

SE-Sync: A certifiably correct algorithm for synchronization over the special Euclidean group^{*}

Journal Title
XX(X):1–34
©The Author(s) 0000
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/



David M. Rosen¹, Luca Carlone², Afonso S. Bandeira³, and John J. Leonard⁴

Abstract

Many important geometric estimation problems naturally take the form of *synchronization over the special Euclidean group*: estimate the values of a set of unknown group elements $x_1, \dots, x_n \in \text{SE}(d)$ given noisy measurements of a subset of their pairwise relative transforms $x_i^{-1}x_j$. Examples of this class include the foundational problems of pose-graph simultaneous localization and mapping (SLAM) (in robotics), camera motion estimation (in computer vision), and sensor network localization (in distributed sensing), among others. This inference problem is typically formulated as a nonconvex maximum-likelihood estimation that is computationally hard to solve in general. Nevertheless, in this paper we present an algorithm that is able to efficiently recover *certifiably globally optimal* solutions of the special Euclidean synchronization problem in a non-adversarial noise regime. The crux of our approach is the development of a semidefinite relaxation of the maximum-likelihood estimation whose minimizer provides an *exact* MLE so long as the magnitude of the noise corrupting the available measurements falls below a certain critical threshold; furthermore, whenever exactness obtains, it is possible to *verify* this fact *a posteriori*, thereby *certifying* the optimality of the recovered estimate. We develop a specialized optimization scheme for solving large-scale instances of this semidefinite relaxation by exploiting its low-rank, geometric, and graph-theoretic structure to reduce it to an equivalent optimization problem defined on a low-dimensional Riemannian manifold, and then design a Riemannian truncated-Newton trust-region method to solve this reduction efficiently. Finally, we combine this fast optimization approach with a simple rounding procedure to produce our algorithm, *SE-Sync*. Experimental evaluation on a variety of simulated and real-world pose-graph SLAM datasets shows that SE-Sync is capable of recovering certifiably globally optimal solutions when the available measurements are corrupted by noise up to an order of magnitude greater than that typically encountered in robotics and computer vision applications, and does so more than an order of magnitude faster than the Gauss-Newton-based approach that forms the basis of current state-of-the-art techniques.

1 Introduction

The preceding decades have witnessed remarkable advances in both the capability and reliability of autonomous robots. As a result, robotics is now on the cusp of transitioning from an academic research project to a ubiquitous technology in industry and in everyday life, with tremendous potential in applications such as transportation (Thrun et al., 2006; Leonard et al., 2008; Urmson et al., 2008), medicine (Burschka et al., 2005; Taylor, 2006; Taylor et al., 2008), and disaster response (Fallon et al., 2015; Johnson et al., 2015; Pratt and Manzo, 2013) to increase productivity, alleviate suffering, and preserve life.

At the same time however, the very applications for which robotics is poised to realize the greatest benefit typically carry a correspondingly high cost of poor performance. For example, in the case of

transportation, the failure of an autonomous vehicle to function correctly may lead to destruction of property, severe injury, or even loss of life. This high

¹Oculus Research, Redmond, WA, USA

²Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, USA

³Department of Mathematics and Center for Data Science, Courant Institute of Mathematical Sciences, New York University, New York, NY, USA

⁴Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA

Corresponding author:

David M. Rosen, Oculus Research, Redmond, WA 91809, USA
Email: david.rosen@oculus.com

^{*}This report is an extended version of a paper originally presented at the 12th International Workshop on the Algorithmic Foundations of Robotics (Rosen et al., 2016).

cost of poor performance presents a serious barrier to the widespread adoption of robotic technology in the high-impact but safety-critical applications that we would most like to address, absent some guarantee of “good behavior” on the part of the autonomous agent(s). While such guarantees (viz. of correctness, optimality, bounded suboptimality, etc.), have long been a feature of algorithms for planning (Russel and Norvig, 2010) and control (Stengel, 1994; Åström and Murray, 2008), to date it appears that the development of practical algorithms with clearly-delineated performance guarantees for robotic *perception* has to a great extent remained an open problem.

This paper presents one such method, *SE-Sync*, for solving the fundamental perceptual problem of pose estimation. Formally, we address the problem of *synchronization over the special Euclidean group*: estimate the values of a set of unknown group elements $x_1, \dots, x_n \in \text{SE}(d)$ given noisy measurements of a subset of their pairwise relative transforms $x_i^{-1}x_j$.¹ This estimation problem lies at the core of many important geometric perception tasks in robotics and computer vision, including pose-graph simultaneous localization and mapping (SLAM) (Grisetti et al., 2010; Stachniss et al., 2016), camera motion and/or orientation estimation (Arrigoni et al., 2016; Hartley et al., 2013; Tron et al., 2016), and sensor network localization (Peters et al., 2015), among others.

In the context of robotics, current state-of-the-art algorithms for solving the special Euclidean synchronization problem as it arises in pose-graph SLAM formulate the problem as an instance of maximum-likelihood estimation under an assumed probability distribution for the measurement noise, and then apply general first- or second-order smooth numerical optimization methods (Nocedal and Wright, 2006) to estimate a critical point of the objective function (Grisetti et al., 2010; Stachniss et al., 2016). While this approach has been crucial in enabling the development of fast and scalable SLAM inference methods (thereby reducing SLAM to a practically-solvable problem), its reliance upon *local* search techniques leaves the resulting algorithms vulnerable to convergence to significantly suboptimal critical points. Indeed, it is not difficult to find even fairly simple real-world examples where estimates recovered by local search methods can be so poor as to be effectively unusable (Fig. 1), even for relatively low levels of measurement noise (Carlone et al., 2015a). Given the crucial role that the estimates supplied by SLAM systems play in enabling the basic functions of mobile robots, this lack of reliability in existing SLAM

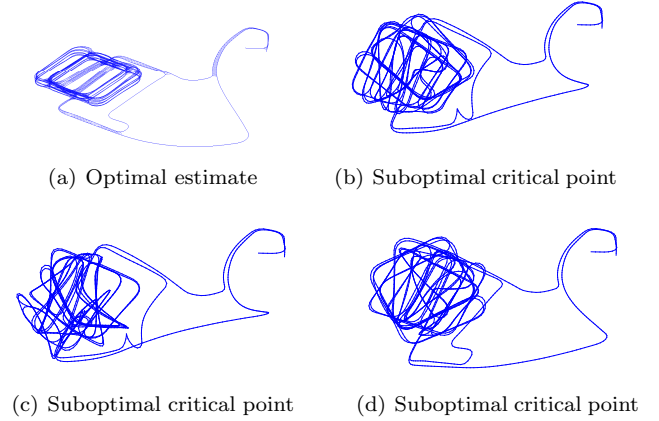


Figure 1. Examples of suboptimal estimates in pose-graph SLAM. This figure shows several estimates for the trajectory of a robot as it enters and explores a multi-level parking garage, obtained as critical points of the corresponding pose-graph SLAM maximum-likelihood estimation. Fig. 1(a) shows the globally optimal trajectory estimate, obtained using SE-Sync, while 1(b), 1(c), and 1(d) are suboptimal critical points of the MLE obtained using local search.

inference methods represents a serious barrier to the development of robust autonomous agents generally.

SE-Sync ameliorates this lack of reliability by enabling the efficient recovery of *provably globally optimal* pose-graph SLAM solutions under realistic operating conditions. Our algorithm exploits a convex *semidefinite relaxation* of the special Euclidean synchronization problem whose minimizer provides an *exact* maximum-likelihood estimate so long as the magnitude of the noise corrupting the available measurements falls below a certain critical threshold; furthermore, whenever exactness obtains, it is possible to *verify* this fact *a posteriori*, thereby *certifying* the optimality of the recovered estimate. SE-Sync thus belongs to the class of *certifiably correct* algorithms (Bandeira, 2016), meaning that it is capable of efficiently solving a generally-intractable problem within a restricted (but still practically-relevant) operational regime, and of *certifying* the correctness of the solutions that it recovers. In the case of our algorithm, experimental evaluation on a variety of simulated and real-world pose-graph SLAM datasets shows that SE-Sync is capable of recovering certifiably globally optimal pose-graph SLAM estimates when the available noise is corrupted by noise up to an order of magnitude greater than that typically encountered in robotics and computer vision applications, and does so more than an order of magnitude faster than the Gauss-Newton-based approach that forms the basis of current state-of-the-art techniques.

The rest of this paper is organized as follows. In the next section we provide an overview of prior work on the special Euclidean synchronization problem, with a particular emphasis on its instantiations in robotics and computer vision applications. In Section 3 we derive the semidefinite relaxation that will form the basis of our approach, and prove that its minimizers provide *exact* maximum-likelihood estimates for sufficiently small noise. Section 5 develops a fast optimization scheme for solving large-scale instances of the semidefinite relaxation efficiently. Finally, Section 6 provides experimental results demonstrating the advantages of our algorithm in comparison the Gauss-Newton-based approach that forms the basis of current state-of-the-art pose-graph SLAM methods. Finally, Section 7 concludes with a summary of this paper’s contributions and a discussion of future research directions.

2 Related work

In this section we provide an overview of prior work on the special Euclidean (and related) synchronization problems, focusing in particular on its instantiations in robotics and computer vision applications.

The standard formulation of the general group synchronization problem formalizes it as an instance of maximum-likelihood estimation under an assumed probability distribution for the measurement noise (Singer, 2011). This formulation is attractive from a theoretical standpoint, due to the powerful analytical framework and strong performance guarantees that maximum-likelihood estimation affords (Ferguson, 1996). However, this formal rigor often comes at the expense of computational tractability, as it is frequently the case that the optimization problem underlying an instance of maximum-likelihood estimation is nonconvex, and therefore computationally hard to solve in general. Unfortunately, this turns out to be the case for the special Euclidean synchronization problem in particular, due to the nonconvexity of $SE(d)$ itself.

Given the fundamental computational hardness of nonconvex optimization, prior work on special Euclidean synchronization (and related estimation problems) has predominantly focused on the development of *approximate* algorithms that can efficiently compute high-quality estimates in practice. These approximate algorithms can be broadly categorized into two classes.

The first class consists of algorithms that are based upon the (heuristic) application of fast *local* search techniques to identify promising estimates. This approach has proven particularly attractive in

robotics and computer vision applications, as the high computational speed of first- and second-order smooth nonlinear programming methods (Nocedal and Wright, 2006), together with their ability to exploit the measurement sparsity that typically occurs in naturalistic problem instances (Dellaert and Kaess, 2006), enables these techniques to scale effectively to large problems while maintaining real-time operation with limited computational resources. Indeed, there now exist a variety of mature algorithms and software libraries implementing this approach that are able to process special Euclidean synchronization problems involving tens to hundreds of thousands of latent states in real time using only a single thread on a commodity processor (Dellaert et al., 2010; Olson et al., 2006; Grisetti et al., 2009; Kümmerle et al., 2011; Kaess et al., 2012; Rosen et al., 2014; Lourakis and Argyros, 2009; Konolige, 2010). However, the restriction to *local* search leaves these methods vulnerable to convergence to significantly suboptimal local minima, even for relatively low levels of measurement noise (Carlone et al., 2015a). Some recent research has attempted to more explicitly tackle the problem of suboptimal convergence, including the pursuit of methods with larger basins of convergence to favorable solutions, (Olson et al., 2006; Grisetti et al., 2009), advanced initialization techniques (Carlone et al., 2015b; Martinec and Pajdla, 2007; Rosen et al., 2015), and the elucidation of deeper insights into the global structure of the SLAM optimization problem (Huang et al., 2010, 2012; Wang et al., 2012). While all of the aforementioned efforts have led to substantial improvements in the robustness of current practical SLAM techniques, ultimately they are still unable to provide any guarantees on the correctness of the solutions that they recover.

