C2}, (*b*) Body graph Γ_{B2} (see text for details).

2.5 Discussion

In this chapter, we looked at the notion of unique registrability through the lens of rigidity theory. Given that there are two families of unknowns inherent in the problem—the global coordinates and the patch transforms—we first addressed the question as to what uniqueness exactly means for the rigid registration problem. We ended up making a mild assumption of non-degeneracy that makes the notion of uniqueness equivalent for both families of unknowns. We then introduced the notion of body graph which allowed us to reformulate the question of unique registrability into a question about graph rigidity. Specifically, we concluded that unique registrability is equivalent to global rigidity of the body graph. This equivalence opened up the possibility of using non-trivial results from rigidity theory. In particular, we showed that the necessary condition of quasi $(d+1)$-connectivity of the correspondence graph, which was conjectured in [5] to be sufficient for unique registrability in $\mathbb{R}^d$, is indeed sufficient for $d = 2$, but fails to be so for $d \geq 3$. The practical utility of these characterizations is that they lead to efficiently testable criteria for unique registrability. To ascertain unique registrability in $\mathbb{R}^2$, we only need to test either quasi 3-connectivity of the correspondence graph or 3-connectivity of the body graph. Two and three vertex-connectivity can be tested efficiently using linear-time algorithms [39]. For $d \geq 3$, unique registrability can be tested simply by testing generic global rigidity of the body graph, and there exists a polynomial-time randomized algorithm for the latter.

2.6 Technical Proofs

In this section, we give proofs for Theorem 20, Corollary 21 and Theorem 24.

2.6.1 Proof of Theorem 20

We will show that unique registrability is equivalent to global rigidity of the body graph framework corresponding to the ground-truth. Then, the assumption of genericity (A2) along with Proposition 7 (genericity of local rigidity) allows us to remove the dependence on any particular framework, thus proving the theorem. We first make some definitions specialized to the registration problem which allow us to express the question of uniqueness registrability in a form amenable to a rigidity theoretic analysis.

Definition 27 (Node-patch framework)**.** *Given a correspondence graph* $\Gamma_C = (\mathcal{S}, \mathcal{P}, \mathcal{E})$*, and a map* $\mathbf{x} : \mathcal{S} \to \mathbb{R}^d$ *that assigns coordinates to the nodes, the pair* $(\Gamma_C, \mathbf{x})$ *is called a node-patch framework.*

Definition 28 (Equivalence of node-patch frameworks)**.** *Two node-patch frameworks* $(\Gamma_C, \mathbf{x})$ *and* $(\Gamma_C, \mathbf{y})$ *are said to be equivalent, denoted by* $(\Gamma_C, \mathbf{x}) \sim (\Gamma_C, \mathbf{y})$*, if* $\mathbf{x}(k) = \mathcal{Q}_i \mathbf{y}(k)$*,* $(k, i) \in \mathcal{E}$*, where* $\mathcal{Q}_i$ *is a rigid transform.*

Definition 29 (Congruence of node-patch frameworks)**.** *Two node-patch frameworks* $(\Gamma_C, \mathbf{x})$ *and* $(\Gamma_C, \mathbf{y})$ *are said to be congruent, denoted by* $(\Gamma_C, \mathbf{x}) \equiv (\Gamma_C, \mathbf{y})$*, if* $\mathbf{x}(k) = \mathcal{Q}\mathbf{y}(k),\ k \in \mathcal{S}$*, where* $\mathcal{Q}$ *is a rigid transform.*

Given a solution $(\mathbf{X}, \mathfrak{R})$ to REG, where $\mathbf{X} = (x_k)_{k=1}^N$, $\mathfrak{R} = (\mathcal{R}_i)_{i=1}^M$, we will denote by $\mathbf{x}$ the map that assigns to node k the coordinate x_k, and say that $(\Gamma_C, \mathbf{x})$ is the node-patch framework corresponding to the solution $(\mathbf{X}, \mathfrak{R})$.

Proposition 30. *Let* $(\mathbf{X}, \mathfrak{R})$ *and* $(\mathbf{Y}, \mathfrak{T})$ *be two solutions to REG. Then the corresponding node-patch frameworks* $(\Gamma_C, \mathbf{x})$ *and* $(\Gamma_C, \mathbf{y})$ *are equivalent.*

Proof. Since $(\mathbf{X}, \mathfrak{R})$ and $(\mathbf{Y}, \mathfrak{T})$ are solutions to REG, we have that $\mathbf{x}(k) = \mathcal{R}_i(x_{k,i})$ and $\mathbf{y}(k) = \mathcal{T}_i(x_{k,i})$, $k \in P_i$, $i \in [1 : M]$. Thus $\mathbf{x}(k) = \mathcal{Q}_i \mathbf{y}(k)$, where $\mathcal{Q}_i = \mathcal{R}_i \circ \mathcal{T}_i^{-1}$. □

Proposition 31. *Let* $(\mathbf{X}, \mathfrak{R})$ *be a solution to REG with the corresponding node-patch framework* $(\Gamma_C, \mathbf{x})$ *and let* $\mathbf{y}$ *be such that* $(\Gamma_C, \mathbf{y}) \sim (\Gamma_C, \mathbf{x})$. *Then there exists some* $\mathfrak{T}$ *for which* $(\mathbf{Y}, \mathfrak{T})$ *is a solution of REG.*

Proof. Indeed, $(\Gamma_C, \mathbf{y}) \sim (\Gamma_C, \mathbf{x})$ implies that there exists rigid transforms $(\mathcal{Q}_i)_{i=1}^M$ such that $\mathbf{y}(k) = \mathcal{Q}_i \mathbf{x}(k)$, $(k, i) \in \mathcal{E}$. Since $(\mathbf{X}, \mathcal{R})$ is a solution to REG, we have $\mathbf{x}(k) = \mathcal{R}_i(x_{k,i})$, $(k, i) \in \mathcal{E}$. Thus, $\mathbf{y}(k) = (\mathcal{Q}_i \circ \mathcal{R}_i)(x_{k,i})$, which shows that $(\mathbf{Y}, \mathcal{T})$ is a solution to REG, where $\mathbf{Y} = (\mathbf{y}(k))_{k=1}^N$ and $\mathcal{T} = (\mathcal{Q}_i \circ \mathcal{R}_i)_{i=1}^M$. □

Foregoing definitions and propositions allow us to express the condition of unique registrability in a compact manner. Namely, let $(\Gamma_C, \bar{\mathbf{x}})$ be the ground-truth node-patch framework. Then, under assumption (A1), REG has a unique solution if and only if for any node-patch framework $(\Gamma_C, \mathbf{y})$ such that $(\Gamma_C, \mathbf{y}) \sim (\Gamma_C, \bar{\mathbf{x}})$, we have $(\Gamma_C, \mathbf{y}) \equiv (\Gamma_C, \bar{\mathbf{x}})$. The next two propositions relate node-patch framework and body graph framework.

Proposition 32. *Two node-patch frameworks* $(\Gamma_C, \mathbf{x})$ *and* $(\Gamma_C, \mathbf{y})$ *are equivalent (Definition 28) if and only if the body graph frameworks* $(\Gamma_B, \mathbf{x})$ *and* $(\Gamma_B, \mathbf{y})$ *are equivalent (Definition 3).*

Proof. Suppose $(\Gamma_C, \mathbf{x}) \sim (\Gamma_C, \mathbf{y})$. Pick an arbitrary edge $(k, l) \in E$ in the body graph $\Gamma_B = (V, E)$. From construction of Γ_B, $(k, l) \in E$ if and only if there is a patch, say P_i, that contains both the nodes k and l. Since $(\Gamma_C, \mathbf{x}) \sim (\Gamma_C, \mathbf{y})$, there exists a rigid transform $\mathcal{Q}_i$ such that $\mathbf{x}(k) = \mathcal{Q}_i \mathbf{y}(k)$ and $\mathbf{x}(l) = \mathcal{Q}_i \mathbf{y}(l)$. This implies that $\mathbf{x}(k) - \mathbf{x}(l) = \mathcal{Q}_i(\mathbf{y}(k) - \mathbf{y}(l))$, from where it follows that $\|\mathbf{x}(k) - \mathbf{x}(l)\| = \|\mathbf{y}(k) - \mathbf{y}(l))\|$. Thus, $(\Gamma_B, \mathbf{x}) \sim (\Gamma_B, \mathbf{y})$.

Conversely, suppose $(\Gamma_B, \mathbf{x}) \sim (\Gamma_B, \mathbf{y})$. Consider an arbitrary patch P_i. Note that any subgraph of Γ_B induced by a patch is a clique. This, along with the assumption that $(\Gamma_B, \mathbf{x}) \sim (\Gamma_B, \mathbf{y})$, implies that $\|\mathbf{x}(k) - \mathbf{x}(l)\| = \|\mathbf{y}(k) - \mathbf{y}(l))\|$ for every $k, l \in P_i$, which, in turn, implies that there exists a rigid transform $\mathcal{Q}_i$ such that $\mathbf{x}(v) = \mathcal{Q}_i \mathbf{y}(v)$, $v \in P_i$. Thus, $(\Gamma_C, \mathbf{x}) \sim (\Gamma_C, \mathbf{y})$. □

Proposition 33. *Two node-patch frameworks* $(\Gamma_C, \mathbf{x})$ *and* $(\Gamma_C, \mathbf{y})$ *are congruent (Definition 29) if and only if the body graph frameworks* $(\Gamma_B, \mathbf{x})$ *and* $(\Gamma_B, \mathbf{y})$ *are congruent (Definition 4).*

The above result easily follows from Definitions 4 and 29. We are now in a position to complete the proof of Theorem 20. Suppose REG has a unique solution. We will show that the body graph framework $(\Gamma_B, \bar{\mathbf{x}})$ is globally rigid. Consider a framework $(\Gamma_B, \mathbf{y}) \sim (\Gamma_B, \bar{\mathbf{x}})$. Then, by Proposition 32, $(\Gamma_C, \mathbf{y}) \sim (\Gamma_C, \bar{\mathbf{x}})$. By Proposition 31, this implies that $(\Gamma_C, \mathbf{y})$ correponds to a solution of REG. Now, since REG has a unique solution, $(\Gamma_C, \mathbf{y}) \equiv (\Gamma_C, \bar{\mathbf{x}})$. Thus, by Proposition 33, $(\Gamma_B, \mathbf{y}) \equiv (\Gamma_B, \bar{\mathbf{x}})$.

Conversely, suppose $(\Gamma_B, \bar{\mathbf{x}})$ is globally rigid. Let $(\mathbf{Y}, \mathcal{T})$ be a solution to REG. By Proposition 30, $(\Gamma_C, \mathbf{y}) \sim (\Gamma_C, \bar{\mathbf{x}})$. Hence, by Proposition 32, $(\Gamma_B, \mathbf{y}) \sim (\Gamma_B, \bar{\mathbf{x}})$. This, by global rigidity of $(\Gamma_B, \bar{\mathbf{x}})$, implies that $(\Gamma_B, \mathbf{y}) \equiv (\Gamma_B, \bar{\mathbf{x}})$. Finally, by Proposition 33, $(\Gamma_C, \mathbf{y}) \equiv (\Gamma_C, \bar{\mathbf{x}})$.

2.6.2 Proof of Theorem 21.

Before proving Theorem 21, we state and prove the following proposition.

Proposition 34. *Given a graph $G = (V, E)$, consider the graph $G' = (V \cup \{v'\}, E')$ obtained by adding a new vertex v' to G and attaching it to a clique $H \subseteq G$, i.e., v' is adjacent to every vertex of H and to no other vertex of G. If G' is generically globally rigid, then G is generically globally rigid.*

Proof. Suppose G is not generically globally rigid. Consider two frameworks $(G, \mathbf{p})$ and $(G, \mathbf{q})$ which are equivalent but not congruent. To these frameworks, add the new vertex v' to get new frameworks $(G', \mathbf{p}')$ and $(G', \mathbf{q}')$ such that the distance between v' and any vertex of the subgraph H is equal in both $(G', \mathbf{p}')$ and $(G', \mathbf{q}')$. Note that this can be done because H is a clique and so the subframeworks induced by H would be congruent in the two frameworks $(G, \mathbf{p})$ and $(G, \mathbf{q})$. Clearly, the new frameworks $(G', \mathbf{p}')$ and $(G', \mathbf{q}')$ are equivalent. But they are not congruent because $(G, \mathbf{p})$ and $(G, \mathbf{q})$ were not congruent to begin with. Thus, G' is not generically globally rigid. □

We now prove Theorem 21. The necessity of 3-connectivity of the body graph Γ_B for unique registrability in $\mathbb{R}^2$ follows from Theorem 20 and Theorem 18. We now establish sufficiency.

Given that the body graph Γ_B is 3-connected, we will prove that Γ_B is generically globally rigid in $\mathbb{R}^2$; this, by Proposition 20, would imply unique registrability in $\mathbb{R}^2$. By assumption (A1), there are at least 3 nodes in each patch. Consider the following cases:

Case 1: Each patch contains at least 4 *nodes.*
Pick an arbitrary edge (k, l) belonging to Γ_B. The fact that there is an edge between vertices k and l implies that there must be a patch, say P_i, which contains the nodes k and l. Since P_i contains at least 4 nodes, we can pick two nodes $\bar{k}$ and $\bar{l}$ belonging to P_i which are distinct from the nodes k and l. Now, P_i induces a clique, say H_i, in Γ_B. This implies that the subgraph of Γ_B induced by the vertex set $\{k, l, \bar{k}, \bar{l}\}$ is K_4, which is an M-circuit (Corollary 14) containing the edge (k, l). The edge (k, l) was chosen arbitrarily, and thus, we have shown that every edge of Γ_B belongs to an M-circuit. Since Γ_B is also 3-connected, we can use Theorem 16 to conclude that Γ_B is M-connected, and hence redundantly rigid. Thus, Γ_B satisfies conditions in Theorem 19, and is hence generically globally rigid in $\mathbb{R}^2$.

Case 2: There are patches with exactly 3 *nodes.*
Suppose there are m patches $P_1, \cdots, P_m$ that contain exactly 3 nodes. Add a new node k_1 exclusively to patch P_1 and call the resulting patch P_1'. The effect of this on the body graph is the addition of a degree-3 vertex k_1 adjacent to the vertices of the clique induced by the 3 nodes in P_1. Call the resulting body graph Γ_B^1. Addition of a degree-k vertex to a k-connected

graph results in a k-connected graph. Thus, Γ_B^1 is 3-connected. We continue inductively: after obtaining Γ_B^i, add a new node k_{i+1} exclusively to patch P_{i+1} to get P'_{i+1} and the resulting body graph Γ_B^{i+1}. Note that we preserve 3-connectivity at every step of the induction. We stop after m steps, i.e., after we have obtained the body graph Γ_B^m. As a result of this inductive procedure, every patch now contains at least 4 nodes. Hence, from the arguments made in *Case 1* above, Γ_B^m is generically globally rigid in $\mathbb{R}^2$. Now, Γ_B^m was obtained from Γ_B^{m-1} by addition of a vertex and attaching it to a clique. Hence, from Proposition 34, Γ_B^{m-1} is generically globally rigid in $\mathbb{R}^2$. Backtracking similarly in an inductive fashion and employing Proposition 34 at every step, we deduce that the original body graph Γ_B is generically globally rigid in $\mathbb{R}^2$.

2.6.3 Proof of Theorem 24.

We first prove Theorem 24.(ii). We are given that every patch has at least k nodes and the body graph Γ_B is k-connected. Let H_i and H_j be the cliques of Γ_B induced by patches P_i and P_j, $i \neq j$. To establish quasi k-connectivity of Γ_C, it suffices to show that there exists k disjoint H_i-H_j paths. Indeed, it is clear from Definition 22 that the existence of k disjoint H_i-H_j paths in Γ_B implies the existence of k $\mathcal{S}$-disjoint paths in Γ_C between P_i and P_j. Add two new vertices a and b to Γ_B such that a is adjacent to every vertex of H_i (and to no other vertex of Γ_B), and b is adjacent to every vertex of H_j (and to no other vertex of Γ_B). Since each patch has at least k nodes, degree(a) $\geq k$ and degree(b) $\geq k$. Addition of a degree-k vertex to a k-connected graph results in a k-connected graph. Thus, the graph obtained after adding a and b to Γ_B is k-connected. This implies that there are at least k independent paths between a and b. Now, each such path has to be of the form $a - v_1 - \cdots - v_r - b$, where $v_1 \in H_i$ and $v_r \in H_j$. This is because a is adjacent only to vertices from H_i and b is adjacent only to vertices from H_j. Removing a and b from every such independent path gives us k disjoint H_i-H_j paths.

We now prove Theorem 24.(i). Assume, without loss of generality, that no two patches are identical. To prove k-connectivity of the body graph $\Gamma_B = (V, E)$, we will show that given arbitrary vertices $a, b \in V$, there exists k independent paths between them. We consider the following cases:

Case 1: a and b do not belong to the same patch.
Suppose $a \in P_i$ and $b \in P_j$, where $i \neq j$. Denote the cliques of Γ_B induced by patches P_i and P_j as H_i and H_j. Since Γ_C is quasi k-connected, there exists k disjoint H_i-H_j paths (Observation 23). Note that a vertex in $V(H_i) \cap V(H_j)$ is also considered an H_i-H_j path. Let $P = v_1 - \cdots - v_r$ be one such path, where $v_1 \in H_i$ and $v_r \in H_j$. Since H_i and H_j are cliques, $(a, v_1) \in E$ and $(v_r, b) \in E$. Thus for each of the k disjoint H_i-H_j paths, we can, if needed, append vertices a and b at the ends to make it of the form $a - \cdots - b$. For instance, if $v_1 \neq a$ and $v_r \neq b$, we modify the path to $a - v_1 - \cdots - v_r - b$. Thus, we have k independent paths between a and b.

Case 2: a and b belong to the same patch.
Suppose a and b belong to patch P_l. Quasi k-connectivity of the correspondence graph implies that each patch has at least k participating nodes (Observation 23). In particular, this means that the clique H_l of Γ_B induced by P_l has at least k vertices. Thus, if a and b belong to P_l, there are at least $k-1$ independent paths within the clique H_l. If P_l has more than k nodes, we thus get k independent paths between a and b, all from within H_l. But suppose P_l has exactly k nodes. We need an additional path between a and b that is independent of the $k-1$ paths we have from within H_l. Since we have exactly k nodes in P_l, each node has to be participating, i.e., each node belongs to at least 2 patches. We consider the following sub-cases:

Sub-case I: There is a patch P_i, $i \neq l$, that contains both a and b.
In this case we get the additional path of the form $a - v - b$, where $v \in P_i$ and $v \notin P_l$, which, clearly, is independent of the $k-1$ paths from within H_l. The assumption that no two patches are identical ensures the existence of the v in question.