As an alternative to local search, a second class of algorithms employs *convex relaxation*: in this approach, one modifies the original estimation problem so as to obtain a convex approximation that can be (*globally*) solved efficiently. Recent work has proposed a wide variety of convex relaxations for special Euclidean synchronization and related estimation problems, including linear (Carlone et al., 2015b; Martinec and Pajdla, 2007), spectral (Singer, 2011; Arrigoni et al., 2016; Cucuringu et al., 2012; Bandeira et al., 2013) and semidefinite (Singer, 2011; Rosen et al., 2015; Wang and Singer, 2013; Özyeşil et al., 2015) formulations, among others. The advantage of these techniques is that the convex surrogates they employ generally capture the global structure of the original problem well enough that their solutions lie near high-quality regions of the search space for

the original estimation problem. However, as these surrogates are typically obtained from the original problem by relaxing constraints, their minimizers are generally infeasible for the original estimation problem, and therefore must be (potentially suboptimally) reprojected onto the original problem's feasible set.²

While SE-Sync is also (strictly speaking) a fast approximation technique based upon convex relaxation, it differs from prior art in that (as we prove in Section 4.2) it admits a (fairly large) regime of operation in which it is able to recover the *exact* optimal solution, together with a *certificate* of that solution's correctness.

3 Problem formulation

3.1 Notation and mathematical preliminaries

Miscellaneous sets: The symbols \mathbb{N} and \mathbb{R} denote the nonnegative integers and the real numbers, respectively, and we write $[n] \triangleq \{1, \dots, n\}$ for $n > 0$ as a convenient shorthand notation for sets of indexing integers. We use $|S|$ to denote the cardinality of a set S .

Differential geometry and Lie groups: We will encounter several smooth manifolds and Lie groups over the course of this paper, and will often make use of both the *intrinsic* and *extrinsic* formulations of the same manifold as convenience dictates; our notation will generally be consistent with that of (Warner, 1983) in the former case and (Guillemin and Pollack, 1974) in the latter. When considering an extrinsic realization $\mathcal{M} \subseteq \mathbb{R}^d$ of a manifold \mathcal{M} as an embedded submanifold of a Euclidean space and a function $f: \mathbb{R}^d \rightarrow \mathbb{R}$, it will occasionally be important for us to distinguish the notions of f considered as a function on \mathbb{R}^d and f considered as a function on the submanifold \mathcal{M} ; in these cases, to avoid confusion we will always reserve ∇f and $\nabla^2 f$ for the gradient and Hessian of f with respect to the usual metric on \mathbb{R}^d , and write $\text{grad } f$ and $\text{Hess } f$ to refer to the Riemannian gradient and Hessian of f considered as a function on \mathcal{M} (equipped with the metric inherited from its embedding) (Boothby, 2003; Kobayashi and Nomizu, 1996).

We let $O(d)$, $SO(d)$, and $SE(d)$ denote the orthogonal, special orthogonal, and special Euclidean groups in dimension d , respectively. For computational convenience we will often identify the (abstract) Lie groups $O(d)$ and $SO(d)$ with their realizations as the matrix groups:

$$O(d) \cong \{R \in \mathbb{R}^{d \times d} \mid R^T R = R R^T = I_d\} \quad (1a)$$

$$SO(d) \cong \{R \in \mathbb{R}^{d \times d} \mid R^T R = R R^T = I_d, \det(R) = +1\}, \quad (1b)$$

and $SE(d)$ with the semidirect product $\mathbb{R}^d \rtimes SO(d)$ with group operations:

$$(t_1, R_1) \cdot (t_2, R_2) = (t_1 + R_1 t_2, R_1 R_2). \quad (2a)$$

$$(t, R)^{-1} = (-R^{-1}t, R^{-1}). \quad (2b)$$

The set of orthonormal k -frames in \mathbb{R}^n ($k \leq n$):

$$\text{St}(k, n) \triangleq \{Y \in \mathbb{R}^{n \times k} \mid Y^T Y = I_k\} \quad (3)$$

is also a smooth compact matrix manifold, called the *Stiefel manifold*, and we equip $\text{St}(k, n)$ with the Riemannian metric induced by its embedding into $\mathbb{R}^{n \times k}$ (Absil et al., 2008, Sec. 3.3.2).

Linear algebra: In addition to the matrix groups defined above, we write $\text{Sym}(d)$ for the set of real $d \times d$ symmetric matrices; $A \succeq 0$ and $A \succ 0$ indicate that $A \in \text{Sym}(d)$ is positive semidefinite and positive definite, respectively. For general matrices A and B , $A \otimes B$ indicates the Kronecker (matrix tensor) product, A^\dagger the Moore-Penrose pseudoinverse, and $\text{vec}(A)$ the vectorization operator that concatenates the columns of A (Horn and Johnson, 1991, Sec. 4.2). We write $e_i \in \mathbb{R}^d$ and $E_{ij} \in \mathbb{R}^{m \times n}$ for the i th unit coordinate vector and (i, j) -th unit coordinate matrix, respectively, and $\mathbb{1}_d \in \mathbb{R}^d$ for the all-1's vector. Finally, $\|\cdot\|_2$, $\|\cdot\|_F$, and $\|\cdot\|_*$ denote the spectral, Frobenius, and nuclear matrix norms, respectively.

We will also frequently consider various $(d \times d)$ -block-structured matrices, and it will be useful to have specialized operators for them. To that end, given square matrices $A_i \in \mathbb{R}^{d \times d}$ for $i \in [n]$, we let $\text{Diag}(A_1, \dots, A_n)$ denote their matrix direct sum. Conversely, given a $(d \times d)$ -block-structured matrix $M \in \mathbb{R}^{dn \times dn}$ with ij -block $M_{ij} \in \mathbb{R}^{d \times d}$, we let $\text{BlockDiag}_d(M)$ denote the linear operator that extracts M 's $(d \times d)$ -block diagonal:

$$\text{BlockDiag}_d(M) \triangleq \begin{pmatrix} M_{11} & & \\ & \ddots & \\ & & M_{nn} \end{pmatrix} \quad (4)$$

and SymBlockDiag_d its corresponding symmetrized form:

$$\text{SymBlockDiag}_d(M) \triangleq \frac{1}{2} \text{BlockDiag}_d(M + M^T). \quad (5)$$

Finally, we let $\text{SBD}(d, n)$ denote the set of symmetric $(d \times d)$ -block-diagonal matrices in $\mathbb{R}^{dn \times dn}$:

$$\text{SBD}(d, n) \triangleq \{\text{Diag}(S_1, \dots, S_n) \mid S_1, \dots, S_n \in \text{Sym}(d)\}. \quad (6)$$

Graph theory: An *undirected graph* is a pair $G = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is a finite set and \mathcal{E} is a set of unordered

pairs $\{i, j\}$ with $i, j \in \mathcal{V}$ and $i \neq j$. Elements of \mathcal{V} are called *vertices* or *nodes*, and elements of \mathcal{E} are called *edges*. An edge $e = \{i, j\} \in \mathcal{E}$ is said to be *incident* to the vertices i and j ; i and j are called the *endpoints* of e . We write $\delta(v)$ for the set of edges incident to a vertex v .

A *directed graph* is a pair $\vec{G} = (\mathcal{V}, \vec{\mathcal{E}})$, where \mathcal{V} is a finite set and $\vec{\mathcal{E}} \subset \mathcal{V} \times \mathcal{V}$ is a set of *ordered* pairs (i, j) with $i \neq j$.³ As before, elements of \mathcal{V} are called *vertices* or *nodes*, and elements of $\vec{\mathcal{E}}$ are called (*directed*) *edges* or *arcs*. Vertices i and j are called the *tail* and *head* of the directed edge $e = (i, j)$, respectively; e is said to *leave* i and *enter* j (we also say that e is *incident* to i and j and that i and j are e 's *endpoints*, as in the case of undirected graphs). We let $t, h: \vec{\mathcal{E}} \rightarrow \mathcal{V}$ denote the functions mapping each edge to its tail and head, respectively, so that $t(e) = i$ and $h(e) = j$ for all $e = (i, j) \in \vec{\mathcal{E}}$. Finally, we again let $\delta(v)$ denote the set of directed edges incident to v , and $\delta^-(v)$ and $\delta^+(v)$ denote the sets of edges leaving and entering vertex v , respectively.

Given an undirected graph $G = (\mathcal{V}, \mathcal{E})$, we can construct a directed graph $\vec{G} = (\mathcal{V}, \vec{\mathcal{E}})$ from it by arbitrarily ordering the elements of each pair $\{i, j\} \in \mathcal{E}$; the graph \vec{G} so obtained is called an *orientation* of G .

A *weighted graph* is a triplet $G = (\mathcal{V}, \mathcal{E}, w)$ comprised of a graph $(\mathcal{V}, \mathcal{E})$ and a weight function $w: \mathcal{E} \rightarrow \mathbb{R}$ defined on the edges of this graph; since \mathcal{E} is finite, we can alternatively specify the weight function w by simply listing its values $\{w_e\}_{e \in \mathcal{E}}$ on each edge. Any unweighted graph can be interpreted as a weighted graph equipped with the *uniform weight function* $w \equiv 1$.

We can associate to a directed graph $\vec{G} = (\mathcal{V}, \vec{\mathcal{E}})$ with $n = |\mathcal{V}|$ and $m = |\vec{\mathcal{E}}|$ the *incidence matrix* $A(\vec{G}) \in \mathbb{R}^{n \times m}$ whose rows and columns are indexed by $i \in \mathcal{V}$ and $e \in \vec{\mathcal{E}}$, respectively, and whose elements are determined by:

$$A(\vec{G})_{ie} = \begin{cases} +1, & e \in \delta^+(i) \\ -1, & e \in \delta^-(i), \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

Similarly, we can associate to an undirected graph G an *oriented incidence matrix* $A(G)$ obtained as the incidence matrix of any of its orientations \vec{G} . We obtain a *reduced (oriented) incidence matrix* $\tilde{A}(G)$ by removing the final row from the (oriented) incidence matrix $A(G)$.

Finally, we can associate to a weighted undirected graph $G = (\mathcal{V}, \mathcal{E}, w)$ with $n = |\mathcal{V}|$ the *Laplacian matrix* $L(G) \in \text{Sym}(n)$ whose rows and columns are indexed

by $i \in \mathcal{V}$, and whose elements are determined by:

$$L(G)_{ij} = \begin{cases} \sum_{e \in \delta(i)} w(e), & i = j, \\ -w(e), & i \neq j \text{ and } e = \{i, j\} \in \mathcal{E}, \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

A straightforward computation shows that the Laplacian of a weighted graph G and the incidence matrix of one of its orientations \vec{G} are related by:

$$L(G) = A(\vec{G})W A(\vec{G})^\top, \quad (9)$$

where $W \triangleq \text{Diag}(w(e_1), \dots, w(e_m))$ is the diagonal matrix containing the weights of G 's edges.

Probability and statistics: We write $\mathcal{N}(\mu, \Sigma)$ for the multivariate Gaussian distribution with mean $\mu \in \mathbb{R}^d$ and covariance matrix $0 \preceq \Sigma \in \text{Sym}(d)$, and Langevin(M, κ) for the isotropic Langevin distribution on $\text{SO}(d)$ with mode $M \in \text{SO}(d)$ and concentration parameter $\kappa \geq 0$ (cf. Appendix A). With reference to a hidden parameter X whose value we wish to infer, we will write \underline{X} for its true (latent) value, \tilde{X} to denote a noisy observation of X , and \hat{X} to denote an estimate of X .

3.2 The special Euclidean synchronization problem

The $\text{SE}(d)$ synchronization problem consists of estimating the values of a set of n unknown group elements $x_1, \dots, x_n \in \text{SE}(d)$ given noisy measurements of m of their pairwise relative transforms $x_{ij} \triangleq x_i^{-1}x_j$ ($i \neq j$). We model the set of available measurements using an undirected graph $G = (\mathcal{V}, \mathcal{E})$ in which the nodes $i \in \mathcal{V}$ are in one-to-one correspondence with the unknown states x_i and the edges $\{i, j\} \in \mathcal{E}$ are in one-to-one correspondence with the set of available measurements, and we assume without loss of generality that G is connected.⁴ We let $\vec{G} = (\mathcal{V}, \vec{\mathcal{E}})$ be the directed graph obtained from G by fixing an orientation, and assume that a noisy measurement \tilde{x}_{ij} of the relative transform x_{ij} is obtained by sampling from the following probabilistic generative model:

$$\begin{aligned} \tilde{t}_{ij} &= t_{ij} + t_{ij}^\epsilon, & t_{ij}^\epsilon &\sim \mathcal{N}(0, \tau_{ij}^{-1} I_d), \\ \tilde{R}_{ij} &= R_{ij} R_{ij}^\epsilon, & R_{ij}^\epsilon &\sim \text{Langevin}(I_d, \kappa_{ij}), \end{aligned} \quad (10)$$

for all $(i, j) \in \vec{\mathcal{E}}$, where $x_{ij} = (t_{ij}, R_{ij})$ is the true (latent) value of x_{ij} .⁵ Finally, we define $\tilde{x}_{ji} \triangleq \tilde{x}_{ij}^{-1}$, $\kappa_{ji} \triangleq \kappa_{ij}$, $\tau_{ji} \triangleq \tau_{ij}$, and $\tilde{R}_{ji} \triangleq \tilde{R}_{ij}^\top$ for all $(i, j) \in \vec{\mathcal{E}}$.

Given a set of noisy measurements \tilde{x}_{ij} sampled from the generative model (10), a straightforward computation shows that a maximum-likelihood estimate $\hat{x}_{\text{MLE}} \in \text{SE}(d)^n$ for the states x_1, \dots, x_n is obtained as a minimizer of:⁶

Problem 1. Maximum-likelihood estimation for SE(d) synchronization.

$$p_{\text{MLE}}^* = \min_{\substack{t_i \in \mathbb{R}^d \\ R_i \in \text{SO}(d)}} \sum_{(i,j) \in \mathcal{E}} \left\{ \kappa_{ij} \|R_j - R_i \tilde{R}_{ij}\|_F^2 + \tau_{ij} \|t_j - t_i - R_i \tilde{t}_{ij}\|_2^2 \right\} \quad (11)$$

Unfortunately, Problem 1 is a high-dimensional non-convex nonlinear program, and is therefore computationally hard to solve in general. Consequently, in this paper our strategy will be to replace this problem with a (convex) *semidefinite relaxation* (Vandenberghe and Boyd, 1996), and then exploit this relaxation to search for solutions of Problem 1.

4 Forming the semidefinite relaxation

In this section we develop the semidefinite relaxation that we will solve in place of the maximum-likelihood estimation Problem 1. Our approach proceeds in two stages. We begin in Section 4.1 by developing a sequence of simplified but equivalent reformulations of Problem 1 with the twofold goal of simplifying its analysis and elucidating some of the structural correspondences between the optimization (11) and several simple graph-theoretic objects that can be constructed from the set of available measurements \tilde{x}_{ij} and the graphs G and \tilde{G} . We then exploit the simplified versions of Problem 1 so obtained to derive the semidefinite relaxation in Section 4.2.

4.1 Simplifying the maximum-likelihood estimation

Our first step is to rewrite Problem 1 in a more standard form for quadratic programs. First, define the *translational* and *rotational weight graphs* $W^\tau \triangleq (\mathcal{V}, \mathcal{E}, \{\tau_{ij}\})$ and $W^\rho \triangleq (\mathcal{V}, \mathcal{E}, \{\kappa_{ij}\})$ to be the weighted undirected graphs with node set \mathcal{V} , edge set \mathcal{E} , and edge weights τ_{ij} and κ_{ij} for $\{i, j\} \in \mathcal{E}$, respectively. The Laplacians of W^τ and W^ρ are then:

$$L(W^\tau)_{ij} = \begin{cases} \sum_{e \in \delta(i)} \tau_e, & i = j, \\ -\tau_{ij}, & \{i, j\} \in \mathcal{E}, \\ 0, & \{i, j\} \notin \mathcal{E}, \end{cases} \quad (12a)$$

$$L(W^\rho)_{ij} = \begin{cases} \sum_{e \in \delta(i)} \kappa_e, & i = j, \\ -\kappa_{ij}, & \{i, j\} \in \mathcal{E}, \\ 0, & \{i, j\} \notin \mathcal{E}. \end{cases} \quad (12b)$$

Similarly, let $L(\tilde{G}^\rho)$ denote the *connection Laplacian* for the rotational synchronization problem determined by the measurements \tilde{R}_{ij} and measurement weights

κ_{ij} for $(i, j) \in \tilde{\mathcal{E}}$; this is the symmetric $(d \times d)$ -block-structured matrix determined by (cf. (Singer and Wu, 2012; Wang and Singer, 2013)):

$$L(\tilde{G}^\rho) \in \text{Sym}(dn)$$

$$L(\tilde{G}^\rho)_{ij} \triangleq \begin{cases} d_i^\rho I_d, & i = j, \\ -\kappa_{ij} \tilde{R}_{ij}, & \{i, j\} \in \mathcal{E}, \\ 0_{d \times d}, & \{i, j\} \notin \mathcal{E}, \end{cases} \quad (13a)$$

$$d_i^\rho \triangleq \sum_{e \in \delta(i)} \kappa_e. \quad (13b)$$

Finally, let $\tilde{V} \in \mathbb{R}^{n \times dn}$ be the $(1 \times d)$ -block-structured matrix with (i, j) -blocks:

$$\tilde{V}_{ij} \triangleq \begin{cases} \sum_{e \in \delta^-(j)} \tau_e \tilde{t}_e^\top, & i = j, \\ -\tau_{ji} \tilde{t}_{ji}^\top, & (j, i) \in \tilde{\mathcal{E}}, \\ 0_{1 \times d}, & \text{otherwise,} \end{cases} \quad (14)$$

and $\tilde{\Sigma}$ the $(d \times d)$ -block-structured block-diagonal matrix determined by:

$$\begin{aligned} \tilde{\Sigma} &\triangleq \text{Diag}(\tilde{\Sigma}_1, \dots, \tilde{\Sigma}_n) \in \text{SBD}(d, n) \\ \tilde{\Sigma}_i &\triangleq \sum_{e \in \delta^-(i)} \tau_e \tilde{t}_e \tilde{t}_e^\top. \end{aligned} \quad (15)$$

Aggregating the rotational and translational states into the block matrices:

$$R \triangleq (R_1 \ \cdots \ R_n) \in \text{SO}(d)^n \subset \mathbb{R}^{d \times dn} \quad (16a)$$

$$t \triangleq \begin{pmatrix} t_1 \\ \vdots \\ t_n \end{pmatrix} \in \mathbb{R}^{dn} \quad (16b)$$

and exploiting definitions (12a)–(15), Problem 1 can be rewritten more compactly in the following standard form:

Problem 2. Maximum-likelihood estimation, QP form.

$$p_{\text{MLE}}^* = \min_{\substack{t \in \mathbb{R}^{dn} \\ R \in \text{SO}(d)^n}} \begin{pmatrix} t \\ \text{vec}(R) \end{pmatrix}^\top (M \otimes I_d) \begin{pmatrix} t \\ \text{vec}(R) \end{pmatrix}, \quad (17a)$$

$$M \triangleq \begin{pmatrix} L(W^\tau) & \tilde{V} \\ \tilde{V}^\top & L(\tilde{G}^\rho) + \tilde{\Sigma} \end{pmatrix}. \quad (17b)$$

Problem 2 is obtained from Problem 1 through a straightforward (although somewhat tedious) manipulation of the objective function (Appendix B.1).

Expanding the quadratic form in (17), we obtain:

$$p_{\text{MLE}}^* = \min_{\substack{t \in \mathbb{R}^{dn} \\ R \in \text{SO}(d)^n}} \left\{ t^\top (L(W^\top) \otimes I_d) t + 2t^\top (\tilde{V} \otimes I_d) \text{vec}(R) \right. \\ \left. + \text{vec}(R)^\top \left((L(\tilde{G}^\rho) + \tilde{\Sigma}) \otimes I_d \right) \text{vec}(R) \right\}. \quad (18)$$

Now observe that for a fixed value of R , (18) reduces to the *unconstrained* minimization of a quadratic form in the translational variable t , for which we can find a closed-form solution. This enables us to analytically eliminate t from the optimization problem (18), thereby obtaining:

Problem 3. Rotation-only maximum-likelihood estimation.

$$p_{\text{MLE}}^* = \min_{R \in \text{SO}(d)^n} \text{tr}(\tilde{Q} R^\top R) \quad (19a)$$

$$\tilde{Q} \triangleq L(\tilde{G}^\rho) + \underbrace{\tilde{\Sigma} - \tilde{V}^\top L(W^\top)^\dagger \tilde{V}}_{\tilde{Q}^\tau} \quad (19b)$$

Furthermore, given any minimizer R^* of (19), we can recover a corresponding optimal value t^* for t according to:

$$t^* = -\text{vec} \left(R^* \tilde{V}^\top L(W^\top)^\dagger \right). \quad (20)$$

The derivation of (19) and (20) from (18) is given in Appendix B.2.

Finally, we derive a simplified expression for the translational data matrix \tilde{Q}^τ appearing in (19b). Let

$$\Omega \triangleq \text{Diag}(\tau_{e_1}, \dots, \tau_{e_m}) \in \text{Sym}(m) \quad (21)$$

denote the diagonal matrix whose rows and columns are indexed by the directed edges $e \in \tilde{\mathcal{E}}$ and whose e th diagonal element gives the precision of the translational observation corresponding to that edge. Similarly, let $\tilde{T} \in \mathbb{R}^{m \times dn}$ denote the $(1 \times d)$ -block-structured matrix with rows and columns indexed by $e \in \tilde{\mathcal{E}}$ and $k \in \mathcal{V}$, respectively, and whose (e, k) -block is given by:

$$\tilde{T}_{ek} \triangleq \begin{cases} -\tilde{t}_{kj}^\top, & e = (k, j) \in \tilde{\mathcal{E}}, \\ 0_{1 \times d}, & \text{otherwise.} \end{cases} \quad (22)$$

Then Problem 3 can be rewritten as:

Problem 4. Simplified maximum-likelihood estimation.

$$p_{\text{MLE}}^* = \min_{R \in \text{SO}(d)^n} \text{tr}(\tilde{Q} R^\top R) \quad (23a)$$

$$\tilde{Q} = L(\tilde{G}^\rho) + \tilde{Q}^\tau \quad (23b)$$

$$\tilde{Q}^\tau = \tilde{T}^\top \Omega^{\frac{1}{2}} \Pi \Omega^{\frac{1}{2}} \tilde{T} \quad (23c)$$

where $\Pi \in \mathbb{R}^{m \times m}$ is the matrix of the orthogonal projection operator $\pi: \mathbb{R}^m \rightarrow \ker(A(\tilde{G})\Omega^{\frac{1}{2}})$ onto the kernel of the weighted incidence matrix $A(\tilde{G})\Omega^{\frac{1}{2}}$ of \tilde{G} . The derivation of (23c) from (19b) is presented in Appendix B.3.

The advantage of expression (23c) for \tilde{Q}^τ versus the original formulation given in (19b) is that the constituent matrices Π , Ω , and \tilde{T} in (23c) each admit particularly simple interpretations in terms of the underlying directed graph \tilde{G} and the translational data $(\tau_{ij}, \tilde{t}_{ij})$ attached to each edge $(i, j) \in \tilde{\mathcal{E}}$; our subsequent development will heavily exploit this structure.