Sub-case II: There is no patch other than P_l that contains both a and b.
Suppose $a \in P_i$ and $b \in P_j$, $i \neq j$. From the quasi k-connectivity assumption, we know there are k disjoint H_i-H_j paths. Moreover, recall that there are exactly k vertices in H_l. Consider the following possibilities:

(i) Suppose every disjoint H_i-H_j path contains a vertex from H_l. This is possible if and only if each path contains exactly one vertex from H_l. In this case, there exists a path of the form $a - v_1 - \cdots - v_r$, such that $v_1, \cdots, v_r \notin H_l$, and $v_r \in H_j$. From completeness of the clique H_j, we can append b to the end of this path to get $a - v_1 - \cdots - v_r - b$. This path is independent of the $k-1$ paths we have from within H_l. Thus we have the required additional path.

(ii) The only other case is when there exists a disjoint H_i-H_j path that has no vertex from H_l. Let that path be $v_1 - \cdots - v_r$ where $v_1 \in H_i$ and $v_r \in H_j$. From completeness of the cliques H_i and H_j, we can append a and b to the ends of this path to get $a - v_1 - \cdots - v_r - b$, which is independent of the $k-1$ paths we have from within H_l. Again, we have the required additional path.

Chapter 3

Tightness of Convex Relaxation

3.1 Introduction

In rigid registration, we wish to compute the global coordinates of a set of N points, given the local coordinates of M overlapping subsets of these points (patches), where each patch has its own local coordinate system. The M local coordinate systems are related to each other by rigid transforms, which are otherwise unknown. Suppose the N points are labelled $[1:N]$, and the M patches are denoted by $\{P_1, \cdots, P_M\}$. If point k belongs to P_i, let $\mathbf{x}_{k,i}$ denote its corresponding local coordinate. Ideally, when the local coordinates $\mathbf{x}_{k,i}$ are exact, the rigid registration problem is to find global coordinates $\mathbf{z}_1, \ldots, \mathbf{z}_N$, and rigid transforms $(\mathbf{O}_i, \mathbf{t}_i)$ such that

$$\mathbf{z}_k = \mathbf{O}_i \mathbf{x}_{k,i} + \mathbf{t}_i, \tag{3.1}$$

for every $k \in P_i$, $i \in [1:M]$. In general, however, the local coordinate measurements may be noisy and we cannot expect (3.1) to hold. Instead, we consider the following least-squares minimization [14] to solve the registration problem:

$$\text{(LS-REG)} \qquad \min_{(\mathbf{O}_i),(\mathbf{t}_i),(\mathbf{z}_k)} \sum_{i=1}^{M} \sum_{k \in P_i} \|\mathbf{z}_k - \big(\mathbf{O}_i \mathbf{x}_{k,i} + \mathbf{t}_i\big)\|^2.$$

The fundamental difficulty of LS-REG stems from the fact that variables $\mathbf{O}_i$ are constrained to be in $\mathbb{O}(d)$, a nonconvex set. In fact, $\mathbb{O}(d)$ is not even connected [43]. E.g., $\mathbb{O}(1) = \{-1, 1\}$, and $\mathbb{O}(2)$ is topologically equivalent to the union of two disjoint circles in $\mathbb{R}^2$. This makes it difficult to apply local optimization methods to solve LS-REG. In particular, if we initialize the method on the wrong component, then we cannot hope to get close to the global optimum.

To combat these issues with the domain of optimization problem LS-REG, we derive an equivalent optimization problem – a rank-constrained semidefinite program – in a higher-dimensional space (i.e., we *lift* the problem to a higher dimension). We then drop the rank constraint to get a convex semidefinite program, called a *convex relaxation* of the rank-constrained problem. Overall, this procedure is an instance of the well-known technique of *Lift and Relax* [44]. Empirically, we observe that the global optimum of the relaxed problem is also a global optimum of the rank-constrained problem when the noise level in the local coordinate measurements is below a certain threshold (Fig. 3.1); in this case, we say that the relaxation is *tight* [1]. In this chapter, we give a mathematical justification of this phenomenon using the theory of Lagrange duality. Our analysis is inspired by the analysis in [1], where the authors explain a similar tightness phenomenon in the context of the phase synchronization problem.

3.1.1 Organization

This chapter is organized as follows. In Section 3.2, we derive a convex relaxation for LS-REG. In Section 3.3, we discuss how duality might help explain the tightness phenomenon of this relaxation, and then derive the main result (Theorem 38) of this chapter. We conclude with a discussion in Section 3.4, which is followed by an Appendix, where we give proofs of some of the results used in Section 3.3.

3.1.2 Notations

$[m:n]$ denotes the integers $\{m, m+1, \ldots, n\}$. $\mathbb{S}^k$ denotes the set of $k \times k$ symmetric matrices. $\mathbb{S}^k_+$ denotes the set of positive semidefinite (PSD) matrices, i.e., $k \times k$ symmetric matrices with nonnegative eigenvalues. $\mathbf{I}_n$ denotes the $n \times n$ identity matrix. Any $\mathbf{X} \in \mathbb{R}^{Md \times Md}$ can be seen as being composed of $M \times M$ blocks, with each block of size $d \times d$; we denote the (i,j)-th block of $\mathbf{X}$ by $[\mathbf{X}]_{ij}$. For a matrix $\mathbf{X}$, $\|\cdot\|$ and $\|\cdot\|_2$ denote the Frobenius and spectral norms; the latter is simply the largest singular value $\sigma_{\max}(\mathbf{X})$. We note that $\|\mathbf{X}\|_2 \leq \|\mathbf{X}\|$. $\mathrm{Tr}(\mathbf{A})$ denotes the trace of $\mathbf{A}$; and $\langle \mathbf{X}, \mathbf{Y} \rangle = \mathrm{Tr}(\mathbf{XY})$ is the inner product between two symmetric matrices $\mathbf{X}$ and $\mathbf{Y}$. For a symmetric matrix $\mathbf{X}$, $\lambda_i(\mathbf{X})$ denotes the i-th eigenvalue, where we assume the eigenvalues to be arranged in ascending order.

3.2 Convex Relaxation

In this section, we obtain a convex relaxation of LS-REG, which has been derived in detail in [14]. Here, we outline the major steps of the derivation. Note that LS-REG can be rewritten as

$$\min_{(\mathbf{O}_i)} \left[\min_{(\mathbf{t}_i),(\mathbf{z}_k)} \sum_{i=1}^{M} \sum_{k \in P_i} \|\mathbf{z}_k - (\mathbf{O}_i \mathbf{x}_{k,i} + \mathbf{t}_i)\|^2 \right]. \tag{3.2}$$

It was observed in [14] that if we fix the orthogonal transforms $(\mathbf{O}_i)$, the minimization problem inside the square brackets in (3.2) is a quadratic convex optimization in $(\mathbf{t}_i)$ and $(\mathbf{z}_k)$, with a closed-form optimum which is linear in $(\mathbf{O}_i)$. That is, if we knew the optimal orthogonal transforms for LS-REG, the optimal values of $(\mathbf{t}_i)$ and $(\mathbf{z}_k)$ could be computed using a simple linear transform. Based on this observation, LS-REG can be reduced to the following optimization involving just the orthogonal transforms:

$$\text{(ORTHO-REG)} \qquad \min_{\mathbf{O}_1,\ldots,\mathbf{O}_M \in \mathbb{O}(d)} \sum_{i,j=1}^{M} \operatorname{Tr}\left(\mathbf{C}_{ij}\mathbf{O}_j^\top \mathbf{O}_i\right).$$

In particular, ORTHO-REG is equivalent to LS-REG in the following sense: Suppose $(\mathbf{O}_i^*)_{i=1}^M$, $(\mathbf{t}_i^*)_{i=1}^M$, $(\mathbf{z}_k^*)_{k=1}^N$ are globally optimal for LS-REG. Then,

$$\sum_{i=1}^{M}\sum_{k\in P_i}\|\mathbf{z}_k^* - (\mathbf{O}_i^*\mathbf{x}_{k,i} + \mathbf{t}_i^*)\|^2 = \min_{\mathbf{O}_1,\ldots,\mathbf{O}_M \in \mathbb{O}(d)} \sum_{i,j=1}^{M} \operatorname{Tr}\left(\mathbf{C}_{ij}\mathbf{O}_j^\top \mathbf{O}_i\right) = \operatorname{Tr}\left(\mathbf{C}_{ij}\mathbf{O}_j^{*\top} \mathbf{O}_i^*\right). \tag{3.3}$$

Note that ORTHO-REG is a quadratic form in $(\mathbf{O}_i)$. Thus, defining $\mathbf{O} = (\mathbf{O}_1 \cdots \mathbf{O}_M) \in \mathbb{R}^{d\times Md}$ and $\mathbf{C} = \left((\mathbf{C}_{ij})_{i,j=1}^M\right) \in \mathbb{S}^{Md}$, we can rewrite ORTHO-REG as

$$\min_{\mathbf{O}\,\in\,\mathcal{D}} \operatorname{Tr}\left(\mathbf{C}\mathbf{O}^\top\mathbf{O}\right), \tag{3.4}$$

where the domain $\mathcal{D} \subset \mathbb{R}^{d\times Md}$ can be interpreted as the set of "row block-vectors", i.e. row "vectors" where each element of a row is a $d \times d$ orthogonal matrix. The data matrix $\mathbf{C}$ is an $Md \times Md$ positive semidefinite matrix [14]. This is because the objective in (3.4) is simply a rewriting of the objective in LS-REG, which, being the sum of norms, is always nonnegative.

We now introduce a new variable, (Gram matrix) $\mathbf{G} = \mathbf{O}^\top\mathbf{O} \in \mathbb{R}^{Md\times Md}$, which allows us to write (3.4) as

$$\min_{\mathbf{G}\,\in\,\mathcal{S}} \operatorname{Tr}\left(\mathbf{C}\mathbf{G}\right), \tag{3.5}$$

where $\mathcal{S} \in \mathbb{R}^{Md\times Md}$ consists of all matrices $\mathbf{G}$ which can be decomposed as $\mathbf{G} = \mathbf{O}^\top\mathbf{O}$ with $\mathbf{O} \in \mathcal{D}$. To get a better characterization of domain $\mathcal{S}$, note that any $\mathbf{G} \in \mathcal{S}$ necessarily satisfies the following properties:

(i) $\mathbf{G} \in \mathbb{S}^{Md}$, i.e. $\mathbf{G}$ is symmetric;

(ii) $\mathbf{G} \in \mathbb{S}_+^{Md}$, i.e. $\mathbf{G}$ is positive semidefinite;

(iii) $[\mathbf{G}]_{ii} = \mathbf{I}_d$, where $[\mathbf{G}]_{ii}$ denotes i-th diagonal block of $\mathbf{G}$, and $\mathbf{I}_d$ is $d \times d$ identity matrix;

(iv) $\operatorname{rank}(\mathbf{G}) = d$.

Conversely, it is not difficult to show (using spectral decomposition) that any $\mathbf{G} \in \mathbb{R}^{Md\times Md}$ satisfying properties (i)-(iv) above, can be decomposed as $\mathbf{G} = \mathbf{O}^\top\mathbf{O}$ with $\mathbf{O} \in \mathcal{D}$. This implies that properties (i)-(iv) fully characterize the set $\mathcal{S}$. In other words, (3.5) can be written as the following optimization problem:

$$\text{(REG-SDP)} \qquad \begin{aligned} \min_{\mathbf{G} \in \mathbb{S}_+^{Md}} \quad & \mathrm{Tr}\,(\mathbf{CG}) \\ \text{subject to} \quad & [\mathbf{G}]_{ii} = \mathbf{I}_d, \quad i \in [1:M], \\ & \mathrm{rank}(\mathbf{G}) = d, \end{aligned}$$

where $\mathbb{S}_+^{Md}$ denotes the set of positive semidefinite symmetric matrices of size $Md \times Md$. This is a standard semidefinite program (SDP) [45], but with an additional rank constraint. In particular, observe that the nonconvex nature of ORTHO-REG is isolated in the rank constraint of REG-SDP. This suggests an obvious convex relaxation of REG-SDP — simply drop the rank constraint to get the following convex semidefinite program:

$$\text{(C-SDP)} \qquad \begin{aligned} \min_{\mathbf{G} \in \mathbb{S}_+^{Md}} \quad & \mathrm{Tr}(\mathbf{CG}) \\ \text{subject to} \quad & [\mathbf{G}]_{ii} = \mathbf{I}_d, \quad i \in [1:M]. \end{aligned}$$

Let $\mathbf{G}^*$ be a global optimum of C-SDP. From the fact that the diagonal blocks of $\mathbf{G}^*$ are $\mathbf{I}_d$, we get that $\mathrm{rank}(\mathbf{G}^*) \geq d$ (see Proposition 40). If $\mathrm{rank}(\mathbf{G}^*) = d$, then clearly $\mathbf{G}^*$ is globally optimum for REG-SDP as well, meaning that we have solved the nonconvex problem by solving its convex relaxation. That is, the convex relaxation is *tight*. Empirically, we notice the well-known phenomena of phase transition for convex relaxations (e.g. see [16, 17]), where below a certain noise threshold in the data, the relaxation remains tight (see Fig. 3.1). In case of C-SDP, the data is the matrix $\mathbf{C}$, as it encodes the information about pairwise relations among the M orthogonal transforms we are to estimate. Data matrix $\mathbf{C}$ ultimately depends on the local coordinates $\mathbf{x}_{k,i}$. If these local coordinate measurements are exact, we say that $\mathbf{C}$ is *clean*; otherwise, we say that $\mathbf{C}$ is noisy.

3.3 Tightness of Convex Relaxation

To begin with, let us clarify what we want to prove in order to explain the tightness behavior of C-SDP. Denote by $\mathbf{C}_0$ the clean data matrix, i.e. the data matrix for the case when the local coordinate measurements are exact. Since the local coordinate measurements are clean, we know that the ground-truth rigid transforms relating patch coordinate systems to the global coordinate system solve the registration problem. Let the local coordinate system for patch P_i be related to the global coordinate system via rigid transform $(\bar{\mathbf{O}}_i, \bar{\mathbf{t}}_i)$. Let $\mathbf{G}_0$ be defined by $[\mathbf{G}_0]_{ij} = \bar{\mathbf{O}}_i^\top \bar{\mathbf{O}}_j$, for $i, j \in [1:M]$. We will see (Proposition 39) that $\mathbf{G}_0$ is a global minimizer for C-SDP when the data matrix is clean. Moreover, $\mathrm{rank}(\mathbf{G}_0) = d$. Thus, C-SDP is tight when the data matrix is clean.

Now, suppose the local coordinate measurements are noisy. Denote the corresponding

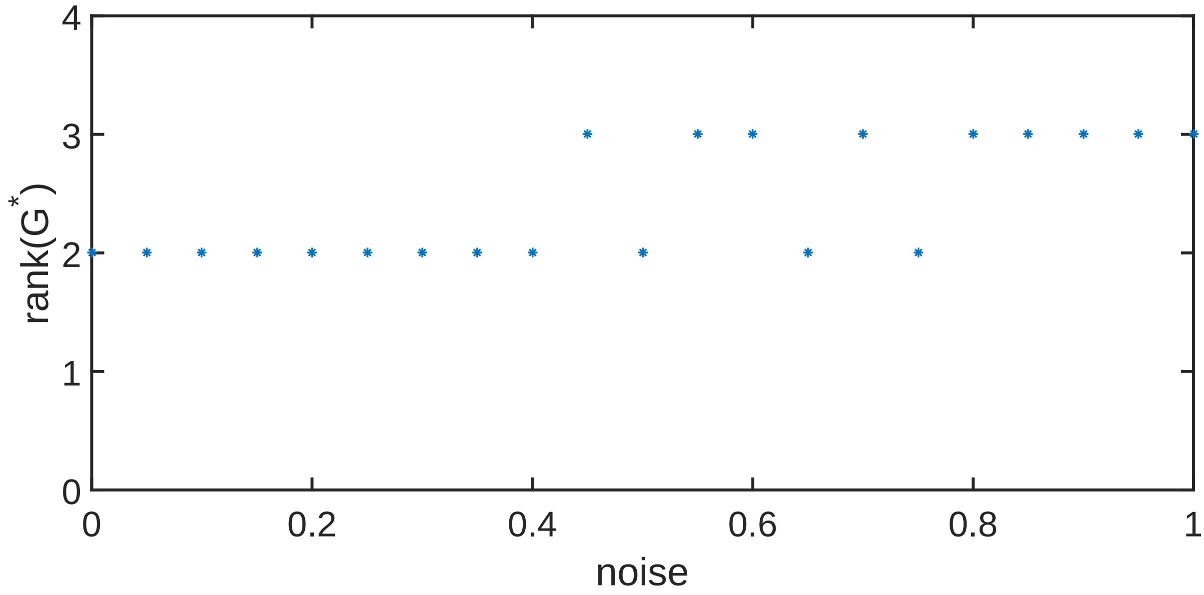

Figure 3.1: Phase transition for tightness of C-SDP ($d = 2$). The plot shows the rank of the global optimum $\mathbf{G}^*$ as a function of the noise level in the data $\mathbf{C}$. Below a certain noise threshold, the rank of $\mathbf{G}^*$ is exactly 2, i.e., the relaxation is tight. Above this threshold, the rank of $\mathbf{G}^*$ exceeds 2, making it infeasible for REG-SDP.

data matrix by $\mathbf{C} = \mathbf{C}_0 + \mathbf{W}$, where the matrix $\mathbf{W}$ is the perturbation because of the noisy measurements. With the noisy data matrix $\mathbf{C}$ ($= \mathbf{C}_0 + \mathbf{W}$), let $\mathbf{G}^*$ be global minimizer for C-SDP. To explain the tightness behavior of C-SDP, we want to prove that $\text{rank}(\mathbf{G}^*) = d$ when the "magnitude" of noise $\mathbf{W}$ is not too large. In short, perturbing the data matrix, $\mathbf{C}_0 \to \mathbf{C}$, perturbs the optimum of C-SDP, $\mathbf{G}_0 \to \mathbf{G}^*$, and we want to show that $\text{rank}(\mathbf{G}^*) = \text{rank}(\mathbf{G}_0) = d$ if the perturbation in the data matrix is not too large.

Stated as above, it is not very clear how to proceed ahead. This is where we bring KKT conditions [45] into play: Loosely speaking, KKT conditions state the necessary and sufficient conditions that must hold at optimum of a convex program, provided some regularity condition is satisfied. Before clarifying how using KKT conditions helps us explain tightness behavior of C-SDP, we explicitly state the KKT conditions for C-SDP in the following lemma.

Lemma 35. $\mathbf{G}^*$ *is a KKT point of C-SDP if there exists a block-diagonal matrix* $\mathbf{\Lambda}^* \in \mathbb{S}^{Md}$ *such that*

(*a*) $\mathbf{G}^* \in \mathbb{S}^{Md}_+$, $[\mathbf{G}^*]_{ii} = \mathbf{I}_d$,

(*b*) $\mathbf{C} + \mathbf{\Lambda}^* \in \mathbb{S}^{Md}_+$, *and*

(*c*) $(\mathbf{C} + \mathbf{\Lambda}^*)\, \mathbf{G}^* = \mathbf{0}$.