4.2 Relaxing the maximum-likelihood estimation

In this subsection, we turn our attention to the development of a convex relaxation that will enable us to recover a global minimizer of Problem 4 in practice. We begin by relaxing the condition that $R \in \text{SO}(d)^n$, obtaining the following:

Problem 5. Orthogonal relaxation of the maximum-likelihood estimation.

$$p_{\text{O}}^* = \min_{R \in \text{O}(d)^n} \text{tr}(\tilde{Q} R^\top R). \quad (24)$$

We immediately have that $p_{\text{O}}^* \leq p_{\text{MLE}}^*$ since $\text{SO}(d)^n \subset \text{O}(d)^n$. However, we expect that this relaxation will often be *exact* in practice: since $\text{O}(d)$ is a disjoint union of two components separated by a distance of $\sqrt{2}$ under the Frobenius norm, and the values \tilde{R}_i that we wish to estimate all lie in $\text{SO}(d)$, the elements R_i^* of an estimate R^* obtained as a minimizer of (24) will still all lie in the $+1$ component of $\text{O}(d)$ so long as the elementwise estimation error in R^* satisfies $\|R_i^* - \tilde{R}_i\|_F < \sqrt{2}$ for all $i \in [n]$. This latter condition will hold so long as the noise perturbing the data matrix \tilde{Q} is not too large (cf. Appendix C.4).⁷

Now we derive the Lagrangian dual of Problem 5, using its *extrinsic* formulation:

$$p_{\text{O}}^* = \min_{R \in \mathbb{R}^{d \times dn}} \text{tr}(\tilde{Q} R^\top R) \quad (25) \\ \text{s.t. } R_i^\top R_i = I_d \quad \forall i = 1, \dots, n.$$

The Lagrangian corresponding to (25) is:

$$\mathcal{L}: \mathbb{R}^{d \times dn} \times \text{Sym}(d)^n \rightarrow \mathbb{R} \\ \mathcal{L}(R, \Lambda_1, \dots, \Lambda_n) = \text{tr}(\tilde{Q} R^\top R) - \sum_{i=1}^n \text{tr}(\Lambda_i (R_i^\top R_i - I_d)) \\ = \text{tr}(\tilde{Q} R^\top R) + \sum_{i=1}^n \text{tr}(\Lambda_i) - \text{tr}(\Lambda_i R_i^\top R_i) \quad (26)$$

where the $\Lambda_i \in \text{Sym}(d)$ are symmetric matrices of Lagrange multipliers for the (symmetric) matrix orthonormality constraints in (25). We can simplify (26) by aggregating the Lagrange multipliers Λ_i into a single direct sum matrix $\Lambda \triangleq \text{Diag}(\Lambda_1, \dots, \Lambda_n) \in \text{SBD}(d, n)$ to yield:

$$\begin{aligned} \mathcal{L}: \mathbb{R}^{d \times dn} \times \text{SBD}(d, n) &\rightarrow \mathbb{R} \\ \mathcal{L}(R, \Lambda) &= \text{tr} \left((\tilde{Q} - \Lambda) R^T R \right) + \text{tr}(\Lambda). \end{aligned} \quad (27)$$

The Lagrangian dual problem for (25) is thus:

$$\max_{\Lambda \in \text{SBD}(d, n)} \left\{ \inf_{R \in \mathbb{R}^{d \times dn}} \text{tr} \left((\tilde{Q} - \Lambda) R^T R \right) + \text{tr}(\Lambda) \right\}, \quad (28)$$

with corresponding dual function:

$$d(\Lambda) \triangleq \inf_{R \in \mathbb{R}^{d \times dn}} \text{tr} \left((\tilde{Q} - \Lambda) R^T R \right) + \text{tr}(\Lambda). \quad (29)$$

However, we observe that since

$$\text{tr} \left((\tilde{Q} - \Lambda) R^T R \right) = \text{vec}(R)^T \left((\tilde{Q} - \Lambda) \otimes I_d \right) \text{vec}(R), \quad (30)$$

then $d(\Lambda) = -\infty$ in (29) unless $(\tilde{Q} - \Lambda) \otimes I_d \succeq 0$, in which case the infimum is attained for $R = 0$. Furthermore, we have $(\tilde{Q} - \Lambda) \otimes I_d \succeq 0$ if and only if $\tilde{Q} - \Lambda \succeq 0$. Therefore, the dual problem (28) is equivalent to the following semidefinite program:

Problem 6. Primal semidefinite relaxation for SE(d) synchronization.

$$\begin{aligned} p_{\text{SDP}}^* &= \max_{\Lambda \in \text{SBD}(d, n)} \text{tr}(\Lambda) \\ \text{s.t. } &\tilde{Q} - \Lambda \succeq 0. \end{aligned} \quad (31)$$

Finally, a straightforward application of the duality theory for semidefinite programs (see Appendix B.4 for details) shows that the dual of Problem 6 is:

Problem 7. Dual semidefinite relaxation for SE(d) synchronization.

$$\begin{aligned} p_{\text{SDP}}^* &= \min_{Z \in \text{Sym}(dn)} \text{tr}(\tilde{Q}Z) \\ \text{s.t. } &Z = \begin{pmatrix} I_d & * & * & \cdots & * \\ * & I_d & * & \cdots & * \\ * & * & I_d & & * \\ \vdots & \vdots & & \ddots & \vdots \\ * & * & * & \cdots & I_d \end{pmatrix} \succeq 0. \end{aligned} \quad (32)$$

At this point, it is instructive to compare the dual semidefinite relaxation (32) with the simplified maximum-likelihood estimation (23). For any $R \in$

$\text{SO}(d)^n$, the product $Z = R^T R$ is positive semidefinite and has identity matrices along its $(d \times d)$ -block-diagonal, and so is a feasible point of (32); in other words, (32) can be regarded as a relaxation of the maximum-likelihood estimation obtained by *expanding* (23)'s *feasible set*. Consequently, if it so happens that a minimizer Z^* of Problem 7 admits a decomposition of the form $Z^* = R^{*T} R^*$ for some $R^* \in \text{SO}(d)^n$, then it is straightforward to verify that this R^* is also a minimizer of Problem 4. More precisely, we have the following:

Theorem 1. Let Z^* be a minimizer of the semidefinite relaxation Problem 7. If Z^* factors as:

$$Z^* = R^{*T} R^*, \quad R^* \in \text{O}(d)^n, \quad (33)$$

then R^* is a minimizer of Problem 5. If additionally $R^* \in \text{SO}(d)^n$, then R^* is also a minimizer of Problem 4, and $x^* = (t^*, R^*)$ (with t^* given by equation (20)) is an optimal solution of the maximum-likelihood estimation Problem 1.

Proof. Weak Lagrangian duality implies that the optimal values of Problems 5 and 6 satisfy $p_{\text{SDP}}^* \leq p_{\text{O}}^*$. But if Z^* admits the factorization (33), then R^* is also a feasible point of (24), and so we must have that $p_{\text{O}}^* \leq \text{tr}(\tilde{Q} R^{*T} R^*) = p_{\text{SDP}}^*$. This shows that $p_{\text{O}}^* = p_{\text{SDP}}^*$, and consequently that R^* is a minimizer of Problem 5, since it attains the optimal value.

Similarly, we have already established that the optimal values of Problems 4 and 5 satisfy $p_{\text{O}}^* \leq p_{\text{MLE}}^*$. But if additionally $R^* \in \text{SO}(d)^n$, then R^* is feasible for Problem 4, and so by the same logic as before we have that $p_{\text{O}}^* = p_{\text{MLE}}^*$ and R^* is a minimizer of Problem 4. The final claim now follows from the optimality of R^* for Problem 4 together with equation (20).

From a practical standpoint, Theorem 1 serves to identify a class of solutions of the (convex) semidefinite relaxation Problem 7 that correspond to *global minimizers* of the nonconvex maximum-likelihood estimation Problem 1. This naturally leads us to consider the following two questions: Under what conditions does Problem 7 admit a solution belonging to this class? And if there does exist such a solution, can we guarantee that we will be able to *recover* it by solving Problem 7 using a numerical optimization method?⁸ These questions are addressed by the following:

Proposition 2. Exact recovery via the semidefinite relaxation Problem 7. Let Q be the matrix of the form (23b) constructed using the true (latent) relative transforms $\underline{x}_{ij} = (t_{ij}, R_{ij})$ in (10). There exists a

constant $\beta \triangleq \beta(Q) > 0$ (depending upon Q) such that, if $\|\tilde{Q} - Q\|_2 < \beta$, then:

- (i) The dual semidefinite relaxation Problem 7 has a unique solution Z^* , and
- (ii) $Z^* = R^{*\top} R^*$, where $R^* \in \text{SO}(d)^n$ is a minimizer of the simplified maximum-likelihood estimation Problem 4.

This result is proved in Appendix C, using an approach adapted from (Bandeira et al., 2016).

In short, Proposition 2 guarantees that as long as the noise corrupting the available measurements \tilde{x}_{ij} in (10) is not too large (as measured by the spectral norm of the deviation of the data matrix \tilde{Q} from its exact latent value Q),⁹ we can recover a global minimizer R^* of Problem 4 (and hence also a global minimizer $x^* = (t^*, R^*)$ of the maximum-likelihood estimation Problem 1 via (20)) by solving Problem 7 using *any* numerical optimization method.

5 The SE-Sync algorithm

In light of Proposition 2, our overall strategy in this paper will be to search for exact solutions of the (hard) maximum-likelihood estimation Problem 1 by solving the (convex) semidefinite relaxation Problem 7. In order to realize this strategy as a practical algorithm, we therefore require (i) a method that is able to solve Problem 7 effectively in large-scale real-world problems, and (ii) a rounding procedure that recovers an *optimal* solution of Problem 1 from a solution of Problem 7 when exactness obtains, and a *feasible approximate solution* otherwise. In this section, we develop a pair of algorithms that fulfill these requirements. Together, these procedures comprise *SE-Sync*, our proposed algorithm for synchronization over the special Euclidean group.

5.1 Solving the semidefinite relaxation

As a semidefinite program, Problem 7 can in principle be solved in polynomial time using interior-point methods (Vandenberghe and Boyd, 1996; Todd, 2001). In practice, however, the high computational cost of general-purpose semidefinite programming algorithms prevents these methods from scaling effectively to problems in which the dimension of the decision variable Z is greater than a few thousand (Todd, 2001). Unfortunately, typical instances of Problem 7 arising in (for example) robotics and computer vision applications are one to two orders of magnitude larger than this maximum effective problem size, and are therefore well beyond the reach of these general-purpose methods. To overcome this limitation, in

this subsection we develop a specialized optimization procedure for solving large-scale instances of Problem 7 efficiently. We first exploit this problem's low-rank, geometric, and graph-theoretic structure to reduce it to an equivalent optimization problem defined on a low-dimensional Riemannian manifold (Boothby, 2003; Kobayashi and Nomizu, 1996), and then design a fast Riemannian optimization method to solve this reduction efficiently.

5.1.1 Simplifying Problem 7

Exploiting low-rank structure: The dominant computational cost when applying general-purpose semidefinite programming methods to solve Problem 7 is the need to store and manipulate expressions involving the (large, dense) matrix variable Z . In particular, the $O(n^3)$ computational cost of multiplying and factoring such expressions quickly becomes intractable as the problem size n increases. On the other hand, in the case that exactness holds, we know that the actual solution Z^* of Problem 7 that we seek has a very concise description in the factored form $Z^* = R^{*\top} R^*$ for $R^* \in \text{SO}(d)^n$. More generally, even in those cases where exactness fails, minimizers Z^* of Problem 7 typically have a rank r not much greater than d , and therefore admit a symmetric rank decomposition $Z^* = Y^{*\top} Y^*$ for $Y^* \in \mathbb{R}^{r \times dn}$ with $r \ll dn$.

In a pair of papers, (Burer and Monteiro, 2003, 2005) proposed an elegant general approach to exploit the fact that large-scale semidefinite programs often admit such low-rank solutions: simply replace every instance of the decision variable Z with a rank- r product of the form $Y^\top Y$ to produce a *rank-restricted* version of the original problem. This substitution has the two-fold effect of (i) dramatically reducing the size of the search space and (ii) rendering the positive semidefiniteness constraint *redundant*, since $Y^\top Y \succeq 0$ for *any* choice of Y . The resulting rank-restricted form of the problem is thus a low-dimensional *nonlinear* program, rather than a *semidefinite* program. In the specific case of Problem 7, this produces:

Problem 8. Rank-restricted semidefinite relaxation, NLP form.

$$\begin{aligned} p_{\text{SDPLR}}^* &= \min_{Y \in \mathbb{R}^{r \times dn}} \text{tr}(\tilde{Q} Y^\top Y) \\ \text{s.t. } &\text{BlockDiag}_d(Y^\top Y) = \text{Diag}(I_d, \dots, I_d). \end{aligned} \quad (34)$$

Provided that Problem 7 admits a solution Z^* with $\text{rank}(Z^*) \leq r$, we can *recover* such a solution from an optimal solution Y^* of Problem 8 according to $Z^* = Y^{*\top} Y^*$.

Exploiting geometric structure: In addition to exploiting Problem 7's low-rank structure, following (Boumal, 2015) we also observe that the specific form of the constraints appearing in Problems 7 and 8 (i.e., that the $d \times d$ block-diagonals of Z and $Y^\top Y$ must be I_d) admits a nice geometric interpretation that can be exploited to further simplify Problem 8. Introducing the block decomposition:

$$Y \triangleq (Y_1 \quad \dots \quad Y_n) \in \mathbb{R}^{r \times dn}, \quad (35)$$

the block-diagonal constraints appearing in (34) are equivalent to:

$$Y_i^\top Y_i = I_d, \quad Y_i \in \mathbb{R}^{r \times d}, \quad (36)$$

i.e., they require that each Y_i be an element of the Stiefel manifold $\text{St}(d, r)$ in (3). Consequently, Problem 8 can be equivalently formulated as an *unconstrained* Riemannian optimization problem on a product of Stiefel manifolds:

Problem 9. Rank-restricted semidefinite relaxation, Riemannian optimization form.

$$p_{\text{SDPLR}}^* = \min_{Y \in \text{St}(d, r)^n} \text{tr}(\tilde{Q} Y^\top Y). \quad (37)$$

This is the optimization problem that we will actually solve in practice.

Exploiting graph-theoretic structure: The reduction of Problem 7 to Problem 9 obviates the need to form or manipulate the large, dense matrix variable Z directly. However the data matrix \tilde{Q} that parameterizes each of Problems 4–9 is also dense and of the same order as Z , and so presents a similar computational difficulty. Accordingly, here we develop an analogous concise description of \tilde{Q} in terms of sparse matrices (and their inverses) associated with the graph \vec{G} .

Equations (23b) and (23c) provide a decomposition of \tilde{Q} in terms of the sparse matrices $L(\vec{G}^\rho)$, \tilde{T} , and Ω , and the dense orthogonal projection matrix Π . However, since Π is also a matrix derived from a sparse graph, we might suspect that it too should admit some kind of sparse description. And indeed, it turns out that Π admits a sparse decomposition as:

$$\underline{A}(\vec{G})\Omega^{\frac{1}{2}} = LQ_1 \quad (38a)$$

$$\Pi = I_m - \Omega^{\frac{1}{2}} \underline{A}(\vec{G})^\top L^{-\top} L^{-1} \underline{A}(\vec{G})\Omega^{\frac{1}{2}} \quad (38b)$$

where equation (38a) is a thin LQ decomposition¹⁰ of the weighted reduced incidence matrix $\underline{A}(\vec{G})\Omega^{\frac{1}{2}}$ of \vec{G} . This result is derived in Appendix B.3. Note that expression (38b) for Π requires only the sparse lower-triangular factor L from (38a), which can be easily and

efficiently obtained (e.g. by applying successive Givens rotations (Golub and Loan, 1996, Sec. 5.2.1) directly to $\underline{A}(\vec{G})\Omega^{\frac{1}{2}}$ itself).

Together, equations (23b), (23c), and (38b) provide a concise description of \tilde{Q} in terms of sparse matrices, as desired. We will exploit this decomposition in Section 5.1.3 to design a fast Riemannian optimization method for solving Problem 9.

5.1.2 The Riemannian Staircase At this point, it is again instructive to compare Problem 9 with the simplified maximum-likelihood estimation Problem 4 and its relaxation Problem 7. Since the (special) orthogonal matrices satisfy condition (3) with $k = n = d$, we have the set of inclusions:

$$\text{SO}(d) \subset \text{O}(d) = \text{St}(d, d) \subset \text{St}(d, d+1) \subset \dots \quad (39)$$

and we can therefore view the set of rank-restricted Riemannian optimization problems (37) as comprising a *hierarchy* of relaxations of the maximum-likelihood estimation (23) that are intermediate between Problem 5 and Problem 7 for $d < r < dn$. However, unlike Problem 7, the various instantiations of Problem 9 are *nonconvex* due to the (re)introduction of the quadratic orthonormality constraints (3). It may therefore not be clear whether anything has really been gained by relaxing Problem 4 to Problem 9, since it appears that we may have simply replaced one difficult nonconvex optimization problem with another. The following remarkable result (Corollary 8 of (Boumal et al., 2016b)) justifies this approach:

Proposition 3. A sufficient condition for global optimality in Problem 9. *If $Y \in \text{St}(d, r)^n$ is a (row) rank-deficient second-order critical point¹¹ of Problem 9, then Y is a global minimizer of Problem 9 and $Z^* = Y^\top Y$ is a solution of the dual semidefinite relaxation Problem 7.*

Proposition 3 immediately suggests an algorithm for recovering solutions Z^* of Problem 7 from Problem 9: simply apply a second-order Riemannian optimization method to search successively higher levels of the hierarchy of relaxations (37) until a *rank-deficient* second-order critical point is obtained.¹² This algorithm, the *Riemannian Staircase* (Boumal, 2015; Boumal et al., 2016b), is summarized as Algorithm 1. We emphasize that while Algorithm 1 may require searching up to $O(dn)$ levels of the hierarchy (37) in the worst case, in practice this a gross overestimate; typically one or two “stairs” suffice.

[†]For example, the second-order Riemannian trust-region method (Boumal et al., 2016a, Algorithm 3).

Algorithm 1 The Riemannian Staircase

Input: An initial point $Y \in \text{St}(d, r_0)^n$, $r_0 \geq d + 1$.
Output: A minimizer Y^* of Problem 9 corresponding to a solution $Z^* = Y^{*\top} Y^*$ of Problem 7.

```

1: function RIEMANNIANSTAIRCASE( $Y$ )
2:   for  $r = r_0, \dots, dn + 1$  do
3:     Starting at  $Y$ , apply a Riemannian
       optimization method† to identify a second-
       order critical point  $Y^* \in \text{St}(d, r)^n$  of
       Problem 9.
4:     if  $\text{rank}(Y^*) < r$  then
5:       return  $Y^*$ 
6:     else
7:       Set  $Y \leftarrow \begin{pmatrix} Y^* \\ 0_{1 \times dn} \end{pmatrix}$ .
8:     end if
9:   end for
10: end function

```

5.1.3 A Riemannian optimization method for Problem 9 Proposition 3 and the Riemannian Staircase (Algorithm 1) provide a means of obtaining *global* minimizers of Problem 7 by *locally* searching for second-order critical points of Problem 9. In this subsection, we design a Riemannian optimization method that will enable us to rapidly identify these critical points in practice.

Equations (23b), (23c), and (38b) provide an efficient means of computing products with \tilde{Q} *without* the need to form \tilde{Q} explicitly by performing a sequence of sparse matrix multiplications and sparse triangular solves. This operation is sufficient to evaluate the objective appearing in Problem 9, as well as its gradient and Hessian-vector products when it is considered as a function on the ambient Euclidean space $\mathbb{R}^{r \times dn}$:

$$F(Y) \triangleq \text{tr}(\tilde{Q}Y^\top Y) \quad (40a)$$

$$\nabla F(Y) = 2Y\tilde{Q} \quad (40b)$$

$$\nabla^2 F(Y)[\dot{Y}] = 2\dot{Y}\tilde{Q}. \quad (40c)$$

Furthermore, there are simple relations between the ambient Euclidean gradient and Hessian-vector products in (40b) and (40c) and their corresponding Riemannian counterparts when $F(\cdot)$ is viewed as a function restricted to the embedded submanifold $\text{St}(d, r)^n \subset \mathbb{R}^{r \times dn}$. With reference to the orthogonal projection operator onto the tangent space of $\text{St}(d, r)^n$ at Y (Edelman et al., 1998, eq. (2.3)):

$$\begin{aligned} \text{Proj}_Y: T_Y(\mathbb{R}^{r \times dn}) &\rightarrow T_Y(\text{St}(d, r)^n) \\ \text{Proj}_Y(X) &= X - Y \text{SymBlockDiag}_d(Y^\top X) \end{aligned} \quad (41)$$

the Riemannian gradient $\text{grad} F(Y)$ is simply the orthogonal projection of the ambient Euclidean

gradient $\nabla F(Y)$ (cf. (Absil et al., 2008, eq. (3.37))):

$$\text{grad} F(Y) = \text{Proj}_Y \nabla F(Y). \quad (42)$$

Similarly, the Riemannian Hessian-vector product $\text{Hess} F(Y)[\dot{Y}]$ can be obtained as the orthogonal projection of the ambient directional derivative of the gradient vector field $\text{grad} F(Y)$ in the direction of \dot{Y} (cf. (Absil et al., 2008, eq. (5.15))). A straightforward computation shows that this is given by:¹³

$$\begin{aligned} \text{Hess} F(Y)[\dot{Y}] &= \text{Proj}_Y \left(D[\text{grad} F(Y)][\dot{Y}] \right) \\ &= \text{Proj}_Y \begin{pmatrix} \nabla^2 F(Y)[\dot{Y}] \\ -\dot{Y} \text{SymBlockDiag}_d(Y^\top \nabla F(Y)) \end{pmatrix}. \end{aligned} \quad (43)$$

Equations (23b), (23c), and (38)–(43) provide an efficient means of computing $F(Y)$, $\text{grad} F(Y)$, and $\text{Hess} F(Y)[\dot{Y}]$. Consequently, we propose to employ a *truncated-Newton trust-region* optimization method (Dembo and Steihaug, 1983; Steihaug, 1983; Nash, 2000) to solve Problem 9; this approach will enable us to exploit the availability of an efficient routine for computing Hessian-vector products $\text{Hess} F(Y)[\dot{Y}]$ to implement a second-order optimization method *without* the need to explicitly form or factor the dense matrix $\text{Hess} F(Y)$. Moreover, truncated-Newton methods comprise the current state of the art for superlinear large-scale unconstrained nonlinear programming (Nocedal and Wright, 2006, Sec. 7.1), and are therefore ideally suited for solving large-scale instances of (37). Accordingly, we will apply the truncated-Newton *Riemannian Trust-Region* (RTR) method (Absil et al., 2007; Boumal et al., 2016a) to efficiently compute high-precision¹⁴ estimates of second-order critical points of Problem 9.