(We call $\mathbf{\Lambda}^*$ *the dual variable corresponding to* $\mathbf{G}^*$*.)*

We defer the proof of this lemma to the Appendix of this chapter. Observe that the identity matrix $\mathbf{I}_{Md}$ is strictly feasible for C-SDP. Thus, by Slater's condition [45], the KKT conditions in Lemma 35 are both necessary and sufficient for $\mathbf{G}^*$ to be a global minimizer of C-SDP [45]. Note that we refer to $\mathbf{\Lambda}^*$ as *the* dual variable corresponding to $\mathbf{G}^*$: this is because, given a global minimizer $\mathbf{G}^*$, there is a unique block-diagonal matrix $\mathbf{\Lambda}^*$ satisying conditions in Lemma 35 (see Proposition 36).

The KKT conditions in Lemma 35 give us an indirect way of proving $\mathrm{rank}(\mathbf{G}^*) = d$. Let $\mathbf{\Lambda}^*$ be as in Lemma 35. Then, $\mathrm{nullity}(\mathbf{C}+\mathbf{\Lambda}^*) \geq d$. This is because: (i) $(\mathbf{C}+\mathbf{\Lambda}^*)\mathbf{G}^* = \mathbf{0}$ is equivalent to the fact that the columns of $\mathbf{G}^*$ lie in the null space of $\mathbf{C} + \mathbf{\Lambda}^*$, and (ii) $\mathrm{rank}(\mathbf{G}^*) \geq d$ is equivalent to the fact that the column space of $\mathbf{G}^*$ has dimension at least d. In fact, this also shows that

$$\mathrm{nullity}(\mathbf{C} + \mathbf{\Lambda}^*) = d \implies \mathrm{rank}(\mathbf{G}^*) = d. \tag{3.6}$$

In other words, to prove that $\mathrm{rank}(\mathbf{G}^*) = d$, it is sufficient to prove that $\mathrm{nullity}(\mathbf{C} + \mathbf{\Lambda}^*) = d$, where $\mathbf{\Lambda}^*$ is the dual variable corresponding to $\mathbf{G}^*$. The pertinent question now is: How do we find a candidate for $\mathbf{\Lambda}^*$? Fortunately, the KKT conditions in Lemma 35 provide an answer: Any block-diagonal matrix $\mathbf{\Lambda}^*$ satisfying conditions in Lemma 35 can be expressed in terms of $\mathbf{G}^*$; the following proposition gives the explicit dependence of $\mathbf{\Lambda}^*$ on $\mathbf{G}^*$ (we defer the proof to the Appendix of this chapter).

Proposition 36. *Let $\mathbf{G}^*$ and $\mathbf{\Lambda}^*$ be as in Lemma 35. Then*

$$\mathbf{\Lambda}^* = -\mathrm{bd}(\mathbf{C}\mathbf{G}^*), \tag{3.7}$$

where, bd *is the linear operator that leaves the diagonal blocks untouched and sets other elements to* 0*.*

3.3.1 Main Result

Before stating the main result of this chapter, we note a standard result from matrix analysis – Weyl's theorem [46] – which quantifies perturbation in the eigenvalues of a (symmetric) matrix in terms of the perturbation in the matrix.

Theorem 37 (Weyl). *Let $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{n\times n}$ be symmetric. Then for $j \in [1:n]$,*

$$\lambda_j(\mathbf{X}) - \|\mathbf{Y}\|_2 \leq \lambda_j(\mathbf{X}+\mathbf{Y}) \leq \lambda_j(\mathbf{X}) + \|\mathbf{Y}\|_2.$$

We now state the main result of this chapter concerning the tightness of C-SDP, and see how the proof of this result exploits duality. Recall: $\mathbf{C}_0$ is the clean data matrix and $\mathbf{G}_0$ is the Gram matrix of the ground-truth rigid transforms ($\mathbf{G}_0$ globally minimizes C-SDP for the clean data matrix, see Proposition 39); $\mathbf{C}$ is the noisy data matrix and $\mathbf{G}^*$ is global minimizer for C-SDP corresponding to the noisy data matrix.

Theorem 38. *Suppose* $\text{nullity}(\mathbf{C}_0) = d$. *Let* $\mathbf{C} = \mathbf{C}_0 + \mathbf{W}$ *be the noisy data matrix. Then there exists some* $\eta > 0$ *such that* $\text{rank}(\mathbf{G}^*) = d$ *(i.e.* C-SDP *is tight) if* $\|\mathbf{W}\| < \eta$.

Proof. Let $\mathbf{\Lambda}^*$ be the dual corresponding to $\mathbf{G}^*$. From our discussion preceding equation (3.6), we know that $\text{nullity}(\mathbf{C} + \mathbf{\Lambda}^*) \geq d$. Moreover, equation (3.6) says that in order to prove $\text{rank}(\mathbf{G}^*) = d$, it is sufficient to prove $\text{nullity}(\mathbf{C} + \mathbf{\Lambda}^*) = d$. In other words, it is enough to prove that $\lambda_{d+1}(\mathbf{C} + \mathbf{\Lambda}^*) > 0$. From Proposition 36, we have that

$$\begin{aligned} \mathbf{C} + \mathbf{\Lambda}^* &= \mathbf{C} - \text{bd}(\mathbf{C}\mathbf{G}^*) \\ &= \mathbf{C}_0 + \mathbf{W} - \text{bd}((\mathbf{C}_0 + \mathbf{W})\mathbf{G}^*) \\ &= \mathbf{C}_0 + \mathbf{W} - \text{bd}(\mathbf{C}_0\mathbf{G}^*) - \text{bd}(\mathbf{W}\mathbf{G}^*) \\ &= \mathbf{C}_0 + \underbrace{\mathbf{W} - \text{bd}(\mathbf{C}_0(\mathbf{G}^* - \mathbf{G}_0)) - \text{bd}(\mathbf{W}\mathbf{G}^*)}_{\mathbf{A}} \end{aligned}$$

To get the last equality, we have added a superfluous term $\text{bd}(\mathbf{C}_0\mathbf{G}_0)$. We can do this because, as we will show in the next subsection (equation (3.10)), $\mathbf{C}_0\mathbf{G}_0 = \mathbf{0}$. Thus, in summary, $\mathbf{C} + \mathbf{\Lambda}^* = \mathbf{C}_0 + \mathbf{A}$, where $\mathbf{A} = \mathbf{W} - \text{bd}(\mathbf{W}\mathbf{G}^*) - \text{bd}(\mathbf{C}_0\mathbf{\Delta})$, and $\mathbf{\Delta} = \mathbf{G}^* - \mathbf{G}_0$ is the perturbation in the optimum of C-SDP due to pertubation in the data matrix. Now, by Weyl's theorem,

$$\lambda_{d+1}(\mathbf{C} + \mathbf{\Lambda}^*) \geq \lambda_{d+1}(\mathbf{C}_0) - \|\mathbf{A}\|_2. \tag{3.8}$$

So, if we want $\lambda_{d+1}(\mathbf{C} + \mathbf{\Lambda}^*) > 0$, we need to upperbound $\|\mathbf{A}\|_2$.

$$\begin{aligned} \|\mathbf{A}\|_2 &= \|\mathbf{W} - \text{bd}(\mathbf{W}\mathbf{G}^*) - \text{bd}(\mathbf{C}_0\mathbf{\Delta})\|_2 \\ &\leq \|\mathbf{W}\| + \|\text{bd}(\mathbf{W}\mathbf{G}^*)\| + \|\text{bd}(\mathbf{C}_0\mathbf{\Delta})\| \\ &\leq \|\mathbf{W}\| + \|\mathbf{W}\mathbf{G}^*\| + \|\mathbf{C}_0\mathbf{\Delta}\| \\ &\leq \|\mathbf{W}\|\,(1 + \|\mathbf{G}^*\|) + \|\mathbf{C}_0\|\,\|\mathbf{\Delta}\|, \end{aligned}$$

where we have used the results that $\|\cdot\|_2 \leq \|\cdot\|$, that $\|\cdot\|$ obeys triangle inequality, that $\|\text{bd}(\cdot)\| \leq \|\cdot\|$, and that $\|\cdot\|$ is sub-multiplicative [46]. We now need an upper bound on $\|\mathbf{\Delta}\|$, which is the magnitude of perturbation in the optimum of C-SDP when the data matrix is perturbed from $\mathbf{C}_0$ to $\mathbf{C} = \mathbf{C}_0 + \mathbf{W}$. We do this in Lemma 41, where we prove that

$$\|\mathbf{\Delta}\| \leq \frac{2M}{\lambda_{d+1}(\mathbf{C}_0)}\|\mathbf{W}\|.$$

Thus, we have

$$\begin{aligned} \|\mathbf{A}\|_2 &\leq \|\mathbf{W}\| \left(1 + \|\mathbf{G}^*\| + \frac{2M}{\lambda_{d+1}(\mathbf{C}_0)}\|\mathbf{C}_0\|\right) \\ &\leq \|\mathbf{W}\| \left(1 + M\sqrt{d} + \frac{2M}{\lambda_{d+1}(\mathbf{C}_0)}\|\mathbf{C}_0\|\right), \end{aligned}$$

where, the fact that $\|\mathbf{G}^*\| \le M\sqrt{d}$ follows from Proposition 40(b). The terms in the paranthesis are constant. Thus, by controlling $\|\mathbf{W}\|$, we can control $\|\mathbf{A}\|_2$. In particular, if

$$\|\mathbf{W}\| < \eta = \frac{\lambda_{d+1}(\mathbf{C}_0)}{\left(1 + M\sqrt{d} + \frac{2M}{\lambda_{d+1}(\mathbf{C}_0)}\|\mathbf{C}_0\|\right)},$$

then $\lambda_{d+1}(\mathbf{C}_0) - \|\mathbf{A}\|_2 > 0$. This, from equation (3.8), implies that $\lambda_{d+1}(\mathbf{C} + \mathbf{\Lambda}^*) > 0$. □

In the following subsections, we tie some loose ends in the proof above. In particular: (*i*) we prove that $\mathbf{C}_0\mathbf{G}_0 = \mathbf{0}$; (*ii*) we justify the assumption made in the main theorem regarding nullity of the clean data matrix, nullity$(\mathbf{C}_0) = d$; (*iii*) we prove stability of C-SDP.

3.3.2 Clean Case

Suppose the local coordinate measurements are exact (noiseless). Let $\bar{\mathbf{z}}_1, \cdots, \bar{\mathbf{z}}_N$ be the ground-truth global coordinates of the N points. Let the ground-truth rigid transform relating the local coordinate system of patch P_i to the global coordinate system be $(\bar{\mathbf{O}}_i, \bar{\mathbf{t}}_i)$, $i \in [1:M]$. Since the local coordinate measurements are noiseless, the ground-truth global coordinates and rigid transforms solve the rigid registration problem exactly. That is,

$$\sum_{i=1}^{M}\sum_{k\in P_i}\|\bar{\mathbf{z}}_k - (\bar{\mathbf{O}}_i\mathbf{x}_{k,i} + \bar{\mathbf{t}}_i)\|^2 = 0 = \mathrm{Tr}\left([\mathbf{C}_0]_{ij}\bar{\mathbf{O}}_j^\top\bar{\mathbf{O}}_i\right) \tag{3.9}$$

Recall that $\mathbf{C}_0$ is the clean data matrix. Also recall that we defined $\mathbf{G}_0$ by $[\mathbf{G}_0]_{ij} = \bar{\mathbf{O}}_i^\top\bar{\mathbf{O}}_j$, for $i, j \in [1:M]$. Now, from (3.9) and the fact that the objective function in REG-SDP is the exact reformulation of the objective function in ORTHO-REG, we get that $\mathrm{Tr}(\mathbf{C}_0\mathbf{G}_0) = 0$. Moreover, since $\mathbf{C}_0 \in \mathbb{S}_+^{Md}$, and trace of product of two positive semidefinite matrices is always nonnegative, we have that $\mathrm{Tr}(\mathbf{C}_0\mathbf{G}) \ge 0 \ \forall\ \mathbf{G} \in \mathbb{S}_+^{Md}$. In particular, since any $\mathbf{G}$ feasible for C-SDP is positive semidefinite, and $\mathrm{Tr}(\mathbf{C}_0\mathbf{G}_0) = 0$, we have proved the following:

Proposition 39. $\mathbf{G}_0$ *is global minimizer for C-SDP when the data matrix is clean.*

Moreover, for arbitrary positive semidefinite matrices $\mathbf{X}$ and $\mathbf{Y}$, the following holds:

$$\mathrm{Tr}(\mathbf{XY}) = 0 \iff \mathbf{XY} = \mathbf{0}.$$

This, along with the fact that $\mathrm{Tr}(\mathbf{C}_0\mathbf{G}_0) = \mathbf{0}$ gives us

$$\mathbf{C}_0\mathbf{G}_0 = \mathbf{0}. \tag{3.10}$$

Assumption on $\mathbf{C}_0$: Since rank$(\mathbf{G}_0) = d$, we deduce from (3.10) that nullity$(\mathbf{C}_0) \ge d$. The assumption we make on $\mathbf{C}_0$ is that nullity$(\mathbf{C}_0) = d$. This is equivalent to the condition that the body graph (defined in Chapter 2) corresponding to the registration problem is *affinely rigid* [14]. For example, this is true when each patch contains all the N points. More generally, $\mathrm{nullity}(\mathbf{C}_0) = d$ when there exists an ordering of the patch indices such that P_1 contains at least $d + 1$ points, and P_i and $P_1 \cup P_2 \cup \cdots \cup P_{i-1}$ have at least $d + 1$ points in common for $i \geq 2$ [14] (this occurs naturally in applications like multiview registration [9]).

3.3.3 Stability of C-SDP

We now state a result that quantifies stability of C-SDP. More specifically, Lemma 41 upper-bounds the perturbation in global minimizer of C-SDP following a perturbation in the clean data matrix $\mathbf{C}_0$.

Before stating and proving Lemma 41, we note some properties of positive semidefinite matrices $\mathbf{G}$ with $[\mathbf{G}]_{ii} = \mathbf{I}_d$ in Proposition 40; we defer the proof of this proposition to the Appendix of this chapter.

Proposition 40. *Suppose* $\mathbf{G} \in \mathbb{S}_+^{Md}$ *and* $[\mathbf{G}]_{ii} = \mathbf{I}_d$. *Then*

(*a*) $\mathrm{rank}(\mathbf{G}) \geq d$;

(*b*) $\sigma_{\max}\left([\mathbf{G}]_{ij}\right) \leq 1, \;\; i, j \in [1:M]$;

(*c*) *If* $\mathrm{rank}(\mathbf{G}) = d$, *then any nonzero eigenvalue of* $\mathbf{G}$ *is* M.

Lemma 41. *Suppose* $\mathrm{nullity}(\mathbf{C}_0) = d$. *Let* $\mathbf{C}$ *be the corresponding noisy data matrix, and let* $\mathbf{W} = \mathbf{C} - \mathbf{C}_0$. *Let* $\mathbf{G}^*$ *be a global optimum for* C-SDP *corresponding to the noisy data* $\mathbf{C}$. *Let* $\Delta = \mathbf{G}^* - \mathbf{G}_0$, *where* $\mathbf{G}_0$ *is the ground-truth solution. Then*

$$\|\Delta\| \leq \frac{2M}{\lambda_{d+1}(\mathbf{C}_0)} \|\mathbf{W}\|.$$

Proof. Since $\mathrm{rank}(\mathbf{G}_0) = d$ we can write the spectral decomposition of the ground truth solution as (see Proposition 40(c))

$$\mathbf{G}_0 = \sum_{i=1}^{d} M \mathbf{s}_i \mathbf{s}_i^\top$$

where $\mathbf{s}_i, \; i \in [1:d]$ are orthonormal. Similarly, we can write the optimal solution corresponding to the perturbed data $\mathbf{C}$ as

$$\mathbf{G}^* = \sum_{i=1}^{Md} \alpha_i \mathbf{g}_i \mathbf{g}_i^\top$$

where $\mathbf{g}_i, \; i \in [1:d]$ are orthonormal.