5.2 Rounding the solution

In the previous subsection, we described an efficient algorithmic approach for computing minimizers $Y^* \in \text{St}(d, r)^n$ of Problem 9 that correspond to solutions $Z^* = Y^{*\top} Y^*$ of Problem 7. However, our ultimate goal is to extract an optimal solution $R^* \in \text{SO}(d)^n$ of Problem 4 from Z^* whenever exactness holds, and a *feasible approximate solution* $\hat{R} \in \text{SO}(d)^n$ otherwise. In this subsection, we develop a rounding procedure satisfying these criteria. To begin, let us consider the case in which exactness obtains; here:

$$Y^{*\top} Y^* = Z^* = R^{*\top} R^* \quad (44)$$

for some optimal solution $R^* \in \text{SO}(d)^n$ of Problem 4. Since $\text{rank}(R^*) = d$, this implies that $\text{rank}(Y^*) = d$ as

Algorithm 2 Rounding procedure for solutions of Problem 9

Input: An optimal solution $Y^* \in \text{St}(d, r)^n$ of Problem 9.

Output: A feasible point $\hat{R} \in \text{SO}(d)^n$.

```

1: function ROUNDSOLUTION( $Y^*$ )
2:   Compute a rank- $d$  truncated singular value
   decomposition  $U_d \Xi_d V_d^\top$  for  $Y^*$  and assign
    $\hat{R} \leftarrow \Xi_d V_d^\top$ .
3:   Set  $N_+ \leftarrow |\{\hat{R}_i \mid \det(\hat{R}_i) > 0\}|$ .
4:   if  $N_+ < \lceil \frac{n}{2} \rceil$  then
5:      $\hat{R} \leftarrow \text{Diag}(1, \dots, 1, -1) \hat{R}$ .
6:   end if
7:   for  $i = 1, \dots, n$  do
8:     Set  $\hat{R}_i \leftarrow \text{NEARESTROTATION}(\hat{R}_i)$ .
9:   end for
10:  return  $\hat{R}$ 
11: end function

```

well. Consequently, letting

$$Y^* = U_d \Xi_d V_d^\top \quad (45)$$

denote a (rank- d) thin singular value decomposition (Golub and Loan, 1996, Sec. 2.5.3) of Y^* , and defining

$$\bar{Y} \triangleq \Xi_d V_d^\top \in \mathbb{R}^{d \times dn}, \quad (46)$$

it follows from substituting (45) into (44) that

$$\bar{Y}^\top \bar{Y} = Z^* = R^{*\top} R^*. \quad (47)$$

Equation (47) implies that the $d \times d$ block-diagonal of $\bar{Y}^\top \bar{Y}$ satisfies $\bar{Y}_i^\top \bar{Y}_i = I_d$ for all $i \in [n]$, i.e. $\bar{Y} \in \text{O}(d)^n$. Similarly, comparing the elements of the first block rows of $\bar{Y}^\top \bar{Y}$ and $R^{*\top} R^*$ in (47) shows that $\bar{Y}_1^\top \bar{Y}_j = R_1^{*\top} R_j^*$ for all $j \in [n]$. Left-multiplying this set of equalities by \bar{Y}_1 and letting $A = \bar{Y}_1 R_1^*$ then gives:

$$\bar{Y} = AR^*, \quad A \in \text{O}(d). \quad (48)$$

Since any product of the form AR^* with $A \in \text{SO}(d)$ is *also* an optimal solution of Problem 4 (by gauge symmetry), equation (48) shows that \bar{Y} as defined in (46) is optimal provided that $\bar{Y} \in \text{SO}(d)$ specifically. Furthermore, if this is not the case, we can always make it so by left-multiplying \bar{Y} by any orientation-reversing element of $\text{O}(d)$, for example $\text{Diag}(1, \dots, 1, -1)$. Thus, equations (45)–(48) give a straightforward means of recovering an optimal solution of Problem 4 from Y^* whenever exactness holds.

Moreover, this procedure can be straightforwardly generalized to the case that exactness fails, thereby producing a convenient rounding scheme. Specifically, we can consider the right-hand side of (46) as taken

Algorithm 3 The SE-Sync algorithm

Input: An initial point $Y \in \text{St}(d, r_0)^n$, $r_0 \geq d + 1$.

Output: A feasible estimate $\hat{x} \in \text{SE}(d)^n$ for the maximum-likelihood estimation Problem 1 and the lower bound p_{SDP}^* for Problem 1's optimal value.

```

1: function SE-SYNC( $Y$ )
2:   Set  $Y^* \leftarrow \text{RIEMANNIANSTAIRCASE}(Y)$ .
3:   Set  $p_{\text{SDP}}^* \leftarrow F(\tilde{Q} Y^{*\top} Y^*)$ .
4:   Set  $\hat{R} \leftarrow \text{ROUNDSOLUTION}(Y^*)$ .
5:   Recover the optimal translational estimates  $\hat{t}$ 
   corresponding to  $\hat{R}$  via (20).
6:   Set  $\hat{x} \leftarrow (\hat{t}, \hat{R})$ .
7:   return  $\{\hat{x}, p_{\text{SDP}}^*\}$ 
8: end function

```

from a rank- d *truncated* singular value decomposition of Y^* (so that \bar{Y} is an orthogonal transform of the best rank- d approximation of Y^*), multiply \bar{Y} by an orientation-reversing element of $\text{O}(d)$ according to whether a majority of its block elements have positive or negative determinant, and then project each of \bar{Y} 's blocks to the nearest rotation matrix.¹⁵ This generalized rounding scheme is formalized as Algorithm 2.

5.3 The complete algorithm

Combining the efficient optimization approach of Section 5.1 with the rounding procedure of Section 5.2 produces *SE-Sync* (Algorithm 3), our proposed algorithm for synchronization over the special Euclidean group.

When applied to an instance of $\text{SE}(d)$ synchronization, SE-Sync returns a feasible point $\hat{x} \in \text{SE}(d)^n$ for the maximum-likelihood estimation Problem 1 and the lower bound $p_{\text{SDP}}^* \leq p_{\text{MLE}}^*$ for Problem 1's optimal value. This lower bound provides an *upper* bound on the suboptimality of *any* feasible point $x = (t, R) \in \text{SE}(d)^n$ as a solution of Problem 1 according to:

$$F(\tilde{Q} R^\top R) - p_{\text{SDP}}^* \geq F(\tilde{Q} R^\top R) - p_{\text{MLE}}^*. \quad (49)$$

Furthermore, in the case that Problem 7 is exact, the estimate $\hat{x} = (\hat{t}, \hat{R}) \in \text{SE}(d)^n$ returned by Algorithm 3 *attains* this lower bound:

$$F(\tilde{Q} \hat{R}^\top \hat{R}) = p_{\text{SDP}}^*. \quad (50)$$

Consequently, verifying *a posteriori* that (50) holds provides a *computational certificate* of \hat{x} 's correctness as a solution of the maximum-likelihood estimation Problem 1. SE-Sync is thus a *certifiably correct* algorithm for $\text{SE}(d)$ synchronization, as claimed.

6 Experimental results

In this section we evaluate SE-Sync’s performance on a variety of special Euclidean synchronization problems drawn from 3D pose-graph simultaneous localization and mapping (SLAM). We consider two versions of the algorithm that differ in their approach to evaluating products with the orthogonal projection matrix Π : the first (SE-Sync-Chol) employs the cached Cholesky factor L in (38a) to evaluate this product by working right to left through the sequence of sparse matrix multiplications and upper- and lower-triangular solves in (38b), while the second (SE-Sync-QR) follows (85), using QR decomposition to solve the linear least-squares problem for w^* . As a basis for comparison, we also evaluate the performance of the Gauss-Newton method (GN), the *de facto* standard approach for solving pose-graph SLAM problems in robotics (Kümmerle et al., 2011; Kaess et al., 2012; Grisetti et al., 2010).

All of the following experiments are performed on a Dell Precision 5510 laptop with an Intel Xeon E3-1505M 2.80 GHz processor and 16 GB of RAM running Ubuntu 16.04. Our experimental implementations of SE-Sync[‡] and the Gauss-Newton method are written in MATLAB, and the former takes advantage of the truncated-Newton RTR method (Absil et al., 2007; Boumal et al., 2016a) supplied by the Manopt toolbox (Boumal et al., 2014a) to implement the Riemannian optimization method of Section 5.1.3. Each optimization algorithm is limited to a maximum of 500 (outer) iterations, and convergence is declared whenever the norm of the (Riemannian) gradient is less than 10^{-2} or the relative decrease in function value between two subsequent (accepted) iterations is less than 10^{-5} ; additionally, each outer iteration of RTR is limited to a maximum of 500 Hessian-vector product operations. The Gauss-Newton method is initialized using the *chordal initialization*, a state-of-the-art method for bootstrapping an initial solution in SLAM and bundle adjustment problems (Carlone et al., 2015b; Martinec and Pajdla, 2007), and we set $r_0 = 5$ in the Riemannian Staircase (Algorithm 1). Finally, since SE-Sync is based upon solving the (convex) semidefinite relaxation Problem 7, it does not require a high-quality initialization in order to reach a globally optimal solution; nevertheless, it can still benefit (in terms of reduced computation time) from being supplied with one. Consequently, in the following experiments we employ two initialization procedures in conjunction with each version of SE-Sync: the first (**rand**) simply samples a point uniformly randomly from $\text{St}(d, r_0)^n$, while the second (**chord**) supplies the same chordal initialization that the Gauss-Newton method

receives, in order to enable a fair comparison of the algorithms’ computational speeds.

6.1 Cube experiments

In this first set of experiments, we are interested in investigating how the performance of SE-Sync is affected by factors such as measurement noise, measurement density, and problem size. To that end, we conduct a set of simulation studies that enable us to interrogate each of these factors individually. Concretely, we revisit the **cube** experiments considered in our previous work (Carlone et al., 2015a); this scenario simulates a robot traveling along a rectilinear path through a regular cubical lattice with a side length of s poses (Fig. 2). An odometric measurement is available between each pair of sequential poses, and measurements between nearby nonsequential poses are available with probability p_{LC} ; the measurement values $\tilde{x}_{ij} = (\tilde{t}_{ij}, \tilde{R}_{ij})$ themselves are sampled according to (10). We fix default values for these parameters at $\kappa = 16.67$ (corresponding to an expected angular RMS error of 10 degrees for the rotational measurements \tilde{R}_{ij} , cf. (62) in Appendix A), $\tau = 75$ (corresponding to an expected RMS error of .20 m for the translational measurements \tilde{t}_{ij}), $p_{LC} = .1$, and $s = 10$ (corresponding to a default problem size of 1000 poses), and consider the effect of varying each of them individually; our complete dataset consists of 50 realizations of the **cube** sampled from the generative model just described for *each* joint setting of the parameters κ , τ , p_{LC} , and s . Results for these experiments are shown in Fig. 3.

Consistent with our previous findings (Carlone et al., 2015a), these results suggest that the exactness of the semidefinite relaxation (32) depends primarily upon the level of noise corrupting the rotational observations \tilde{R}_{ij} in (10). Furthermore, we see from Fig. 3(a) that in these experiments, exactness obtains for rotational noise with a root-mean-square angular error up to about 15 degrees; this is roughly an order of magnitude greater than the level of noise affecting sensors typically deployed in robotics and computer vision applications, which provides strong empirical evidence that SE-Sync is capable of recovering certifiably globally optimal solutions of pose-graph SLAM problems under “reasonable” operating conditions.¹⁶

In addition to its ability to recover certifiably optimal solutions, these experiments also show that

[‡]Available online: <https://github.com/david-m-rosen/SE-Sync>.

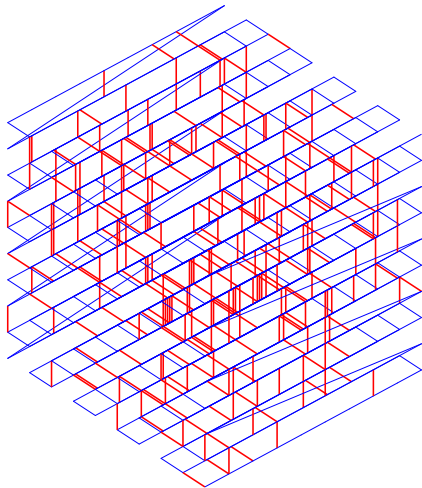


Figure 2. The ground truth configuration of an instance of the cube dataset with $s = 10$ and $p_{LC} = .1$. The robot's trajectory (with associated odometric measurements) is drawn in blue, and loop closure observations in red.

SE-Sync is considerably faster than the Gauss-Newton approach that underpins current state-of-the-art pose-graph SLAM algorithms (Kümmerle et al., 2011; Kaess et al., 2012; Grisetti et al., 2010; Rosen et al., 2014). Indeed, examining the plots in the center column of Fig. 3 reveals that in almost all of our experimental scenarios, SE-Sync is able to both compute *and* certify a globally optimal solution in less time than the Gauss-Newton method requires to simply *obtain* an estimate when supplied with a high-quality chordal initialization. Moreover, restricting attention only to those scenarios in which SE-Sync also employs the chordal initialization (in order to enable a fair comparison with the Gauss-Newton method), we see that both SE-Sync-Chol and SE-Sync-QR are consistently many-fold faster than Gauss-Newton, and that this speed differential increases superlinearly with the problem size (Fig. 3(d)).

Given that SE-Sync performs direct *global* optimization, whereas Gauss-Newton is a purely *local* search technique, it may be somewhat counterintuitive that the former is consistently so much faster than the latter. However, we can attribute this improved performance to two critical design decisions that distinguish these techniques. First, SE-Sync makes use of the *exact* Hessian (cf. Section 5.1.3), whereas Gauss-Newton, by construction, implicitly uses an approximation whose quality degrades in the presence of either large measurement residuals *or* strong nonlinearities in the underlying objective function (cf. e.g. (Rosen et al., 2014, Sec. III-B)), both of which are typical features of SLAM problems. This implies that the quadratic model function (cf. e.g. (Nocedal and

Wright, 2006, Chp. 2)) that SE-Sync employs better captures the shape of the underlying objective than the one used by Gauss-Newton, so that the former is capable of computing higher-quality update steps. Second, and more significantly, SE-Sync makes use of a truncated-Newton method (RTR), which avoids the need to explicitly form or factor the (Riemannian) Hessian $\text{Hess } F(Y)$; instead, at each iteration this approach *approximately* solves the Newton equations using a truncated conjugate gradient algorithm (Golub and Loan, 1996, Chp. 10),¹⁷ and only computes this approximate solution as accurately as is necessary to ensure adequate progress towards a critical point with each applied update. The result is that RTR requires only a few sparse matrix-vector multiplications in order to obtain each update step; moreover, equations (23b), (23c), and (38) show that the constituent matrices involved in these products are *constant*, and can therefore be precomputed and cached at the beginning of the algorithm. In contrast, the Gauss-Newton method must recompute and refactor the Jacobian at *each* iteration, which is considerably more expensive.

6.2 SLAM benchmark datasets

The experiments in the previous section made extensive use of simulated cube datasets to investigate the effects of measurement noise, measurement density, and problem size on SE-Sync's performance. In this next set of experiments, we evaluate SE-Sync on a suite of larger and more heterogeneous 3D pose-graph SLAM benchmarks that better represent the distribution of problems encountered in real-world SLAM applications. The first three of these (the **sphere**, **torus**, and **grid** datasets) are also synthetic (although generated using an observation model different from (10)), while the latter three (the **garage**, **cubicle**, and **rim** datasets) are large-scale real-world examples (Fig. 4). For the purpose of these experiments, we restrict attention to the version of SE-Sync employing QR factorization and the chordal initialization (SE-Sync-QR + chord), and once again compare it with the Gauss-Newton method (GN + chord). Results for these experiments are shown in Table 1.

On each of these examples, both SE-Sync and the Gauss-Newton method converged to the same (globally optimal) solution; however, consistent with our findings in Section 6.1, SE-Sync did so considerably faster, outperforming Gauss-Newton in terms of computational speed by a factor of between 3.3 (on **garage**, the smallest dataset) and 15.7 (on **rim**, the largest dataset). These results further support our claim that SE-Sync provides an effective means of recovering globally optimal pose-graph SLAM

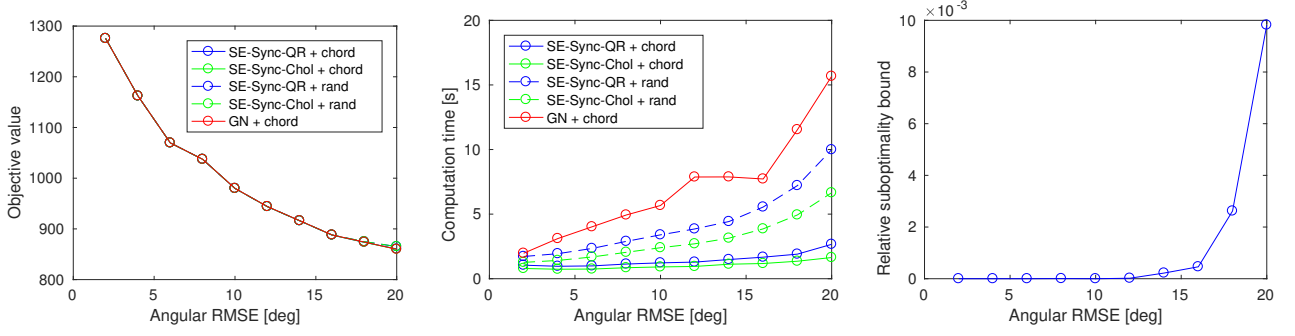
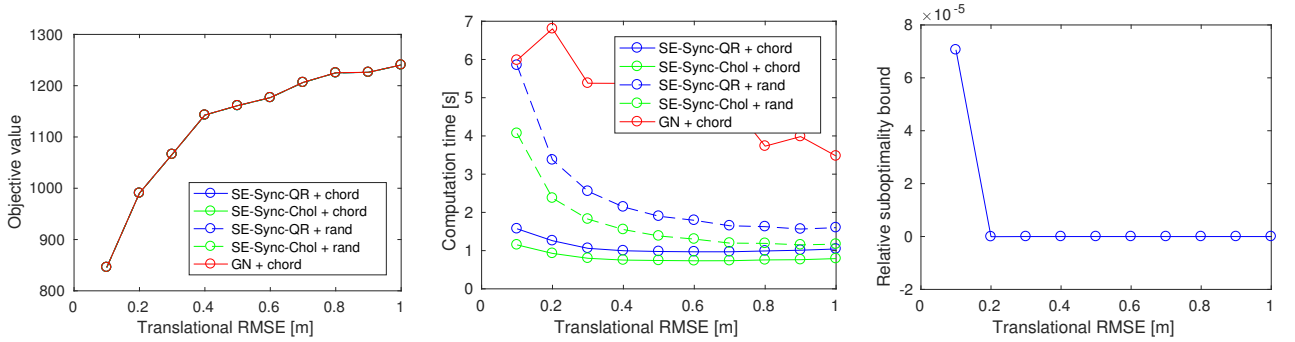
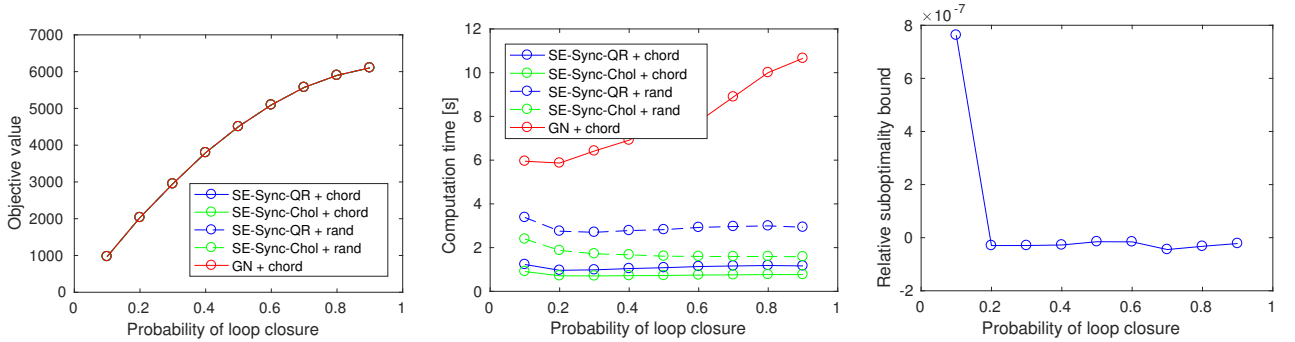
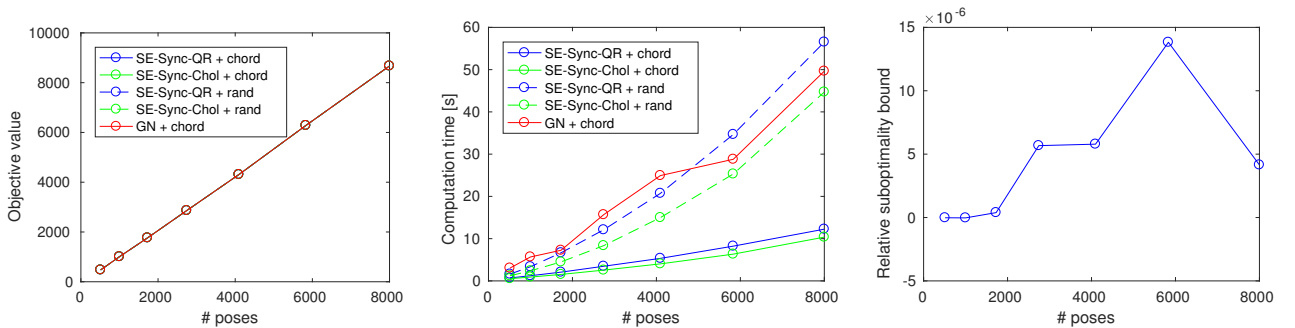
(a) Varying κ while holding $\tau = 75$, $p_{LC} = .1$, $s = 10$.(b) Varying τ while holding $\kappa = 16.67$, $p_{LC} = .1$, $s = 10$.(c) Varying p_{LC} while holding $\kappa = 16.67$, $\tau = 75$, $s = 10$.(d) Varying s while holding $\kappa = 16.67$, $\tau = 75$, $p_{LC} = .1$.

Figure 3. Results for the cube experiments. These figures plot the mean of the objective values (left column) and elapsed computation times (center column) attained by the Gauss-Newton and SE-Sync algorithms, as well as the upper bound $(F(\hat{Q}\hat{R}^T\hat{R}) - p_{SDP}^*)/p_{SDP}^*$ for the relative suboptimality of the solution recovered by SE-Sync (right column), for 50 realizations of the cube dataset as functions of the measurement precisions κ (first row) and τ (second row), the loop closure probability p_{LC} (third row), and the problem size (fourth row).