Let $\mathbf{P}$ be the orthoprojector on kernel($\mathbf{C}_0$). Since $\mathbf{C}_0\mathbf{G}_0 = \mathbf{0}$, rank($\mathbf{G}_0$) $= d$, and nullity($\mathbf{C}_0$) $= d$, we deduce that

$$\mathbf{P} = \sum_{i=1}^{d} \mathbf{s}_i \mathbf{s}_i^\top = \frac{1}{M}\mathbf{G}_0.$$

Clearly, $\mathbf{R} = \mathbf{I} - \mathbf{P}$ is the orthoprojector on range($\mathbf{C}_0$), and

$$\mathbf{g}_i = \mathbf{P}\mathbf{g}_i + \mathbf{R}\mathbf{g}_i.$$

Let $\mathbf{h}_i := \mathbf{R}\mathbf{g}_i$. As will be apparent, we want to lowerbound $\langle \mathbf{C}_0, \mathbf{G}^* \rangle$. Now,

$$\begin{aligned} \langle \mathbf{C}_0, \mathbf{G}^* \rangle &= \sum_{i=1}^{Md} \alpha_i \left(\mathbf{g}_i^\top \mathbf{C}_0 \mathbf{g}_i\right) \\ &= \sum_{i=1}^{Md} \alpha_i \left(\mathbf{h}_i^\top \mathbf{C}_0 \mathbf{h}_i\right) \\ &\geq \lambda_{d+1}(\mathbf{C}_0) \sum_{i=1}^{Md} \alpha_i \, \|\mathbf{h}_i\|^2. \end{aligned} \tag{3.11}$$

Thus, to get a lowerbound on $\langle \mathbf{C}_0, \mathbf{G}^* \rangle$, we just have to lowerbound $\sum_{i=1}^{Md} \alpha_i \, \|\mathbf{h}_i\|^2$. To do this, we first reformulate $\|\mathbf{h}_i\|^2$ as,

$$\begin{aligned} \|\mathbf{h}_i\|^2 &= \|\mathbf{R}\mathbf{g}_i\|^2 \\ &= \mathbf{g}_i^\top \mathbf{R}\mathbf{g}_i \\ &= \mathbf{g}_i^\top (\mathbf{I} - \mathbf{P})\mathbf{g}_i \\ &= \|\mathbf{g}_i\|^2 - \frac{1}{M}(\mathbf{g}_i^\top \mathbf{G}_0 \mathbf{g}_i) \\ &= 1 - \frac{1}{M}\left(\mathbf{g}_i^\top \mathbf{G}_0 \mathbf{g}_i\right), \end{aligned}$$

where the second equality holds because $\mathbf{R}^2 = \mathbf{R}$, since $\mathbf{R}$ is an orthoprojector. Thus,

$$\begin{aligned} \sum_{i=1}^{Md} \alpha_i \|\mathbf{h}_i\|^2 &= \sum_{i=1}^{Md} \alpha_i - \frac{1}{M} \sum_{i=1}^{Md} \alpha_i \left(\mathbf{g}_i^\top \mathbf{G}_0 \mathbf{g}_i\right) \\ &= Md - \frac{1}{M}\langle \mathbf{G}_0, \mathbf{G}^* \rangle. \end{aligned} \tag{3.12}$$

Now, to upperbound $\langle \mathbf{G}_0, \mathbf{G}^* \rangle$, we note that

$$\begin{aligned} \langle \mathbf{G}_0, \mathbf{G}^* \rangle &= \frac{1}{2}\|\mathbf{G}_0\|^2 + \frac{1}{2}\|\mathbf{G}^*\|^2 - \frac{1}{2}\|\boldsymbol{\Delta}\|^2 \\ &\leq \frac{1}{2}M^2 d + \frac{1}{2}M^2 d - \frac{1}{2}\|\boldsymbol{\Delta}\|^2 = M^2 d - \frac{1}{2}\|\boldsymbol{\Delta}\|^2 \end{aligned} \tag{3.13}$$

where, to get the first inequality, we use the fact that any singular value of $[\mathbf{G}^*]_{ij}$ is at most 1 for every $1 \leq i, j \leq M$ (Proposition 40). Combining (3.12) and (3.13), we get

$$\sum_{i=1}^{Md} \alpha_i \|\mathbf{h}_i\|^2 \geq \frac{1}{2M}\|\boldsymbol{\Delta}\|^2.$$

Plugging this in (3.11) gives us,

$$\left\langle \mathbf{C}_0, \mathbf{G}^* \right\rangle \geq \frac{\lambda_{d+1}(\mathbf{C}_0)}{2M}\|\boldsymbol{\Delta}\|^2.$$

Now,

$$\begin{aligned} -\|\mathbf{W}\| \, \|\boldsymbol{\Delta}\| &\leq \left\langle \mathbf{W}, \boldsymbol{\Delta} \right\rangle \\ &= \left(\left\langle \mathbf{C}, \mathbf{G}^* \right\rangle - \left\langle \mathbf{C}, \mathbf{G}_0 \right\rangle\right) + \left(\left\langle \mathbf{C}_0, \mathbf{G}_0 \right\rangle - \left\langle \mathbf{C}_0, \mathbf{G}^* \right\rangle\right) \\ &\leq -\left\langle \mathbf{C}_0, \mathbf{G}^* \right\rangle \end{aligned}$$

where the first inequality is Cauchy-Schwarz (with inner product between two symmetric matrices defined as the trace of their product), and the last inequality is due to the fact that the term in first paranthesis is negative by optimality of $\mathbf{G}^*$, and that the first term in second paranthesis is 0. So,

$$\begin{aligned} \|\mathbf{W}\| \, \|\boldsymbol{\Delta}\| &\geq \left\langle \mathbf{C}_0, \mathbf{G}^* \right\rangle \\ &\geq \frac{\lambda_{d+1}(\mathbf{C}_0)}{2M}\|\boldsymbol{\Delta}\|^2 \end{aligned}$$

which finally gives us,

$$\|\boldsymbol{\Delta}\| \leq \frac{2M}{\lambda_{d+1}(\mathbf{C}_0)}\|\mathbf{W}\|.$$

□

3.4 Discussion

In this chapter, we derived a rank-constrained semidefinite program REG-SDP that isolates the computational difficulty of the rigid registration problem in the rank-constraint. We then derived a convex relaxation C-SDP by dropping the rank-constraint from REG-SDP. The main contribution of the chapter was to give a theoretical justification of the empirically-observed phenomenon that the relaxation remains tight when noise in the data is below a certain threshold. Along the way, we proved a stability result for C-SDP. In the next chapter, we will see how the tightness behavior of C-SDP sheds light on convergence behavior of an algorithm that, instead of working with the relaxation C-SDP, directly attacks the rank-constrained program REG-SDP.

3.5 Appendix

3.5.1 Proof of Lemma 35

KKT conditions for the convex program C-SDP are just the conditions of primal feasibility, dual feasibility, and complementary slackness [45]. Condition (a) is the primal feasibility condition, condition (b) is the dual feasibility condition, and condition (c) is the complementary slackness condition that follows from strong duality [47]. Now, to see why $\mathbf{\Lambda}^*$ is block-diagonal, we write C-SDP as a standard semidefinite program [47]

$$\begin{aligned} \min_{\mathbf{G}} \quad & \mathrm{Tr}(\mathbf{C}\mathbf{G}) \\ \text{s.t.} \quad & \mathrm{Tr}(\mathbf{A}_k\mathbf{G}) = b_k, \;\; k = [1:m], \\ & \mathbf{G} \in \mathbb{S}_+^{Md}. \end{aligned}$$

Here $\mathbf{A}_k \in \mathbb{S}^{Md}$, and $\mathrm{Tr}(\mathbf{A}_k\mathbf{G}) = b_k,\ k = [1:m]$, collectively encode the condition that $[\mathbf{G}]_{ii} = \mathbf{I}_d$. From [47], $\mathbf{\Lambda}^*$ is of the form $\sum_{k=1}^m y_k\mathbf{A}_k$, for some scalars y_k, and it is not difficult to see that $\sum_{k=1}^m y_k\mathbf{A}_k$ forms a symmetric block-diagonal matrix in this case.

3.5.2 Proof of Proposition 36

Proposition 35 tells us that $(\mathbf{C} + \mathbf{\Lambda}^*)\,\mathbf{G}^* = \mathbf{0}$. That is,

$$\mathbf{\Lambda}^*\mathbf{G}^* = -\mathbf{C}\mathbf{G}^*.$$

Proposition 35 also tells us that $\mathbf{\Lambda}^*$ is a symmetric block-diagonal matrix, and that $[\mathbf{G}^*]_{ii} = \mathbf{I}_d$. Using these facts, and comparing the diagonal blocks of the left-hand side and the right-hand side, we get

$$[\mathbf{\Lambda}^*]_{ii} = -[\mathbf{C}\mathbf{G}^*]_{ii}.$$

Thus,

$$\mathbf{\Lambda}^* = -\mathrm{bd}(\mathbf{C}\mathbf{G}^*),$$

where, bd is the linear operator that leaves the diagonal blocks untouched and sets other elements to 0.

3.5.3 Proof of Proposition 40

(a) Consider the first d columns of $\mathbf{G}$. Since $[\mathbf{G}]_{11} = \mathbf{I}_d$, we conclude that the first d columns of $\mathbf{G}$ are linearly independent.

(b) Let $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^d$, $\|\mathbf{x}_i\| = \|\mathbf{x}_j\| = 1$. Regard any vector in $\mathbb{R}^{Md}$ as consisting of M vectors in $\mathbb{R}^d$ stacked vertically. Construct $\mathbf{x} \in \mathbb{R}^{Md}$ with i-th block as $\mathbf{x}_i$, j-th block as $-\mathbf{x}_j$, and every other block as $\mathbf{0}$. Then

$$\mathbf{x}^\top \mathbf{G}\mathbf{x} = \|\mathbf{x}_i\|^2 + \|\mathbf{x}_j\|^2 - 2\mathbf{x}_i^\top [\mathbf{G}]_{ij}\mathbf{x}_j,$$

where we have used the fact that $[\mathbf{G}]_{ji} = [\mathbf{G}]_{ij}^\top$. Since $\mathbf{G} \in \mathbb{S}_+^{Md}$, we get that

$$\|\mathbf{x}_i\|^2 + \|\mathbf{x}_j\|^2 - 2\mathbf{x}_i^\top [\mathbf{G}]_{ij}\mathbf{x}_j \geq 0.$$

This gives us

$$\mathbf{x}_i^\top [\mathbf{G}]_{ij}\mathbf{x}_j \leq \frac{\|\mathbf{x}_i\|^2 + \|\mathbf{x}_j\|^2}{2} = 1.$$

Similarly, by replacing $-\mathbf{x}_j$ with $\mathbf{x}_j$ in the j-th block of $\mathbf{x}$, we get

$$\mathbf{x}_i^\top [\mathbf{G}]_{ij}\mathbf{x}_j \geq -1.$$

Putting these together, we have

$$\left|\mathbf{x}_i^\top [\mathbf{G}]_{ij}\mathbf{x}_j\right| \leq 1.$$

Unit vectors $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^{Md}$ were arbitrary, and thus, the result follows.

(c) Consider the spectral decomposition of $\mathbf{G}$,

$$\mathbf{G} = \sum_{i=1}^{d} \alpha_i \mathbf{v}_i \mathbf{v}_i^\top$$

where $\alpha_i > 0$ are the non-zero eigenvalues, and $\mathbf{v}_i$ are the corresponding orthonormal eigenvectors. Let

$$\mathbf{B} = \left[\sqrt{\alpha_1}\mathbf{v}_1 \cdots \sqrt{\alpha_d}\mathbf{v}_d\right]^\top.$$

Regard $\mathbf{B}$ as a block-row, where each element is of size $d \times d$. Notationally, $\mathbf{B} = [\mathbf{B}_1 \cdots \mathbf{B}_M]$, where $\mathbf{B}_i \in \mathbb{R}^{d\times d}$. Thus,

$$[\mathbf{G}]_{ij} = \mathbf{B}_i^\top \mathbf{B}_j.$$

In particular, since $[\mathbf{G}]_{ii} = \mathbf{I}_d$, we get that

$$\mathbf{B}_i^\top \mathbf{B}_i = \mathbf{I}_d.$$

Now, suppose $\mathbf{v}_1 = (\mathbf{v}_{11}, \cdots, \mathbf{v}_{1M})$, where $\mathbf{v}_{1j} \in \mathbb{R}^d$. Note that $\sqrt{\alpha_1}\mathbf{v}_{1j}^\top$ forms the first row of $\mathbf{B}_j$. From orthogonality of $\mathbf{B}_j$, we get that, for every j,

$$\|\sqrt{\alpha_1}\mathbf{v}_{1j}\|^2 = 1.$$

Or,

$$\|\mathbf{v}_{1j}\|^2 = \frac{1}{\alpha_1}, \quad j \in [1 : M].$$

Now, since $\|\mathbf{v}_1\|^2 = 1$, we have

$$\sum_{j=1}^{M} \|\mathbf{v}_{1j}\|^2 = 1,$$

from where we get that $\alpha_1 = M$. Similarly, $\alpha_2 = \cdots = \alpha_d = M$.

Chapter 4

Convergence of Nonconvex ADMM

4.1 Introduction

In the last chapter, we saw that the least-squares estimator LS-REG for the rigid registration problem can be formulated as a rank-constrained semidefinite program REG-SDP, which can then be relaxed into a convex semidefinite program C-SDP by dropping the rank-constraint. When the relaxation C-SDP is not tight, we have to resort to some kind of "rounding" to obtain a solution feasible for the rank-constrained program REG-SDP [14]. This motivates the question: Could we come up with a theoretically-sound solver that directly attacks the nonconvex problem REG-SDP?

In a recent work [9], it was shown that by formally applying the alternating direction method of multipliers (ADMM), we can derive a computationally-efficient iterative solver REG-ADMM to solve REG-SDP. The solver was empirically shown to have robust performance in the context of multiview registration. Unlike convex programs, convergence analysis of ADMM for nonconvex programs is still in its infancy, and the existing handful of results cannot be applied to derive theoretical convergence of REG-ADMM. In this chapter, we investigate convergence properties of REG-ADMM and our main findings are as follows. We prove that if the REG-ADMM iterates converge, they do so to a stationary (KKT) point of REG-SDP. Moreover, for clean measurements, we give an explicit formula for ρ, for which REG-ADMM is guaranteed to converge to the global optimum (with arbitrary initializations). If the noise is low, we can still show that the iterates converge to the global optimum, provided they are initialized sufficiently close to the optimum. On the other hand, if the noise is high, we point out that the iterates can oscillate if ρ is less than some threshold, irrespective of the initialization. We present simulation results to support our theoretical predictions.

4.1.1 Alternating Direction Method of Multipliers

We start with a brief overview of the alternating direction method of multipliers (ADMM). A more detailed exposition can be found in [18]. ADMM is an iterative method which can be applied to solve a convex program of the following form:

$$\text{(CVX)} \qquad \begin{aligned} \min_{\mathbf{x} \in \mathcal{P},\, \mathbf{z} \in \mathcal{Q}} \quad & f(\mathbf{x}) + g(\mathbf{z}) \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z} - \mathbf{c} = \mathbf{0}, \end{aligned}$$

where, $\mathcal{P}$ and $\mathcal{Q}$ are closed, convex sets in $\mathbb{R}^n$ and $\mathbb{R}^m$; f and g are closed, proper, convex functions; $\mathbf{A} \in \mathbb{R}^{p\times n}$, $\mathbf{B} \in \mathbb{R}^{p\times m}$ and $\mathbf{c} \in \mathbb{R}^p$. For a fixed $\rho > 0$, define the *augmented* Lagrangian function [18] for CVX as

$$\mathcal{L}_\rho(\mathbf{x}, \mathbf{z}, \boldsymbol{\lambda}) = f(\mathbf{x}) + g(\mathbf{z}) + \left\langle \boldsymbol{\lambda}, \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z} - \mathbf{c} \right\rangle + \frac{\rho}{2}\|\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z} - \mathbf{c}\|^2, \tag{4.1}$$

where, $\boldsymbol{\lambda} \in \mathbb{R}^p$ is the Lagrange multiplier (a.k.a. dual variable) corresponding to the equality constraint, $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product, and $\|\cdot\|$ denotes the Euclidean vector norm. The ADMM algorithm, initialized with $(\mathbf{z}^0, \boldsymbol{\lambda}^0)$, performs the following updates for $k \geq 0$:

$$\begin{aligned} \mathbf{x}^{k+1} &= \operatorname*{argmin}_{\mathbf{x} \in \mathcal{P}} \mathcal{L}_\rho(\mathbf{x}, \mathbf{z}^k, \boldsymbol{\lambda}^k); \\ \mathbf{z}^{k+1} &= \operatorname*{argmin}_{\mathbf{z} \in \mathcal{Q}} \mathcal{L}_\rho(\mathbf{x}^{k+1}, \mathbf{z}, \boldsymbol{\lambda}^k); \\ \boldsymbol{\lambda}^{k+1} &= \boldsymbol{\lambda}^k + \rho\,(\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z} - \mathbf{c}). \end{aligned} \tag{4.2}$$

This algorithm is especially relevant when the subproblems involving minimization of the augmented Lagrangian with respect to a primal variable (while keeping the other variables fixed) have a simple closed-form solution. The ADMM algorithm can be shown to converge to global optimum of the convex program CVX under very mild assumptions [18].

4.1.2 ADMM for Rigid Registration

We now derive the ADMM-based algorithm proposed in [9] to solve the following problem:

$$\text{(REG-SDP)} \qquad \begin{aligned} \min_{\mathbf{G} \in \mathbb{S}_+^{Md}} \quad & \operatorname{Tr}(\mathbf{C}\mathbf{G}) \\ \text{subject to} \quad & [\mathbf{G}]_{ii} = \mathbf{I}_d, \quad i \in [1:M], \\ & \operatorname{rank}(\mathbf{G}) = d. \end{aligned}$$

In its original formulation, it is not clear how to use ADMM to solve REG-SDP. To remedy this, we "split" the variable $\mathbf{G}$ and express REG-SDP as:

$$\begin{aligned} \min_{\mathbf{G}, \mathbf{H} \in \mathbb{S}^{Md}} \quad & \operatorname{Tr}(\mathbf{C}\mathbf{G}) \\ \text{subject to} \quad & \mathbf{G} \in \Omega,\ \mathbf{H} \in \Theta, \\ & \mathbf{G} - \mathbf{H} = \mathbf{0}. \end{aligned} \tag{4.3}$$

Here, Ω denotes the (closed, *nonconvex*) set of positive semidefinite matrices whose rank is at most d, and Θ denotes the (closed, convex) set of symmetric matrices whose $d \times d$ diagonal blocks are $\mathbf{I}_d$. Notationally, $\Omega = \{\mathbf{X} \in \mathbb{S}_+^{Md} : \text{rank}(\mathbf{X}) \le d\}$, and $\Theta = \{\mathbf{X} \in \mathbb{S}^{Md} : [\mathbf{X}]_{ii} = \mathbf{I}_d\}$. The equivalence of this reformulation to REG-SDP follows from the fact that the rank of any $\mathbf{H} \in \Theta$ is at least d, since its diagonal blocks are $\mathbf{I}_d$. On the other hand, any $\mathbf{G} \in \Omega$ has rank at most d. Thus, if $\mathbf{H} = \mathbf{G}$, the rank of $\mathbf{G}$ (and $\mathbf{H}$) would have to be d. The reason we define Ω to contain matrices of rank *at most d* rather than matrices of rank exactly d is to ensure that Ω is a closed set; as will become clear, closedness of Ω is needed to ensure existence of (Euclidean) projection on Ω (see equation (4.5)). The splitting of the constraint into Ω and Θ, with an additional linear (consensus) constraint, makes the problem algorithmically amenable to ADMM. For some fixed $\rho > 0$, the augmented Lagrangian for (4.3) is

$$\mathcal{L}_\rho(\mathbf{G}, \mathbf{H}, \mathbf{\Lambda}) = \text{Tr}\,(\mathbf{C}\mathbf{G}) + \text{Tr}(\mathbf{\Lambda}(\mathbf{G} - \mathbf{H})) + \frac{\rho}{2}\|\mathbf{G} - \mathbf{H}\|^2, \tag{4.4}$$

where the symmetric matrix $\mathbf{\Lambda} \in \mathbb{R}^{Md \times Md}$ is the dual variable for the constraint $\mathbf{G} - \mathbf{H} = \mathbf{0}$, and $\|\cdot\|$ denotes the matrix Frobenius norm. The ADMM algorithm in [9], initialized with some $\mathbf{H}^0$ and $\mathbf{\Lambda}^0$, involves the following updates for $k \ge 0$:

$$\text{(REG-ADMM)} \qquad \begin{aligned} \mathbf{G}^{k+1} &= \underset{\mathbf{G} \in \Omega}{\text{argmin}}\ \mathcal{L}_\rho\left(\mathbf{G}, \mathbf{H}^k, \mathbf{\Lambda}^k\right); \\ \mathbf{H}^{k+1} &= \underset{\mathbf{H} \in \Theta}{\text{argmin}}\ \mathcal{L}_\rho\left(\mathbf{G}^{k+1}, \mathbf{H}, \mathbf{\Lambda}^k\right); \\ \mathbf{\Lambda}^{k+1} &= \mathbf{\Lambda}^k + \rho\left(\mathbf{G}^{k+1} - \mathbf{H}^{k+1}\right). \end{aligned}$$

As observed in [9], the first two sub-problems can be expressed as matrix projections onto Θ and Ω respectively. Namely,

$$\begin{aligned} \mathbf{G}^{k+1} &= \Pi_\Omega\left(\mathbf{H}^k - \rho^{-1}\left(\mathbf{C} + \mathbf{\Lambda}^k\right)\right), \\ \mathbf{H}^{k+1} &= \Pi_\Theta\left(\mathbf{G}^k + \rho^{-1}\mathbf{\Lambda}^k\right), \end{aligned} \tag{4.5}$$

where $\Pi_\Omega(\cdot)$ and $\Pi_\Theta(\cdot)$ denotes Euclidean projection [1] onto Ω and Θ. The former can be computed simply by setting the diagonal blocks of the input matrix to $\mathbf{I}_d$. On the other hand, the latter can be computed efficiently by computing its top d eigenvalues and retaining the ones that are positive [48]. More precisely, if $\mathbf{A}$ has a spectral decomposition $\mathbf{A} = \mu_1 \mathbf{u}_1 \mathbf{u}_1^\top +$

[1] Euclidean projection $\Pi_\Gamma(\mathbf{A})$ of a matrix $\mathbf{A} \in \mathbb{S}^n$ on a set $\Gamma \subset \mathbb{S}^n$ is defined as

$$\Pi_\Gamma(\mathbf{A}) = \underset{\mathbf{X} \in \Gamma}{\text{argmin}}\ \|\mathbf{X} - \mathbf{A}\|^2$$

$\cdots + \mu_{Md}\mathbf{u}_{Md}\mathbf{u}_{Md}^\top$, where $\mu_1 \geq \cdots \geq \mu_{Md}$, then

$$\Pi_\Omega(\mathbf{A}) = \sum_{i=1}^{d} \max(\mu_i, 0)\, \mathbf{u}_i\mathbf{u}_i^\top.$$

Note that even though nonconvexity of Ω places REG-ADMM outside the scope of the results on standard ADMM, we empirically find that it has good convergence properties. It was shown in [9] that the performance of REG-ADMM is comparable to state-of-the-art methods for the registration of three-dimensional multiview scans. In fact, this is consistent with recent works where the empirical success of ADMM for nonconvex problems have been reported [19, 20, 21]. Our goal in this chapter is to study theoretical convergence of REG-ADMM. For some of our results, we will be leveraging convergence properties of a closely related ADMM that solves the following convex relaxation of REG-SDP obtained by dropping the rank-constraint:

$$\text{(C-SDP)} \qquad \begin{aligned} &\min_{\mathbf{G}\in\mathbb{S}_+^{Md}} && \mathrm{Tr}(\mathbf{CG}) \\ &\text{subject to} && [\mathbf{G}]_{ii} = \mathbf{I}_d, \quad i \in [1:M]. \end{aligned}$$

An ADMM algorithm to solve C-SDP was proposed in [49]. Initialized with some $\mathbf{H}^0$ and $\boldsymbol{\Lambda}^0$, this involves the following updates for $k \geq 0$:

$$\text{(C-ADMM)} \qquad \begin{aligned} \mathbf{G}^{k+1} &= \underset{\mathbf{G}\in\mathbb{S}_+^{Md}}{\operatorname{argmin}}\ \mathcal{L}_\rho\left(\mathbf{G}, \mathbf{H}^k, \boldsymbol{\Lambda}^k\right); \\ \mathbf{H}^{k+1} &= \underset{\mathbf{H}\in\Theta}{\operatorname{argmin}}\ \mathcal{L}_\rho\left(\mathbf{G}^{k+1}, \mathbf{H}, \boldsymbol{\Lambda}^k\right); \\ \boldsymbol{\Lambda}^{k+1} &= \boldsymbol{\Lambda}^k + \rho\left(\mathbf{G}^{k+1} - \mathbf{H}^{k+1}\right). \end{aligned}$$

The expression for the augmented Lagrangian $\mathcal{L}_\rho$ in C-ADMM is identical to that in (4.4). Note that C-ADMM is similar to REG-ADMM, with the only difference being in the $\mathbf{G}$-update step. Namely, the minimization is over the nonconvex set Ω in REG-SDP, while it is over the closed convex set $\mathbb{S}_+^{Md}$ in C-SDP. This difference turns out to be crucial for convergence: C-ADMM converges to global optimum of C-SDP for any $\rho > 0$, and with arbitrary initialization [49].

While C-ADMM enjoys strong convergence guarantees, it has typical drawbacks of a convex relaxation. First, the rank of the global optimum of C-SDP is not guaranteed to be d, i.e., it might not even be feasible for REG-SDP. If the rank is greater than d, we have to "round" the solution of C-SDP to a rank-d matrix [14], which will generally be suboptimal for REG-SDP. Second, because the $\mathbf{G}$-update requires us to optimize over the entire PSD cone $\mathbb{S}_+^{Md}$, C-ADMM requires full eigendecomposition of an $Md \times Md$ matrix at every iteration [49]. On the other hand, REG-ADMM directly attacks the nonconvex problem REG-SDP, and thus

C-ADMM	REG-ADMM
$\mathbf{G}^{k+1} = \Pi_{\mathbb{S}_+^{Md}}\left(\mathbf{H}^k - \frac{1}{\rho}\left(\mathbf{C} + \mathbf{\Lambda}^k\right)\right)$	$\mathbf{G}^{k+1} = \Pi_{\Omega}\left(\mathbf{H}^k - \frac{1}{\rho}\left(\mathbf{C} + \mathbf{\Lambda}^k\right)\right)$
Compute Md eigenvalues ($Md \sim$ thousands)	Compute top d eigenvalues ($d \sim 2, 3$)
Converges to global minimum for *any* $\rho > 0$, with *any* initialization	Convergence behavior dependent on data noise, initialization, ρ (see Fig. 4.1)

Table 4.1: Table comparing C-ADMM (for solving C-SDP) and REG-ADMM (for solving REG-SDP). The only difference in the algorithms is the $\mathbf{G}$-update step, as observed in the first row of the table. The second row notes the saving in computation that results from using REG-ADMM to address the unrelaxed problem directly. The third row notes the contrast in the convergence behavior.

obviates the need for any rounding. Moreover, REG-ADMM requires computation of only the top d eigenvectors at each iteration. But these benefits come at a cost—it is usually difficult to derive theoretical guarantees for nonconvex optimization. As far as we know, convergence guarantees for REG-ADMM do not follow from existing results on nonconvex ADMM (see the discussion in Section 4.1.5 for further details).

4.1.3 Numerical Experiments

To understand the challenges involved in the convergence analysis of REG-ADMM, we look at some simulation results for the sensor network localization problem. We consider a two-dimensional network with ten nodes. There are three patches, where each patch contains all the ten nodes. The local coordinate system for each patch is obtained by arbitrarily rotating the global coordinate system. Moreover, we perturb the local coordinates using iid Gaussian noise. Our goal is to estimate the patch rotations up to a global transform. We set up REG-SDP for this problem and solve it using REG-ADMM. Figure 4.1 shows the dependence of REG-ADMM on ρ at different noise levels, where by "noise level" we mean the variance of the Gaussian noise. Notice that even when the local coordinates measurements are clean, REG-ADMM may get stuck in a local minimum depending on ρ. This dependence of the limit point on ρ is observed both at low and high noise levels. This is in contrast with C-ADMM, where the iterates converge to a global minimum for any positive ρ [49]. Also observe that when the noise is relatively large (Fig. 4.1c), the iterates of the algorithm may oscillate without converging if ρ is small. Such non-attenuating oscillations are not observed when the noise is low (see Figure 4.1b).

4.1.4 Contribution

The foregoing simulation results provide the main motivation for our investigations, namely, we wish to theoretically justify the observed behavior of REG-ADMM at different noise

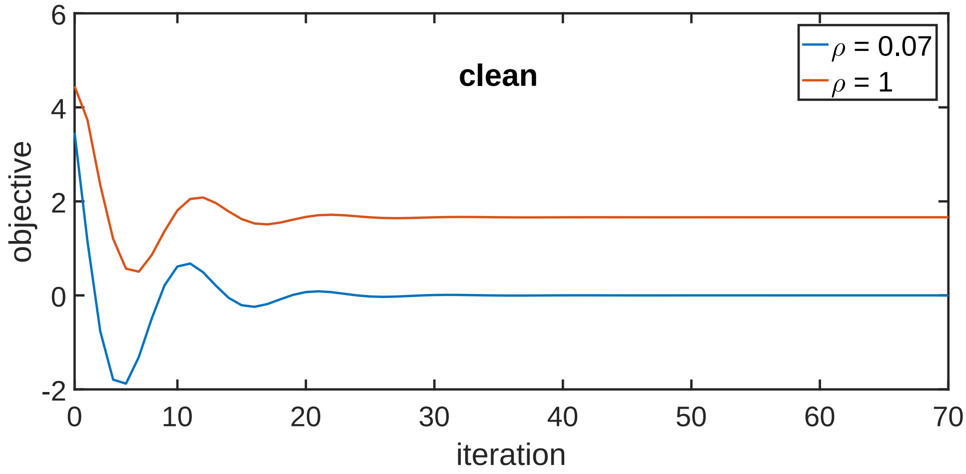

(a) Even with clean measurements, the iterates get stuck in a local minimum when $\rho = 1$ (the optimum value is zero in this case).

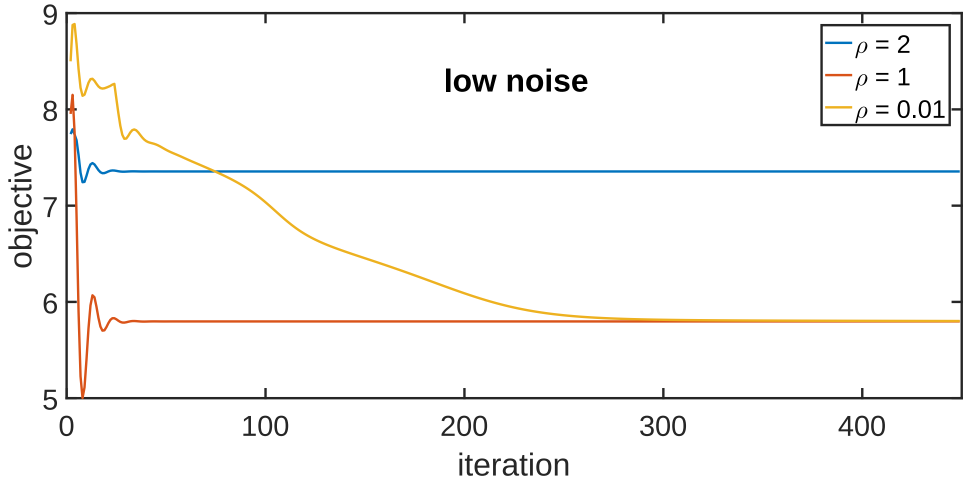

(b) There are no oscillations for small ρ when the noise is low.

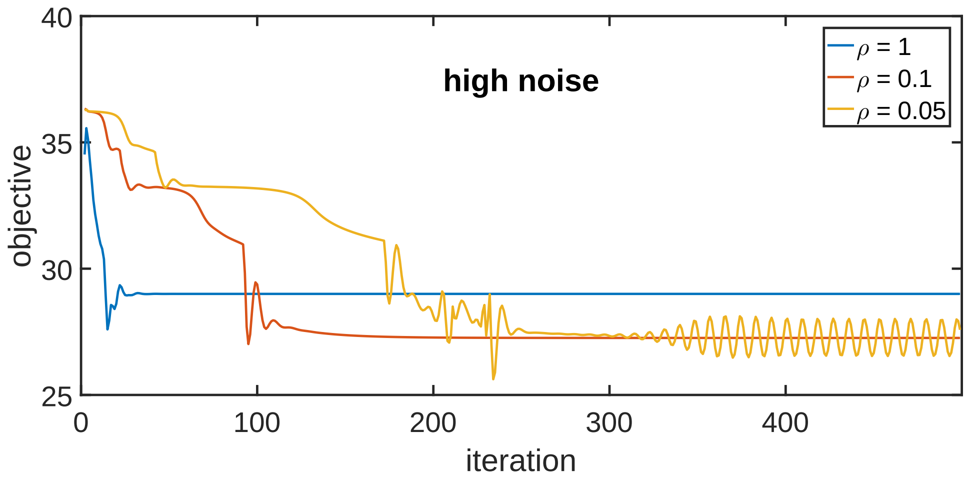

(c) The iterates oscillate for small ρ when the noise is high.

Figure 4.1: Simulation results at different noise levels (see the main text for a description of the experiment).

regimes. We first clarify what we mean by "noise". As mentioned previously, the information about pairwise relations among the M orthogonal transforms are encoded by the matrix $\mathbf{C}$, which ultimately depends on the local coordinates $\mathbf{x}_{k,i}$. If these measurements are exact, we say that $\mathbf{C}$ is clean; otherwise, we say that $\mathbf{C}$ is noisy. Furthermore, we make a distinction between *low* and *high* noise. In this context, we bring in the notion of tightness of the convex relaxation C-SDP. Recall that C-SDP is derived by dropping the rank constraint. Let $\mathbf{G}^*$ denote a global optimum of C-SDP. If $\text{rank}(\mathbf{G}^*) = d$, then clearly $\mathbf{G}^*$ is global optimum of REG-SDP as well, meaning that we have solved the original nonconvex problem. In this case, we say that the relaxation is *tight* [1]. In Chapter 3, we showed existence of a noise threshold, below which C-SDP remains tight. If the noise level in the data matrix $\mathbf{C}$ is below (resp. above) this threshold, we say that the noise is low (resp. high). An informal account of our main findings is as follows:

1. For any arbitrary data matrix $\mathbf{C}$, we prove that if the REG-ADMM iterates are asymptotically feasible and the dual iterates remain bounded, then any limit point of the iterates is a stationary (KKT) point of ORTHO-REG.

2. At low noise, we show that the REG-ADMM iterates converge to the global optimum, provided they are initialized sufficiently close to the optimum.

3. Recall that even when the data matrix is clean, REG-ADMM might get stuck in a local optimum (Figure 1.6a). If the data matrix is clean, then, for arbitrary initialization, we compute ρ for which REG-ADMM converges to the global optimum of REG-SDP (Corollary 47).

4. At high noise, we give a duality-based explanation of why the iterates exhibit non-attenuating oscillations when ρ is small, and why no such oscillations are observed at low noise. This suggests that for highly noisy measurements, ρ should not be set to a very small value.

The novelty of our analysis lies in the fact that we exploit the phenomenon of tightness of convex relaxation to prove convergence of the nonconvex ADMM. We contrast this approach with existing works on nonconvex ADMM (which rely on assumptions that do not apply to REG-ADMM) in the following subsection (Section 4.1.5).

4.1.5 Related Work

The rank-restricted subset Ω of the PSD cone is nonconvex, which implies that standard convergence result for ADMM [18] does not directly apply to REG-ADMM. However, we do leverage the convergence of convex ADMM for analyzing the convergence of REG-ADMM when the noise is low. The theoretical convergence of ADMM for nonconvex problems has been studied in [50, 51, 52]. However, a crucial working assumption common to these results does not hold in our case. More precisely, observe that we can rewrite REG-SDP as

$$\begin{aligned} \min_{\mathbf{G},\mathbf{H}\in\mathbb{S}^{Md}} \quad & \mathrm{Tr}(\mathbf{CG}) + \iota_\Omega(\mathbf{G}) + \iota_\Theta(\mathbf{H}) \\ \text{subject to} \quad & \mathbf{G} - \mathbf{H} = \mathbf{0}, \end{aligned} \tag{4.6}$$

where ι_Γ is the indicator function associated with a feasible set Γ [18], namely, $\iota_\Gamma(\mathbf{Y}) = 0$ if $\mathbf{Y} \in \Gamma$, and $\iota_\Gamma(\mathbf{Y}) = \infty$ otherwise. Notice that, because of the indicator functions, the objective function in (4.6) is non-differentiable in *both* $\mathbf{G}$ and $\mathbf{H}$. This violates a regularity assumption common in existing analyses of nonconvex ADMM, namely, that the objective must be differentiable in *at least* one variable. In these works, convergence results are obtained by proving a monotonic decrease in the augmented Lagrangian. This involves: (i) bounding successive difference in dual variables by successive difference in primal variables, which is where the assumption of smoothness is used; (ii) requiring that the parameter ρ is above a certain threshold. In particular, it is not clear whether this thresholding of the value of ρ is fundamental to convergence, or just an artifact of the analysis.

We do not make such smoothness assumptions in our analysis. We can afford to do this since we are analyzing a special class of problems, as opposed to the more general setups in [50, 51, 52]. Instead of showing a monotonic decrease in the augmented Lagrangian, our analysis relies on the phenomenon of tightness of convex relaxation. This provides more insights into the convergence behavior of the algorithm. For instance, our explanation in Section 4.2.4 shows that the instability of the algorithm (in the high-noise regime) for low values of ρ is fundamental, while suggesting why this instability is not observed in the low-noise regime.

4.1.6 Organization

This chapter is organized as follows. In Section 4.2, we state the main results. This section is divided into four subsections dealing with duality, general convergence result, convergence in low-noise regime (which includes, as a special case, convergence when the data matrix is clean), and oscillations in high-noise regime. We give proofs of the results of this section in Section 4.3, before concluding with a discussion in Section 4.4. The chapter ends with an Appendix containing a proof of an auxiliary result used in Section 4.3.

4.1.7 Notations

The notations used in this chapter are as mentioned in Chapter 3, Section 3.1.2.

4.2 Convergence Analysis

In this section we state and discuss the main results of the chapter. To improve readability, the technical proofs are deferred to section 4.3.

4.2.1 Duality

We start by discussing some results on Lagrange duality, which plays an important role throughout our analysis. To obtain stationarity conditions that optimum orthogonal transforms for the rigid registration problem must satisfy, we work with an equivalent formulation of REG-ADMM. Recall, from Section 3.2 in Chapter 3, that REG-ADMM is just a reformulation of the following optimization problem:

$$\text{(ORTHO-REG)} \qquad \min_{\mathbf{O}_1,\ldots,\mathbf{O}_M \in \mathbb{O}(d)} \sum_{i,j=1}^{M} \mathrm{Tr}\left(\mathbf{C}_{ij}\mathbf{O}_j^\top \mathbf{O}_i\right).$$

To derive KKT conditions for ORTHO-REG we write it as a standard nonlinear program [53]:

$$\begin{aligned} &\min_{\mathbf{O}_1,\cdots,\mathbf{O}_M \in \mathbb{R}^{d\times d}} && \sum_{i,j=1}^{M} \mathrm{Tr}\left([\mathbf{C}]_{ij}\mathbf{O}_j^\top \mathbf{O}_i\right) \\ &\text{subject to} && \mathbf{I}_d - \mathbf{O}_i^\top \mathbf{O}_i = \mathbf{0}, \qquad i \in [1:M]. \end{aligned} \tag{4.7}$$

It is not difficult to check that the gradients of the constraints in (4.7) are linearly independent. As a result, the KKT conditions necessarily hold at any local minimum of (4.7) [53]. The Lagrangian [53] for (4.7) is

$$\mathcal{L}(\mathbf{O}_1,\cdots,\mathbf{O}_M,\mathbf{\Lambda}_1,\cdots,\mathbf{\Lambda}_M) = \sum_{i,j=1}^{M} \mathrm{Tr}\left([\mathbf{C}]_{ij}\mathbf{O}_j^\top \mathbf{O}_i\right) + \sum_{i=1}^{M} \mathrm{Tr}(\mathbf{\Lambda}_k\left(\mathbf{I}_d - \mathbf{O}_i^\top \mathbf{O}_i\right)), \tag{4.8}$$

where the symmetric matrix $\mathbf{\Lambda}_i \in \mathbb{S}^d$ is the Lagrange multiplier for the i-th equality constraint in (4.7), $i \in [1:M]$. The following lemma characterizes any KKT point of ORTHO-REG.

Lemma 42. *The variables $\mathbf{O}_1^*, \ldots, \mathbf{O}_M^* \in \mathbb{R}^{d\times d}$ constitute a KKT point of* ORTHO-REG *if and only if, for $i \in [1:M]$,*

(a) $[\mathbf{G}^*]_{ii} = \mathbf{I}_d$, *and*

(b) $[\mathbf{C}\mathbf{G}^*]_{ii} = [\mathbf{G}^*\mathbf{C}]_{ii}$,

where $\mathbf{G}^*$ is the Gram matrix of $(\mathbf{O}_i^*)_{i=1}^M$, whose (i,j)-th block is $[\mathbf{G}^*]_{ij} = \mathbf{O}_i^{*\top}\mathbf{O}_j^*$. In this case, we say that $\mathbf{G}^*$ is a KKT point of ORTHO-REG.

We now consider the dual variables generated by C-ADMM and REG-ADMM. For both these algorithms, we will assume that the dual initialization $\mathbf{\Lambda}^0 \in \mathbb{S}^{Md}$ is block-diagonal.

Proposition 43. *Let the dual initialization* $\Lambda^0 \in \mathbb{S}^{Md}$ *for* REG-ADMM *(or* C-ADMM*) be block-diagonal. Then*

(i) Λ^k *is block-diagonal at every* $k \geq 1$*;*

(ii) The $\mathbf{H}$*-update step in* (4.5) *reduces to*

$$\mathbf{H}^{k+1} = \Pi_\Theta\left(\mathbf{G}^{k+1}\right). \tag{4.9}$$

4.2.2 General Convergence Result

We now state some convergence results wherein we do not assume anything about the data matrix $\mathbf{C}$. Lemma 44 essentially says that any fixed point of REG-ADMM is a KKT point of ORTHO-REG. Theorem 45 establishes convergence for REG-ADMM under some assumptions that empirically seem to hold if ρ is not too small.

Lemma 44. *Let* $\{\mathbf{G}^k, \mathbf{H}^k, \Lambda^k\}_{k=1}^\infty$ *be the iterates generated by* REG-ADMM. *Suppose there exists a subsequence* $\{\mathbf{H}^{k_l}\}_{l=1}^\infty \subset \{\mathbf{H}^k\}_{k=1}^\infty$ *such that* $\|\mathbf{G}^{k_l+1} - \mathbf{H}^{k_l}\| \to 0$*, and* $\mathbf{H}^{k_l} \to \mathbf{H}^*$ *as* $l \to \infty$*. Then* $\mathbf{H}^*$ *is a KKT point of* ORTHO-REG.

Lemma 44 implies that if REG-ADMM converges, it does so to a KKT point of ORTHO-REG. By convergence of REG-ADMM, we mean that the algorithm "stabilizes"; that is, *asymptotically*, the variables stop getting updated. Thus, if REG-ADMM converges, we have $\|\mathbf{G}^{k+1} - \mathbf{H}^k\| \to 0$ as $k \to \infty$, and this implies convergence to a KKT point by Lemma 44. We now state a result saying that if REG-ADMM generates primal iterates that are asymptotically feasible, and if the dual iterates remain bounded, then there is a subsequence that converges to a KKT point of ORTHO-REG.

Theorem 45. *Suppose the following conditions hold for the iterates generated by* REG-ADMM:

A1. $\|\mathbf{G}^k - \mathbf{H}^k\| \to \mathbf{0}$ *as* $k \to \infty$ *(asymptotic feasibility).*

A2. The dual iterates $\{\Lambda^k\}_{k=0}^\infty$ *are bounded.*

Then there exists a subsequence $\{\mathbf{H}^{k_l}\}_{l=1}^\infty$ *that converges to a KKT point of* ORTHO-REG.

For sufficiently large ρ, we empirically find that A1 and A2 do indeed hold. The requirement that ρ be sufficiently large is in line with other results on nonconvex ADMM [21, 51]. One way to justify our assumption about ρ is that, since ρ acts like a penalty parameter in (4.4), the algorithm has more "incentive" to push $\|\mathbf{G}^k - \mathbf{H}^k\|$ to 0 if ρ is large, thus leading to feasibility. Moreover, if this decay to 0 is fast enough, we will have $\sum_{k=1}^\infty \|\mathbf{G}^k - \mathbf{H}^k\| < \infty$, implying that Λ^k remains bounded.

4.2.3 Convergence in Low-Noise Regime

When noise in the data matrix is low enough for the convex relaxation C-SDP to be tight, we say that we are in the *low-noise regime.* Theorem 38 in Chapter 3 establishes the validity of this notion; that is, we do have a non-zero noise level, $\eta > 0$, below which C-SDP is tight. Recall that we use $\mathbf{C}_0$ to denote the clean data matrix, and $\mathbf{C} = \mathbf{C}_0 + \mathbf{W}$ to denote the noisy data matrix; that is, $\mathbf{W}$ is the perturbation to the clean data matrix due to noisy local coordinate measurements. We now state our result on convergence of REG-ADMM in the low-noise regime.

Theorem 46. *Let* $\mathbf{C}_0, \mathbf{W}$ *be as in Theorem 38. Suppose* $\|\mathbf{W}\| < \eta$*, which implies that the convex relaxation is tight, i.e.,* $\mathrm{rank}(\mathbf{G}^*) = d$*, where* $\mathbf{G}^*$ *is a global optimum for* C-SDP*. Let* $\mathbf{\Lambda}^*$ *be the dual variable corresponding to* $\mathbf{G}^*$ *(see Lemma 35). Suppose* $\mathbf{H}^0, \mathbf{\Lambda}^0, \rho$ *are such that*

$$\|\mathbf{H}^0 - \mathbf{G}^*\|^2 + \frac{1}{\rho^2}\|\mathbf{\Lambda}^0 - \mathbf{\Lambda}^*\|^2 \le \frac{1}{\rho^2}\,\lambda_{d+1}^2(\mathbf{C} + \mathbf{\Lambda}^*). \tag{4.10}$$

Then REG-ADMM *converges to the global optimum, that is,* $\mathbf{H}^k \longrightarrow \mathbf{G}^*$ *and* $\mathbf{\Lambda}^k \longrightarrow \mathbf{\Lambda}^*$ *as* $k \longrightarrow \infty$.

While proving Theorem 46 in Section 4.3.3, we will see that when the conditions in Theorem 46 are satisfied, REG-ADMM and C-ADMM generate identical iterates. This allows us to infer the convergence of REG-ADMM from that of C-ADMM. Thus, we reap the convergence benefits of C-ADMM, while retaining the computational advantages of REG-ADMM (see Table 4.1).

Observe that Theorem 46 implies a tradeoff, namely, if we intialize the dual sufficiently close to the optimal, we can be lax with the primal initialization. This principle is brought to the fore in the clean case, where we know that the dual optimum $\mathbf{\Lambda}^* = \mathbf{0}$ (this is because, in the clean case, $\mathbf{G}_0$ is the primal global optimum, and $\mathbf{C}_0\mathbf{G}_0 = \mathbf{0}$; see (3.10) and Proposition 36).

Corollary 47. *Let* $\mathrm{nullity}(\mathbf{C}_0) = d$ *and* $\mathbf{\Lambda}^0 = \mathbf{0}$*. Then given any primal initialization* $\mathbf{H}^0$*,* REG-ADMM *converges to global optimum provided*

$$\rho \le \frac{\lambda_{d+1}(\mathbf{C}_0)}{\sqrt{\|\mathbf{H}^0\|^2 + 2M\sqrt{d}\,\|\mathbf{H}^0\| + M^2 d}}.$$

In Fig. 4.1a, REG-ADMM gets stuck in a local optimum for $\rho = 1$, but converges to the global optimum if ρ is set using Corollary 47. The result also sheds light on the robustness of the algorithm to primal initializations, as observed in [9]. Note that we have an upper bound on ρ, in contrast to existing results in the literature which prescribe a lower bound. This phenomenon is peculiar to low-noise data. In contrast, small values of ρ in the high-noise regime may lead to non-attenuating oscillations (Fig. 4.1c).

4.2.4 Oscillations in High-Noise Regime

Before discussing oscillations in the high-noise regime, we will derive a necessary condition that a fixed-point of REG-ADMM must satisfy. We first clarify what we mean by a "fixed-point".

Definition 48. *Suppose $\mathcal{T}$ is an iterative procedure initialized with $\boldsymbol{z}^0$, and working as follows: $\boldsymbol{z}^{k+1} = \mathcal{T}(\boldsymbol{z}^k)$. Then $\boldsymbol{z}^{k_0}$ is a* fixed point *of $\mathcal{T}$ if $\boldsymbol{z}^{k_0} = \mathcal{T}(\boldsymbol{z}^{k_0})$.*

Suppose $(\mathbf{H}^*, \boldsymbol{\Lambda}^*)$ is a fixed-point of REG-ADMM with parameter ρ. Then, if we were to initialize REG-ADMM with $\mathbf{H}^0 = \mathbf{H}^*$, $\boldsymbol{\Lambda}^0 = \boldsymbol{\Lambda}^*$, we would have $\mathbf{H}^1 = \mathbf{H}^0\ (= \mathbf{H}^*)$, $\boldsymbol{\Lambda}^1 = \boldsymbol{\Lambda}^0\ (= \boldsymbol{\Lambda}^*)$. In other words, the algorithm would make no progress. The following proposition specifies a condition involving $\mathbf{H}^*$ and $\boldsymbol{\Lambda}^*$ that must necessarily hold in this case.

Proposition 49. *Suppose $(\mathbf{H}^*, \boldsymbol{\Lambda}^*)$ is a fixed-point of REG-ADMM algorithm with parameter ρ. Then $(\mathbf{C} + \boldsymbol{\Lambda}^*)\mathbf{H}^* = \mathbf{0}$.*

Let us now focus on the high-noise regime. Suppose REG-ADMM is initialized with $(\mathbf{H}^0, \boldsymbol{\Lambda}^0)$ such that $\mathbf{H}^0$ is feasible for REG-SDP, $\boldsymbol{\Lambda}^0$ is block-diagonal, and $(\mathbf{C} + \boldsymbol{\Lambda}^0)\mathbf{H}^0 = \mathbf{0}$. In the high-noise regime (i.e., when the convex relaxation is not tight), $\mathbf{C} + \boldsymbol{\Lambda}^0$ must have at least one negative eigenvalue (say $-\mu^2$). Indeed, if $\mathbf{C} + \boldsymbol{\Lambda}^0$ were positive semidefinite, it would follow from Lemma 35 that the convex relaxation is tight. Now, the $\mathbf{G}$-update is

$$\mathbf{G}^1 = \Pi_\Omega\left(\mathbf{H}^0 - \rho^{-1}(\mathbf{C} + \boldsymbol{\Lambda}^0)\right).$$

From Proposition 40, we know that the d non-zero eigenvalues of $\mathbf{H}^0$ are M. Now, observe that $-\rho^{-1}(\mathbf{C} + \boldsymbol{\Lambda}^0)$ has a positive eigenvalue $\rho^{-1}\mu^2$. If ρ is sufficiently small, we would have $\rho^{-1}\mu^2 > M$. Then, since $\mathbf{G}^1$ is determined by the top d eigenvalues of $\mathbf{H}^0 - \rho^{-1}(\mathbf{C} + \boldsymbol{\Lambda}^0)$, we would have that the top eigenvalue of $\mathbf{G}^1$ is strictly bigger than M, making $\mathbf{G}^1$ infeasible for REG-SDP (since, by Proposition 40, any $\mathbf{G}$ feasible for REG-SDP necessarily has all non-zero eigenvalues equal to M). This would imply that $\mathbf{H}^1 \neq \mathbf{G}^1$, and consequently, that $\boldsymbol{\Lambda}^1 \neq \boldsymbol{\Lambda}^0$. That is, for small value of ρ, we see that REG-ADMM does not stabilize even when $\mathbf{H}^0$ and $\boldsymbol{\Lambda}^0$ satisfy $(\mathbf{C} + \boldsymbol{\Lambda}^0)\mathbf{H}^0 = \mathbf{0}$, a property that any candidate for a fixed-point of REG-ADMM must satisfy. Put differently, there can be no fixed-point to which REG-ADMM converges if ρ is sufficiently small.

Observe that the argument above depended on the existence of a negative eigenvalue of $\mathbf{C} + \boldsymbol{\Lambda}^0$. This argument does not hold in the low-noise regime because we can simultaneously have the properties that $(\mathbf{C} + \boldsymbol{\Lambda}^0)\mathbf{H}^0 = \mathbf{0}$, and that all the eigenvalues of $\mathbf{C} + \boldsymbol{\Lambda}^0$ are nonnegative (which holds when $\mathbf{H}^0$ is global optimum; see condition (b) in Lemma 35). This suggests why the instability is not observed for low values of ρ in the low-noise regime.

4.3 Technical Proofs

In this section, we give proofs of the results stated in Section 4.2. The subsections here are given the same title as subsections in Section 4.2 to facilitate the correspondence between the results and the proofs.

4.3.1 Duality

Proof of Lemma 42. For a minimization problem with equality constraints, KKT conditions amount to primal feasibility, and stationarity of Lagrangian (4.8) with respect to the primal variables [53] (i.e., the partial derivative of (4.8) with respect to $\mathbf{O}_i$ should vanish). Primal feasibility gives us condition (a). Vanishing of partial derivative of (4.8) with respect to $\mathbf{O}_i$ gives us

$$\mathbf{O}_i^*\mathbf{\Lambda}_i^* = -\sum_{j=1}^{M}\mathbf{O}_j^*[\mathbf{C}]_{ji}, \quad i \in [1:M]. \tag{4.11}$$

Left multiplying by $\mathbf{O}_i^{*\top}$ and using the primal feasibility condition that $\mathbf{O}_i^{*\top}\mathbf{O}_i^* = \mathbf{I}_d$, we obtain

$$\mathbf{\Lambda}_i^* = -\sum_{j=1}^{M}\mathbf{O}_i^{*\top}\mathbf{O}_j^*[\mathbf{C}]_{ji} = -\sum_{j=1}^{M}[\mathbf{G}^*]_{ij}[\mathbf{C}]_{ji} = -[\mathbf{G}^*\mathbf{C}]_{ii}.$$

Also, note that $\mathbf{\Lambda}_i^{*\top} = -[\mathbf{C}\mathbf{G}^*]_{ii}$. Since $\mathbf{\Lambda}_i^*$ is symmetric, condition (b) follows immediately. Conversely, given conditions (a) and (b) on $\mathbf{G}^*$, it is not difficult to see that $\mathbf{\Lambda}_i$'s defined as $\mathbf{\Lambda}_i = -[\mathbf{C}\mathbf{G}^*]_{ii}$, and $\mathbf{O}_i$'s defined such that $\mathbf{O}_j^\top\mathbf{O}_k = [\mathbf{G}^*]_{jk}$ $\forall j,k \in [1:M]$, together satisfy the KKT conditions. □

Proof of Proposition 43. We prove the proposition for REG-ADMM by induction. The proof for C-ADMM is exactly the same since the $\mathbf{H}$-update and $\mathbf{\Lambda}$-update steps are identical for REG-ADMM and C-ADMM. Clearly, the proposition holds for $k = 0$. Assume that the proposition holds for $k = k_0$. We will show that it then has to hold for $k = k_0 + 1$. Consider the $\mathbf{H}$-update step in (4.5)

$$\mathbf{H}^{k_0+1} = \Pi_\Theta\left(\mathbf{G}^{k_0+1} + \rho^{-1}\mathbf{\Lambda}^{k_0}\right).$$

We know that Θ is the set of symmeteric matrices for which $d \times d$ diagonal blocks are $\mathbf{I}_d$. It is clear that projection of a matrix on Θ is obtained by setting the diagonal blocks of the matrix to $\mathbf{I}_d$. In other words, $\Pi_\Theta(\cdot)$ affects only the diagonal blocks of its argument. By induction hypothesis, $\mathbf{\Lambda}^{k_0}$ is block-diagonal, and thus adding it to $\mathbf{G}_0^k$ in the $\mathbf{H}$-update step does not affect the projection on Θ, or,

$$\mathbf{H}^{k_0+1} = \Pi_\Theta\left(\mathbf{G}^{k_0+1}\right).$$

Now, consider the $\mathbf{\Lambda}$-update step,

$$\mathbf{\Lambda}^{k_0+1} = \mathbf{\Lambda}^{k_0} + \rho\left(\mathbf{G}^{k_0+1} - \mathbf{H}^{k_0+1}\right).$$

Since $\Pi_\Theta(\cdot)$ affects only the diagonal blocks of its arguments, it is clear that the off-diagonal blocks of $\mathbf{H}^{k_0+1}$ and $\mathbf{G}^{k_0+1}$ are the same. Thus, $\mathbf{G}^{k_0+1} - \mathbf{H}^{k_0+1}$ is a block-diagonal matrix, which along with the hypothesis that $\mathbf{\Lambda}^{k_0}$ is block-diagonal, implies that $\mathbf{\Lambda}^{k_0+1}$ is block-diagonal. □

4.3.2 General Convergence Result

Proof of Lemma 44. Since $\mathbf{G}^{k_l+1}$ is formed from top-d eigendecomposition of $\mathbf{H}^{k_l} - \rho^{-1}\left(\mathbf{C} + \mathbf{\Lambda}^{k_l}\right)$ (see (4.5)), the range space of $\left(\mathbf{H}^{k_l} - \rho^{-1}\left(\mathbf{C} + \mathbf{\Lambda}^{k_l}\right) - \mathbf{G}^{k_l+1}\right)$ is orthogonal to range space of $\mathbf{G}^{k_l+1}$, which gives us

$$\left(\mathbf{H}^{k_l} - \rho^{-1}\left(\mathbf{C} + \mathbf{\Lambda}^{k_l}\right) - \mathbf{G}^{k_l+1}\right)\mathbf{G}^{k_l+1} = \mathbf{0}. \tag{4.12}$$

Or,

$$\left(\mathbf{H}^{k_l} - \mathbf{G}^{k_l+1}\right)\mathbf{G}^{k_l+1} = \rho^{-1}\left(\mathbf{C} + \mathbf{\Lambda}^{k_l}\right)\mathbf{G}^{k_l+1}.$$

By hypothesis, the left hand side of the equation above goes to $\mathbf{0}$ as $l \longrightarrow \infty$. Thus,

$$\left(\mathbf{C} + \mathbf{\Lambda}^{k_l}\right)\mathbf{G}^{k_l+1} \longrightarrow \mathbf{0} \text{ as } l \longrightarrow \infty. \tag{4.13}$$

Now, $\|\mathbf{G}^{k_l+1} - \mathbf{H}^{k_l}\| \longrightarrow 0$, which, together with (4.13) (and from continuity of linear operators) implies that, as $l \longrightarrow \infty$,

$$\begin{aligned} &\left(\mathbf{C} + \mathbf{\Lambda}^{k_l}\right)\mathbf{H}^{k_l} \longrightarrow \mathbf{0} \\ \Rightarrow \quad & \mathbf{C}\mathbf{H}^{k_l} \longrightarrow -\mathbf{\Lambda}^{k_l}\mathbf{H}^{k_l} \\ \Rightarrow \quad & [\mathbf{C}\mathbf{H}^{k_l}]_{ii} \longrightarrow -[\mathbf{\Lambda}^{k_l}\mathbf{H}^{k_l}]_{ii}. \end{aligned}$$

Note that $[\mathbf{\Lambda}^{k_l}\mathbf{H}^{k_l}]_{ii}$ is symmetric (since $\mathbf{\Lambda}^{k_l}$ is block-diagonal, and $[\mathbf{H}^{k_l}]_{ii} = \mathbf{I}_d$). Thus, we have that $[\mathbf{C}\mathbf{H}^{k_l}]_{ii}$ is closer and closer to being symmetric. That is, as $l \longrightarrow \infty$,

$$\operatorname{dist}\left([\mathbf{C}\mathbf{H}^{k_l}]_{ii}, \mathbb{S}^d\right) \longrightarrow 0 \quad \forall\, i \in [1:M],$$

where 'dist' stands for distance. Now, by continuity of linear operator $\mathbf{C}$,

$$[\mathbf{C}\mathbf{H}^{k_l}]_{ii} \longrightarrow [\mathbf{C}\mathbf{H}^*]_{ii}.$$

Thus,

$$\operatorname{dist}\left([\mathbf{C}\mathbf{H}^*]_{ii}, \mathbb{S}^d\right) = 0.$$

Since $\mathbb{S}^d \subset \mathbb{R}^{d\times d}$ is a closed set, this means that $[\mathbf{CH}^*]_{ii}$ is symmetric. By Lemma 42, we have our result. □

Proof of Theorem 45. Because $\mathbf{G}^{k+1}$ minimizes $\mathcal{L}_\rho(\mathbf{G}, \mathbf{H}^k, \mathbf{\Lambda}^k)$ over Ω,

$$\mathcal{L}_\rho(\mathbf{G}^k, \mathbf{H}^k, \mathbf{\Lambda}^k) - \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^k, \mathbf{\Lambda}^k) \geq 0.$$

Similarly, because $\mathbf{H}^{k+1}$ minimizes $\mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}, \mathbf{\Lambda}^k)$ over Θ,

$$\mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^k, \mathbf{\Lambda}^k) - \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^k) \geq 0. \tag{4.14}$$

In fact, we can say more for the $\mathbf{H}$-update. We have

$$\begin{aligned}
&\mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^k, \mathbf{\Lambda}^k) - \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^k) \\
&= \mathrm{Tr}(\mathbf{\Lambda}^k(\mathbf{H}^k - \mathbf{H}^{k+1})) + \frac{\rho}{2}\|\mathbf{G}^{k+1} - \mathbf{H}^k\|^2 - \frac{\rho}{2}\|\mathbf{G}^{k+1} - \mathbf{H}^{k+1}\|^2
\end{aligned} \tag{4.15}$$

Now, $\mathrm{Tr}(\mathbf{\Lambda}^k(\mathbf{H}^k - \mathbf{H}^{k+1})) = 0$ because $\mathbf{\Lambda}^k$ is block-diagonal and the diagonal blocks of $(\mathbf{H}^k - \mathbf{H}^{k+1})$ are $\mathbf{0}$. Thus,

$$\begin{aligned}
&\mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^k, \mathbf{\Lambda}^k) - \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^k) \\
&= \frac{\rho}{2}\|\mathbf{G}^{k+1} - \mathbf{H}^k\|^2 - \frac{\rho}{2}\|\mathbf{G}^{k+1} - \mathbf{H}^{k+1}\|^2 \\
&= \frac{\rho}{2}\|\mathbf{H}^{k+1} - \mathbf{H}^k\|^2
\end{aligned} \tag{4.16}$$

where the last equality is by Pythagoras theorem, since Θ is an affine space. So, we have seen that the Lagrangian *decreases* during the $\mathbf{G}$-update and $\mathbf{H}$-update steps, and that we can quantify the decrease during the $\mathbf{H}$-update step. Consider now the $\mathbf{\Lambda}$-update step,

$$\begin{aligned}
&\mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^k) - \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^{k+1}) \\
&= -\rho\|\mathbf{G}^{k+1} - \mathbf{H}^{k+1}\|^2.
\end{aligned} \tag{4.17}$$

This means that the Lagrangian *increases* during the $\mathbf{\Lambda}$-update step. So,

$$\begin{aligned}
&\mathcal{L}_\rho(\mathbf{G}^k, \mathbf{H}^k, \mathbf{\Lambda}^k) - \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^{k+1}) \\
&= \mathcal{L}_\rho(\mathbf{G}^k, \mathbf{H}^k, \mathbf{\Lambda}^k) - \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^k, \mathbf{\Lambda}^k) + \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^k, \mathbf{\Lambda}^k) - \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^k) \\
&\quad + \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^k) - \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^{k+1}) \\
&\geq \frac{\rho}{2}\left(\|\mathbf{H}^{k+1} - \mathbf{H}^k\|^2 - 2\|\mathbf{G}^{k+1} - \mathbf{H}^{k+1}\|^2\right),
\end{aligned} \tag{4.18}$$

where the inequality results from equations (4.14), (4.16), and (4.17).
Let $\{\mathbf{G}^k, \mathbf{H}^k, \mathbf{\Lambda}^k\}_{k=1}^{\infty}$ be the iterates generated by the algorithm. Then for a fixed $\epsilon_0 > 0$, they will satisfy one of the following mutually exclusive cases:

Case 1: $\|\mathbf{H}^{k+1} - \mathbf{H}^k\| \leq (2 + \epsilon_0)\, \|\mathbf{G}^{k+1} - \mathbf{H}^{k+1}\|$ infinitely often.

Case 2: $\|\mathbf{H}^{k+1} - \mathbf{H}^k\| \leq (2 + \epsilon_0)\, \|\mathbf{G}^{k+1} - \mathbf{H}^{k+1}\|$ finitely often.

Suppose Case 1 holds. Choose a subsequence $\left(\mathbf{H}^{k_j}, \mathbf{H}^{k_j+1}\right)_{j=1}^{\infty}$ such that

$$\|\mathbf{H}^{k_j+1} - \mathbf{H}^{k_j}\| \leq (2 + \epsilon_0)\, \|\mathbf{G}^{k_j+1} - \mathbf{H}^{k_j+1}\|. \tag{4.19}$$

Now, from Assumption A1,

$$\|\mathbf{G}^{k_j+1} - \mathbf{H}^{k_j+1}\| \to 0 \text{ as } j \to \infty. \tag{4.20}$$

This, from (4.19), implies that

$$\|\mathbf{H}^{k_j+1} - \mathbf{H}^{k_j}\| \to 0 \text{ as } j \to \infty. \tag{4.21}$$

Equations (4.20) and (4.21), together with Pythagoras theorem (since Θ is an affine space) imply that

$$\|\mathbf{G}^{k_j+1} - \mathbf{H}^{k_j}\| \to 0 \text{ as } j \to \infty.$$

Now, $\{\mathbf{H}^{k_j}\}_{j=1}^{\infty}$ is bounded, because the algorithm tends to feasibility, and the feasible set $\Omega \cap \Theta$ is bounded. Thus, we can choose a convergent subsequence $\{\mathbf{H}^{k_l}\}_{l=1}^{\infty} \subset \{\mathbf{H}^{k_j}\}_{j=1}^{\infty}$. Let $\mathbf{H}^{k_l} \to \mathbf{H}^*$ as $l \to \infty$. By Lemma 44, $\mathbf{H}^*$ is a KKT point.

Suppose Case 2 holds. Then there exists a K_0 such that

$$\|\mathbf{H}^{k+1} - \mathbf{H}^k\| > (2 + \epsilon_0)\|\mathbf{G}^{k+1} - \mathbf{H}^{k+1}\| \;\forall\, k \geq K_0.$$

Thus, from (4.18), we get that

$$\mathcal{L}_\rho(\mathbf{G}^k, \mathbf{H}^k, \mathbf{\Lambda}^k) - \mathcal{L}_\rho(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^{k+1}) > \frac{\rho}{2}\epsilon_0\, \|\mathbf{G}^{k+1} - \mathbf{H}^k\|^2 \quad \forall\, k \geq K_0. \tag{4.22}$$

This means that the Lagrangian decreases monotonically for $k \geq K_0$. Now, $\left\langle \mathbf{C}, \mathbf{G}^k \right\rangle \geq 0 \;\forall\, k$, because $\mathbf{C}, \mathbf{G}^k \in \mathbb{S}_+^{Md}$. Also, $\frac{\rho}{2}\|\mathbf{G}^k - \mathbf{H}^k\|^2 \geq 0$. Thus, using expression (4.4) for the Lagrangian, we get

$$\begin{aligned}\mathcal{L}_\rho(\mathbf{G}^k, \mathbf{H}^k, \mathbf{\Lambda}^k) &\geq \left\langle \mathbf{\Lambda}^k, \mathbf{G}^k - \mathbf{H}^k \right\rangle \\ &\geq -\|\mathbf{\Lambda}^k\|\, \|\mathbf{G}^k - \mathbf{H}^k\|\end{aligned}$$

where the second inequality is Cauchy-Schwarz. It now follows from Assumptions A1 and A2 that the Lagrangian is lower bounded. So, we have established that the Lagrangian decreases (strictly) monotonically, and is lower bounded. Thus, from (4.22), $\|\mathbf{G}^{k+1} - \mathbf{H}^k\|^2 \to 0$ as $k \to \infty$. Existence of convergent subsequence converging to a KKT point now follows as in Case 1. □

4.3.3 Convergence in Low-Noise Regime

We start with a result on C-ADMM, which we will leverage to get results for REG-ADMM in the low-noise regime. Suppose $\mathbf{A}^k = \mathbf{H}^k - \rho^{-1}(\mathbf{C} + \mathbf{\Lambda}^k)$, $\mathbf{A}^* = \mathbf{H}^* - \rho^{-1}(\mathbf{C} + \mathbf{\Lambda}^*)$, where $\mathbf{H}^*$ is a global minimum for the convex relaxation C-SDP, and $\mathbf{\Lambda}^*$ is the corresponding optimal dual as in Lemma 35.

Lemma 50. *For* C-ADMM, $\|\mathbf{A}^{k+1} - \mathbf{A}^*\|^2 \leq \|\mathbf{A}^k - \mathbf{A}^*\|^2$, *for every* k.

We are now ready to prove results on convergence of REG-ADMM in the low-noise regime.

Proof of Theorem 46. Since $\|\mathbf{W}\| < \eta$, we have that nullity$(\mathbf{C} + \mathbf{\Lambda}^*) = d$ (see proof of Theorem 38), which, in particular implies that $\lambda_{d+1}(\mathbf{C} + \mathbf{\Lambda}^*) > 0$. Recall that, according to our notation, $\lambda_{d+1}(\mathbf{C} + \mathbf{\Lambda}^*)$ is the $(d+1)$-th eigenvalue of $\mathbf{C} + \mathbf{\Lambda}^*$, where we assume the eigenvalues to be arranged in an ascending order. We now deduce the eigenvalues of $\mathbf{A}^*$. From KKT condition (a) in Lemma 35, we have that

$$(\mathbf{C} + \mathbf{\Lambda}^*)\, \mathbf{G}^* = \mathbf{0}.$$

In particular, this means that $\mathbf{C} + \mathbf{\Lambda}^*$ and $\mathbf{G}^*$ are simultaneously diagonalizable, and since their product is $\mathbf{0}$, their range spaces are orthogonal. Now, since rank$(\mathbf{G}^*) = d$, the non-zero eigenvalues of $\mathbf{G}^*$ are all M (Proposition 40). Moreover, from KKT condition (b) in Lemma 35, we have that $\mathbf{C} + \mathbf{\Lambda}^* \in \mathbb{S}_+^{Md}$. Putting everything together, we have that

- $\mathbf{A}^*$ has d positive eigenvalues, each of them equal to M.
- $\mathbf{A}^*$ has $(M-1)d$ negative eigenvalues, which are the negative of the non-zero eigenvalues of $\rho^{-1}(\mathbf{C} + \mathbf{\Lambda}^*)$.

Let

$$\mathcal{B}^* := \{\mathbf{A} \in \mathbb{S}^{Md} : \|\mathbf{A} - \mathbf{A}^*\| < \rho^{-1}\lambda_{d+1}(\mathbf{C} + \mathbf{\Lambda}^*)\}.$$

By Weyl's theorem (Theorem 37), we deduce that for any $\mathbf{A} \in \mathcal{B}^*$, the top d eigenvalues of $\mathbf{A}$ would lie in the interval $\big(M - \rho^{-1}\lambda_{d+1}(\mathbf{C} + \mathbf{\Lambda}^*), M + \rho^{-1}\lambda_{d+1}(\mathbf{C} + \mathbf{\Lambda}^*)\big)$, and other eigenvalues of $\mathbf{A}$ would lie in the interval $\big(-\rho^{-1}\left(\lambda_{Md}(\mathbf{C} + \mathbf{\Lambda}^*) + \lambda_{d+1}(\mathbf{C} + \mathbf{\Lambda}^*)\right), 0\big)$. In other words, for any $\mathbf{A} \in \mathcal{B}^*$, only the top d eigenvalues of $\mathbf{A}$ would be nonnegative. Let

$$\mathbf{A}^k = \mathbf{H}^k - \rho^{-1}\left(\mathbf{C} + \mathbf{\Lambda}^k\right).$$

Suppose $\mathbf{A}^k \in \mathcal{B}^*$. Then, because only the top d eigenvalues of $\mathbf{A}^k$ are nonnegative, we have

$$\mathbf{G}^{k+1} = \Pi_\Omega\left(\mathbf{A}^k\right) = \Pi_{\mathbb{S}_+^{Md}}\left(\mathbf{A}^k\right).$$

That is, projection on the nonconvex set Ω is same as projection on the convex set $\mathbb{S}_+^{Md}$.

Now, from Lemma 50 we infer that, if $\mathbf{A}^k \in \mathcal{B}^*$, then $\mathbf{A}^{k+1} \in \mathcal{B}^*$. Thus, every subsequent projection on Ω is equivalent to projection on $\mathbb{S}_+^{Md}$. Note that condition (4.10) in the hypothesis of the theorem implies that $\mathbf{A}^0 \in \mathcal{B}^*$. Thus, the iterates generated by REG-ADMM initialized with $\mathbf{H}_0, \mathbf{\Lambda}_0$ is the same as the iterates generated by C-ADMM initialized with $\mathbf{H}_0, \mathbf{\Lambda}_0$. Now, since C-ADMM converges to global optimum [49], we deduce that REG-ADMM converges to global optimum. □

Proof of Corollary 47. For the clean case, $\mathbf{\Lambda}^* = \mathbf{0}$. Thus, with $\mathbf{\Lambda}^0 = \mathbf{0}$, the condition in Theorem 46 reduces to

$$\|\mathbf{H}^0 - \mathbf{G}^*\|^2 \leq \frac{1}{\rho^2}\lambda_{d+1}^2(\mathbf{C}_0).$$

Now,

$$\|\mathbf{H}^0 - \mathbf{G}^*\|^2 \leq \|\mathbf{H}^0\|^2 + \|\mathbf{G}^*\|^2 + 2\|\mathbf{H}^0\|\,\|\mathbf{G}^*\|.$$

Since $\|\mathbf{G}^*\| = M\sqrt{d}$, we have

$$\|\mathbf{H}^0 - \mathbf{G}^*\|^2 \leq \|\mathbf{H}^0\|^2 + 2M\sqrt{d}\,\|\mathbf{H}^0\| + M^2 d.$$

Thus, the algorithm would converge to the global optimum if

$$\|\mathbf{H}^0\|^2 + 2M\sqrt{d}\,\|\mathbf{H}^0\| + M^2 d \leq \frac{1}{\rho^2}\lambda_{d+1}^2(\mathbf{C}_0),$$

which proves the result. □

4.3.4 Oscillations in High-Noise Regime

Proof of Proposition 49. Suppose we initialize REG-ADMM with $\mathbf{H}^0 = \mathbf{H}^*$, $\mathbf{\Lambda}^0 = \mathbf{\Lambda}^*$. Then, by the conditions that $\mathbf{\Lambda}^1 = \mathbf{\Lambda}^0$, and that $\mathbf{H}^1 = \mathbf{H}^0$, we get that $\mathbf{G}^1 = \mathbf{H}^0$. Now, since $\mathbf{G}^1$ is formed from top-d eigendecomposition of $\mathbf{H}^0 - \rho^{-1}\left(\mathbf{C} + \mathbf{\Lambda}^0\right)$ (see (4.5)), the range space of $\left(\mathbf{H}^0 - \rho^{-1}\left(\mathbf{C} + \mathbf{\Lambda}^0\right) - \mathbf{G}^1\right)$ is orthogonal to range space of $\mathbf{G}^1$, which gives us

$$\left(\mathbf{H}^0 - \rho^{-1}\left(\mathbf{C} + \mathbf{\Lambda}^0\right) - \mathbf{G}^1\right)\mathbf{G}^1 = \mathbf{0}. \tag{4.23}$$

Or,

$$\rho^{-1}\left(\mathbf{C} + \mathbf{\Lambda}^0\right)\mathbf{G}^1 = \left(\mathbf{H}^0 - \mathbf{G}^1\right)\mathbf{G}^1.$$

Since $\mathbf{G}^1 = \mathbf{H}^0$ $(= \mathbf{H}^*)$, the right hand side of the preceding equation is $\mathbf{0}$, and we obtain the desired result. □

4.4 Discussion

In this chapter, we analyzed the convergence behavior of REG-ADMM in different noise regimes. Existing results on nonconvex ADMM do not apply to REG-ADMM as they rely on certain smoothness assumptions that are not satisfied by REG-SDP. We bypassed these assumptions, and exploited the tightness phenomenon of convex relaxation to guide our analysis. We started with a general convergence result saying that if the primal iterates generated by REG-ADMM are asymptotically feasibile and the dual iterates are bounded, then a subsequence converges to a stationary (KKT) point of REG-SDP. To further refine the result, we looked at the behavior of REG-ADMM when the noise is low. In particular, we defined precisely what is meant by "low" noise by invoking tightness of the convex relaxation C-SDP below a certain noise threshold. We then proved that, by initializing the primal and dual variables sufficiently close to the optimum, the iterates of REG-ADMM are guaranteed to converge to the global optimum. By applying this result to the clean case, we showed that given *any* primal initialization, we can explicitly compute values of ρ for which the algorithm converges to the global optimum. For high noise, we showed that for sufficiently small ρ, the iterates generated by REG-ADMM do not stabilize, even if initialization of REG-ADMM satisfies necessary property of a fixed-point. Thus, for highly noisy measurements, ρ should be set to a relatively large value to ensure convergence of REG-ADMM.

4.5 Appendix

4.5.1 Proof of Lemma 50

The proof of this lemma essentially follows the convergence proof for convex ADMM presented in [18]. Recall that the C-ADMM is trying to solve the following optimization problem (which is a reformulation of C-SDP):

$$\begin{aligned} \min_{\mathbf{G},\mathbf{H}\in\mathbb{S}^{Md}} \quad & \mathrm{Tr}(\mathbf{C}\mathbf{G}) \\ \text{subject to} \quad & \mathbf{G}\in\mathbb{S}^{Md}_{+},\ \mathbf{H}\in\Theta, \\ & \mathbf{G}-\mathbf{H}=\mathbf{0}, \end{aligned} \tag{4.24}$$

where, $\Theta = \{\mathbf{X}\in\mathbb{S}^{Md} : [\mathbf{X}]_{ii} = \mathbf{I}_d\}$.

Now, note that

$$\begin{aligned} &\|\mathbf{A}^k-\mathbf{A}^*\|^2 - \|\mathbf{A}^{k+1}-\mathbf{A}^*\|^2 \\ &= \|\mathbf{H}^k-\mathbf{H}^*\|^2 - \|\mathbf{H}^{k+1}-\mathbf{H}^*\|^2 + \rho^{-2}\left(\|\mathbf{\Lambda}^k-\mathbf{\Lambda}^*\|^2 - \|\mathbf{\Lambda}^{k+1}-\mathbf{\Lambda}^*\|^2\right) \end{aligned} \tag{4.25}$$

where we have used the fact that, for every k,

$$\left\langle \mathbf{H}^k-\mathbf{H}^*, \mathbf{\Lambda}^k-\mathbf{\Lambda}^*\right\rangle = 0,$$

because $\mathbf{H}^k-\mathbf{H}^*$ has $\mathbf{0}$ on the diagonal blocks, and $\mathbf{\Lambda}^k-\mathbf{\Lambda}^*$ is block-diagonal. Further algebraic manipulation gives us

$$\begin{aligned}
&\|\mathbf{A}^k - \mathbf{A}^*\|^2 - \|\mathbf{A}^{k+1} - \mathbf{A}^*\|^2 \\
&= \|\mathbf{H}^k - \mathbf{H}^{k+1}\|^2 + \|\mathbf{G}^{k+1} - \mathbf{H}^{k+1}\|^2 - 2\left\langle \mathbf{H}^k - \rho^{-1}(\mathbf{C} + \mathbf{\Lambda}^k) - \mathbf{G}^{k+1}, \mathbf{H}^* - \mathbf{G}^{k+1} \right\rangle \\
&- 2\rho^{-1}\left(\left\langle \mathbf{C}, \mathbf{H}^* - \mathbf{G}^{k+1} \right\rangle - \left\langle \mathbf{\Lambda}^*, \mathbf{G}^{k+1} - \mathbf{H}^{k+1} \right\rangle \right)
\end{aligned} \tag{4.26}$$

Observe that the first two terms in (4.26) are always nonnegative. Nonnegativity of the third term in (4.26) follows from the convex projection property (projection on $\mathbb{S}_+^{Md}$), i.e.,

$$\left\langle \mathbf{H}^k - \rho^{-1}(\mathbf{C} + \mathbf{\Lambda}^k) - \mathbf{G}^{k+1}, \mathbf{H}^* - \mathbf{G}^{k+1} \right\rangle \leq 0.$$

To see why the fourth term in (4.26) is nonnegative, observe that the Lagrangian for the convex program (4.24) is given by

$$\mathcal{L}(\mathbf{G}, \mathbf{H}, \mathbf{\Lambda}) = \left\langle \mathbf{C}, \mathbf{G} \right\rangle + \left\langle \mathbf{\Lambda}, \mathbf{G} - \mathbf{H} \right\rangle.$$

We have already seen in the proof of Lemma 35 that strong duality holds for C-SDP. Thus, from the saddle point property of the Lagrangian at global optimum [45], we get

$$\mathcal{L}(\mathbf{G}^*, \mathbf{H}^*, \mathbf{\Lambda}^*) \leq \mathcal{L}(\mathbf{G}^{k+1}, \mathbf{H}^{k+1}, \mathbf{\Lambda}^*),$$

which gives us

$$\left\langle \mathbf{C}, \mathbf{H}^* - \mathbf{G}^{k+1} \right\rangle - \left\langle \mathbf{\Lambda}^*, \mathbf{G}^{k+1} - \mathbf{H}^{k+1} \right\rangle \leq 0.$$

Chapter 5

Conclusion

In this thesis, we investigated theoretical and algorithmic issues that arise in the context of the rigid registration problem. We started by examining well-posedness of the problem in Chapter 2. We saw that (assuming the local coordinate measurements are exact) existence of a solution to the registration problem is guaranteed, since we have an underlying ground-truth. The more pertinent question was that of uniqueness (upto congruence) of the ground-truth solution. To investigate the question of unique registrability, we reformulated it into a question about rigidity of a graph (body graph). This allowed us to use results from graph rigidity theory, using which we obtained a linear-time-testable criterion for establishing unique registrability for planar networks. Furthermore, we resolved a conjecture on unique registrability posed in [5] by deriving its equivalent formulation in terms of the body graph; this helped us prove the conjecture for planar networks, and to confute it (through counterexamples) for three and higher dimensional networks.

Next, we turned to the scenario when the local coordinate measurements are noisy, in which case we resort to least-squares minimization LS-REG to solve the registration problem. In its original formulation, LS-REG has to be solved over a nonconvex and disconnected domain. To combat this, we first reformulated LS-REG as a rank-constrained semidefinite program REG-SDP, and then dropped the rank constraint from REG-SDP to obtain a computationally tractable convex relaxation C-SDP. The important question here was: Do we lose anything by relaxing a nonconvex program REG-SDP to a convex program C-SDP? In Chapter 3, we investigated this question, and showed that when the local coordinate measurements are not too noisy, global minimizer for the convex program C-SDP is also global minimizer for the nonconvex program REG-SDP. In this case, we do not lose anything by solving the relaxed problem (and we say that the relaxation is tight).

Finally, we considered an ADMM-based iterative solver (REG-ADMM) that can directly attack the rank-constrained semidefinite program REG-SDP. We saw that directly solving REG-SDP not only obviates the need for any sub-optimal rounding step that needs to be performed when the relaxation C-SDP is not tight, but also leads to appreciable computational savings over solving C-SDP. However, existing results on ADMM-based solvers for nonconvex programs relied on assumptions that made them inapplicable for convergence analysis of

REG-ADMM. Furthermore, numerical experiments showed that convergence behavior of REG-ADMM crucially depended on noise in the local coordinate measurements (data) and value of a solver parameter (ρ). In Chapter 4, we first showed that any fixed point of REG-ADMM is a stationary point for the registration problem. We then proved local convergence of iterates to global minimizer of REG-ADMM when the data is not too noisy; as a corollary, we derived explicit values of ρ for which REG-ADMM is guaranteed to converge to global minimizer when the data is clean. Finally, we gave a rigorous justification for the instability of REG-ADMM observed for low values of ρ when noise in the data is high.

Future Directions. Our approach to analysis of REG-ADMM provides a tidy framework within which stronger convergence results can be obtained. More precisely, we have identified "regimes" based on noise level in the data, which neatly classifies the variety of fundamentally-different convergence properties exhibited by REG-ADMM iterates. For instance, we never observe oscillations when noise in the data is low and value of parameter ρ is small; this is in sharp contrast to the situation when noise in the data is high, in which case the iterates exhibit non-attenuating oscillations whenever ρ is set below a certain threshold. This identification of different noise regimes allows us formulate precise conjectures to guide future investigations. We end this dissertation by listing two conjectures whose resolution would be of immediate practical relevance:

- In the low-noise regime, the iterates of REG-ADMM converge to global minimizer of REG-SDP if $0 < \rho < \rho_l$ for some *explicitly computable* ρ_l that depends on the data matrix $\mathbf{C}$ and the initialization. Note that we have proved such a result when the data matrix is clean (Corollary 47).
- In the high-noise regime, the iterates of REG-ADMM converge (i.e. they do not oscillate indefinitely) if $\rho > \rho_h$ for some *explicitly computable* ρ_h that depends on the data matrix $\mathbf{C}$ and the initialization. Our result (Theorem 45) would then guarantee stationarity of the point to which REG-ADMM converges.

Bibliography

[1] A. S. Bandeira, N. Boumal, and A. Singer, "Tightness of the maximum likelihood semidefinite relaxation for angular synchronization," *Mathematical Programming*, vol. 163, no. 1-2, pp. 145–167, 2017.

[2] G. Mao, B. Fidan, and B. D. O. Anderson, "Wireless sensor network localization techniques," *Computer Networks*, vol. 51, no. 10, pp. 2529–2553, 2007.

[3] Y. Shang, W. Rumi, Y. Zhang, and M. Fromherz, "Localization from connectivity in sensor networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 15, no. 11, pp. 961–974, 2004.

[4] K. Chaudhury, Y. Khoo, and A. Singer, "Large-scale sensor network localization via rigid subnetwork registration," *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 2849–2853, 2015.

[5] R. Sanyal, M. Jaiswal, and K. N. Chaudhury, "On a registration-based approach to sensor network localization," *IEEE Transactions on Signal Processing*, vol. 65, no. 20, pp. 5357–5367, 2017.

[6] I. Borg and P. Groenen, "Modern multidimensional scaling: theory and applications," *Journal of Educational Measurement*, vol. 40, no. 3, pp. 277–280, 2003.

[7] M. Cucuringu, Y. Lipman, and A. Singer, "Sensor network localization by eigenvector synchronization over the Euclidean group," *ACM Transactions on Sensor Networks*, vol. 8, no. 3, p. 19, 2012.

[8] S. Krishnan, P. Y. Lee, J. B. Moore, and S. Venkatasubramanian, "Global registration of multiple 3D point sets via optimization-on-a-manifold," *Proc. Eurographics Symposium on Geometry Processing*, pp. 187–196, 2005.

[9] S. Miraj Ahmed and K. N. Chaudhury, "Global multiview registration using non-convex ADMM," *Proc. IEEE International Conference on Image Processing*, pp. 987–991, 2017.

[10] Z. Zhang and H. Zha, "Principal manifolds and nonlinear dimensionality reduction via tangent space alignment," *SIAM Journal on Scientific Computing*, vol. 26, no. 1, pp. 313–338, 2004.

[11] V. M. Govindu and A. Pooja, "On averaging multiview relations for 3D scan registration," *IEEE Transactions on Image Processing*, vol. 23, no. 3, pp. 1289–1302, 2014.

[12] X. Fang and K.-C. Toh, "Using a distributed SDP approach to solve simulated protein molecular conformation problems," *Distance Geometry*, pp. 351–376, 2013.

[13] M. Cucuringu, A. Singer, and D. Cowburn, "Eigenvector synchronization, graph rigidity and the molecule problem," *Information and Inference: A Journal of the IMA*, vol. 1, no. 1, pp. 21–67, 2012.

[14] K. N. Chaudhury, Y. Khoo, and A. Singer, "Global registration of multiple point clouds using semidefinite programming," *SIAM Journal on Optimization*, vol. 25, no. 1, pp. 468–501, 2015.

[15] M. X. Goemans and D. P. Williamson, "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *Journal of the ACM (JACM)*, vol. 42, no. 6, pp. 1115–1145, 1995.

[16] D. G. Mixon, "Phase transitions in phase retrieval," *Excursions in Harmonic Analysis*, pp. 123–147, 2015.

[17] D. L. Donoho, A. Maleki, and A. Montanari, "The noise-sensitivity phase transition in compressed sensing," *IEEE Transactions on Information Theory*, vol. 57, no. 10, pp. 6920–6941, 2011.

[18] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[19] R. Chartrand and B. Wohlberg, "A nonconvex ADMM algorithm for group sparsity with sparse groups," *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 6009–6013, 2013.

[20] S. Diamond, R. Takapoui, and S. Boyd, "A general system for heuristic minimization of convex functions over non-convex sets," *Optimization Methods and Software*, vol. 33, no. 1, pp. 165–193, 2017.

[21] M. Hong, Z.-Q. Luo, and M. Razaviyayn, "Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems," *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 3836–3840, 2015.

[22] J. Aspnes, T. Eren, D. Goldenberg, A. Morse, A. Whiteley, R. Yang, B. Anderson, and P. Belhumeur, "A theory of network localization," *IEEE Transactions on Mobile Computing*, vol. 5, no. 12, pp. 1663–1678, 2006.

[23] I. Shames and T. Summers, "Rigid network design via submodular set function optimization," *IEEE Transactions on Network Science and Engineering*, vol. 2, no. 3, pp. 84–96, 2015.

[24] T. Eren, "Combinatorial measures of rigidity in wireless sensor and robot networks," *IEEE 54th Annual Conference on Decision and Control (CDC)*, pp. 6109–6114, 2015.

[25] S. J. Gortler, C. Gotsman, L. Liu, and D. P. Thurston, "On affine rigidity," *Journal of Computational Geometry*, vol. 4, no. 1, pp. 160–181, 2013.

[26] R. Diestel, *Graph theory.* Springer-Verlag Berlin and Heidelberg, 2000.

[27] L. Asimow and B. Roth, "The rigidity of graphs," *Transactions of the American Mathematical Society*, vol. 245, pp. 279–289, 1978.

[28] ——, "The rigidity of graphs, II," *Journal of Mathematical Analysis and Applications*, vol. 68, no. 1, pp. 171–190, 1979.

[29] R. Connelly, "Generic global rigidity," *Discrete & Computational Geometry*, vol. 33, no. 4, pp. 549–563, 2005.

[30] S. J. Gortler, A. D. Healy, and D. P. Thurston, "Characterizing generic global rigidity," *American Journal of Mathematics*, vol. 132, no. 4, pp. 897–939, 2010.

[31] B. Jackson and T. Jordán, "Connected rigidity matroids and unique realizations of graphs," *Journal of Combinatorial Theory, Series B*, vol. 94, no. 1, pp. 1–29, 2005.

[32] J. B. Saxe, "Embeddability of weighted graphs in k-space is strongly NP-Hard," *Proc. of 17th Allerton Conference in Communications, Control and Computing, Monticello, IL*, pp. 480–489, 1979.

[33] T. G. Abbott, "Generalizations of Kempe's Universality Theorem," *Master's Thesis, Massachusetts Institute of Technology*, 2008.

[34] H. Gluck, "Almost all simply connected closed surfaces are rigid," *Lecture Notes in Math*, vol. 438, pp. 225–239, 1975.

[35] G. Laman, "On graphs and rigidity of plane skeletal structures," *Journal of Engineering Mathematics*, vol. 4, no. 4, pp. 331–340, 1970.

[36] T. Jordán and Z. Szabadka, "Operations preserving the global rigidity of graphs and frameworks in the plane," *Computational Geometry*, vol. 42, no. 6-7, pp. 511–521, 2009.

[37] B. Hendrickson, "Conditions for unique graph realizations," *SIAM Journal on Computing*, vol. 21, no. 1, pp. 65–84, 1992.

[38] D. J. Jacobs and B. Hendrickson, "An algorithm for two-dimensional rigidity percolation: the pebble game," *Journal of Computational Physics*, vol. 137, no. 2, pp. 346–365, 1997.

[39] D. Jungnickel, *Graphs, Networks and Algorithms.* Springer, 2008.

[40] T. Jordán, C. Király, and S. Tanigawa, "Generic global rigidity of body-hinge frameworks," *Journal of Combinatorial Theory, Series B*, vol. 117, pp. 59–76, 2016.

[41] S. Frank and J. Jiang, "New classes of counterexamples to Hendrickson's global rigidity conjecture," *Discrete & Computational Geometry*, vol. 45, no. 3, pp. 574–591, 2011.

[42] R. Connelly and W. J. Whiteley, "Global rigidity: The effect of coning," *Discrete & Computational Geometry*, vol. 43, no. 4, pp. 717–735, 2010.

[43] P.-A. Absil, R. Mahony, and R. Sepulchre, *Optimization Algorithms on Matrix Manifolds.* Princeton University Press, 2009.

[44] E. J. Candes, T. Strohmer, and V. Voroninski, "Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming," *Communications on Pure and Applied Mathematics*, vol. 66, no. 8, pp. 1241–1274, 2013.

[45] S. Boyd and L. Vandenberghe, *Convex Optimization.* Cambridge University Press, 2004.

[46] R. Bhatia, *Matrix Analysis.* Springer Science & Business Media, 2013, vol. 169.

[47] F. Alizadeh, J.-P. A. Haeberly, and M. L. Overton, "Complementarity and nondegeneracy in semidefinite programming," *Mathematical Programming*, vol. 77, no. 1, pp. 111–128, 1997.

[48] J. Dattorro, *Convex optimization & Euclidean distance geometry.* Lulu.com, 2010.

[49] R. Sanyal, S. M. Ahmed, M. Jaiswal, and K. N. Chaudhury, "A scalable ADMM algorithm for rigid registration," *IEEE Signal Processing Letters*, vol. 24, no. 10, pp. 1453–1457, 2017.

[50] M. Hong, Z.-Q. Luo, and M. Razaviyayn, "Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems," *SIAM Journal on Optimization*, vol. 26, no. 1, pp. 337–364, 2016.

[51] Y. Wang, W. Yin, and J. Zeng, "Global convergence of ADMM in nonconvex nonsmooth optimization," *Journal of Scientific Computing*, pp. 1–35, 2015.

[52] T.-H. Chang, M. Hong, W.-C. Liao, and X. Wang, "Asynchronous distributed ADMM for large-scale optimization – part I: algorithm and convergence analysis," *IEEE Transactions on Signal Processing*, vol. 64, no. 12, pp. 3118–3130, 2016.

[53] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific Belmont, 1999.