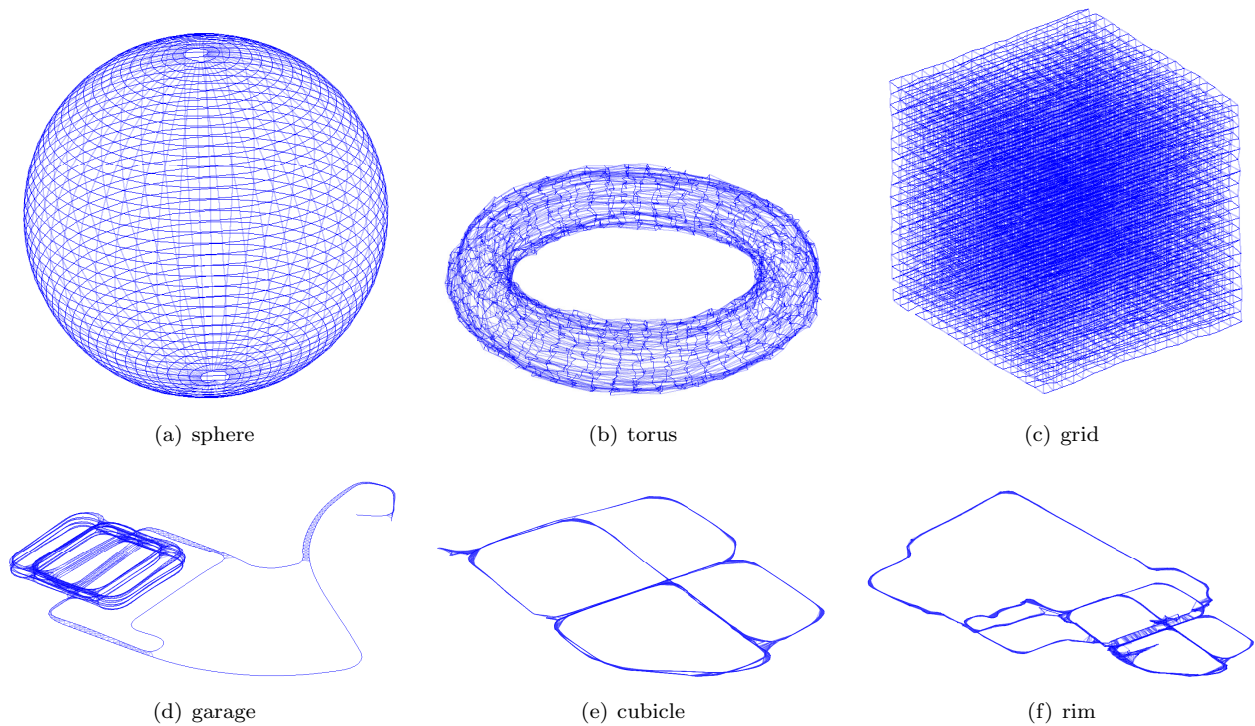


Figure 4. Globally optimal solutions for the SLAM benchmark datasets listed in Table 1.

	# Poses	# Measurements	Gauss-Newton		SE-Sync		
			Objective value	Time [s]	Objective value	Time [s]	Max. suboptimality
sphere	2500	4949	1.687×10^3	14.98	1.687×10^3	2.81	1.410×10^{-11}
torus	5000	9048	2.423×10^4	31.94	2.423×10^4	5.67	7.276×10^{-12}
grid	8000	22236	8.432×10^4	130.35	8.432×10^4	22.37	4.366×10^{-11}
garage	1661	6275	1.263×10^0	17.81	1.263×10^0	5.33	2.097×10^{-11}
cubicle	5750	16869	7.171×10^2	136.86	7.171×10^2	13.08	1.603×10^{-11}
rim	10195	29743	5.461×10^3	575.42	5.461×10^3	36.66	5.639×10^{-11}

Table 1. Results for the SLAM benchmark datasets.

solutions under real-world operating conditions, and does so considerably faster than current state-of-the-art Gauss-Newton-based alternatives.

7 Conclusion

In this paper we presented SE-Sync, a certifiably correct algorithm for synchronization over the special Euclidean group. Our algorithm is based upon the development of a semidefinite relaxation of the special Euclidean synchronization problem whose minimizer provides an *exact, globally optimal* solution so long as the magnitude of the noise corrupting the available measurements falls below a certain critical threshold, and employs a specialized, structure-exploiting Riemannian optimization method to solve large-scale instances of this semidefinite relaxation

efficiently. Experimental evaluation on a variety of simulated and real-world pose-graph SLAM datasets shows that SE-Sync is capable of recovering globally optimal solutions when the available measurements are corrupted by noise up to an order of magnitude greater than that typically encountered in robotics and computer vision applications, and does so more than an order of magnitude faster than the Gauss-Newton-based approach that forms the basis of current state-of-the-art techniques.

In addition to enabling the computation of certifiably correct solutions under *nominal* operating conditions, we believe that SE-Sync may also be extended to support provably robust and statistically efficient estimation in the case that some fraction of the available measurements \tilde{x}_{ij} in (10) are contaminated by outliers. Our basis for this belief

is the observation that Proposition 2 together with the experimental results of Section 6 imply that, under typical operating conditions, the maximum-likelihood estimation Problem 1 is equivalent to a *low-rank* convex program with a *linear* observation model and a *compact* feasible set (Problem 7); in contrast to general nonlinear estimation, this class of problems enjoys a beautiful geometric structure (Chandrasekaran et al., 2012) that has already been shown to enable remarkably robust recovery, even in the presence of gross contamination (Candès et al., 2011; Zhou et al., 2010). We intend to investigate this possibility in future research.

Finally, while the specific relaxation (32) underpinning SE-Sync was obtained by exploiting the well-known Lagrangian duality between quadratically-constrained quadratic programs and semidefinite programs (Luo et al., 2010), recent work in real algebraic geometry has revealed the remarkable fact that the much broader class of (*rational*) *polynomial optimization problems*¹⁸ also admits a hierarchy of semidefinite relaxations that is likewise frequently exact (or can be made arbitrarily sharp) when applied to real-world problem instances (Lasserre, 2001; Parrilo, 2003; Lasserre, 2006; Laurent, 2009; Bugarin et al., 2016; Nie et al., 2006; Nie and Demmel, 2008; Waki et al., 2006). Given the broad generality of this latter class of models, SE-Sync’s demonstration that it is indeed possible to solve surprisingly large (but suitably structured) semidefinite programs with temporal and computational resources typical of those available in embedded systems suggests the further possibility of designing a broad class of practically-effective certifiably correct algorithms for robust machine perception based upon structured semidefinite programming relaxations. It is our hope that this report will encourage further investigation of this exciting possibility within the machine perception community.

A The isotropic Langevin distribution

In this appendix we provide a brief overview of the isotropic Langevin distribution on $\text{SO}(d)$, with a particular emphasis on the important special cases $d = 2, 3$.

The *isotropic Langevin distribution* on $\text{SO}(d)$ with mode $M \in \text{SO}(d)$ and concentration parameter $\kappa \geq 0$, denoted $\text{Langevin}(M, \kappa)$, is the distribution determined by the following probability density function (with respect to the Haar measure on $\text{SO}(d)$ (Warner, 1983)):

$$p(X; M, \kappa) = \frac{1}{c_d(\kappa)} \exp(\kappa \text{tr}(M^T X)), \quad (51)$$

where $c_d(\kappa)$ is a normalization constant (Chiuso et al., 2008; Boumal et al., 2014b). Note that the product $M^T X = M^{-1} X \in \text{SO}(d)$ appearing in (51) is the relative rotation sending M to X .

In general, given any $Z \in \text{SO}(d)$, there exists some $U \in \text{O}(d)$ such that:

$$U^T Z U = \begin{cases} \begin{pmatrix} R(\theta_1) & & \\ & \ddots & \\ & & R(\theta_k) \end{pmatrix}, & d \bmod 2 = 0, \\ \begin{pmatrix} R(\theta_1) & & \\ & \ddots & \\ & & R(\theta_k) \\ & & & 1 \end{pmatrix}, & d \bmod 2 = 1, \end{cases} \quad (52)$$

where $k \triangleq \lfloor d/2 \rfloor$, $\theta_i \in [-\pi, \pi)$ for all $i \in [k]$, and

$$R(\theta) \triangleq \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix} \in \text{SO}(2); \quad (53)$$

this *canonical decomposition* corresponds to the fact that every rotation $Z \in \text{SO}(d)$ acts on \mathbb{R}^d as a set of elementary rotations (53) of mutually-orthogonal 2-dimensional subspaces. Since $\text{tr}(\cdot)$ is a class function, it follows from (52) and (53) that the trace appearing in the isotropic Langevin density (51) can be equivalently expressed as:

$$\text{tr}(M^T X) = d \bmod 2 + 2 \sum_{i=1}^k \cos(\theta_i), \quad (54)$$

where the θ_i ’s are the rotation angles for each of $M^T X$ ’s elementary rotations. Note, however, that while the right-hand side of (54) depends upon the *magnitudes* of these elementary rotations, it does *not* depend upon the *orientation* of their corresponding subspaces; this is the sense in which the Langevin density (51) is “isotropic”.

For the special cases $d = 2, 3$, the normalization constant $c_d(\kappa)$ appearing in (51) admits the following simple closed forms (cf. equations (4.6) and (4.7) of (Boumal et al., 2014b)):

$$c_2(\kappa) = I_0(2\kappa), \quad (55a)$$

$$c_3(\kappa) = \exp(\kappa) (I_0(2\kappa) - I_1(2\kappa)), \quad (55b)$$

where $I_n(z)$ denotes the *modified Bessel function of the first kind* (Olver et al., 2016, eq. (10.32.3)):

$$I_n(z) \triangleq \frac{1}{\pi} \int_0^\pi e^{z \cos(\theta)} \cos(n\theta) d\theta, \quad n \in \mathbb{N}. \quad (56)$$

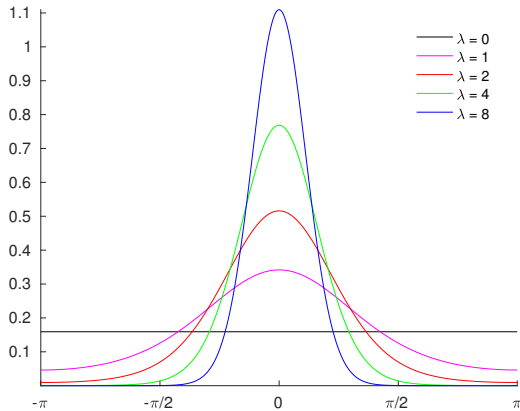


Figure 5. The von Mises distribution. This plot shows the densities (58) of several von Mises distributions on the circle (here identified with $[-\pi, \pi)$) with common mode $\mu = 0$ and varying concentrations λ . For $\lambda = 0$ the von Mises distribution reduces to the uniform distribution, and its density asymptotically converges to the Gaussian $\mathcal{N}(0, \lambda^{-1})$ as $\lambda \rightarrow \infty$.

Furthermore, in these dimensions every rotation acts on a single 2-dimensional subspace, and hence is described by a single rotation angle. Letting $\theta \triangleq \angle(M^\top X)$, it follows from (51) and (54) that the distribution over θ induced by $\text{Langevin}(M, \kappa)$ has a density satisfying:

$$p(\theta; \kappa) \propto \exp(2\kappa \cos(\theta)), \quad \theta \in [-\pi, \pi). \quad (57)$$

Now recall that the *von Mises distribution* on the circle (here identified with $[-\pi, \pi)$) with mode $\mu \in [-\pi, \pi)$ and concentration parameter $\lambda \geq 0$, denoted $\text{vM}(\mu, \lambda)$, is the distribution determined by the following probability density function:

$$p(\theta; \mu, \lambda) = \frac{\exp(\lambda \cos(\theta - \mu))}{2\pi I_0(\lambda)}. \quad (58)$$

This distribution plays a role in circular statistics analogous to that of the Gaussian distribution on Euclidean space (Fisher, 1953). For $\lambda = 0$ it reduces to the uniform distribution on S^1 , and becomes increasingly concentrated at the mode μ as $\lambda \rightarrow \infty$ (cf. Fig. 5). In fact, considering the asymptotic expansion (Olver et al., 2016, 10.40.1):

$$I_0(\lambda) \sim \frac{\exp(\lambda)}{\sqrt{2\pi\lambda}}, \quad \lambda \rightarrow \infty, \quad (59)$$

and the second-order Taylor series expansion of the cosine function:

$$\cos(\theta - \mu) \sim 1 - \frac{1}{2}(\theta - \mu)^2, \quad \theta \rightarrow \mu, \quad (60)$$

Algorithm 4 A sampler for the isotropic Langevin distribution on $\text{SO}(d)$ in dimensions 2 and 3

Input: Mode $M \in \text{SO}(d)$ with $d \in \{2, 3\}$, concentration parameter $\kappa \geq 0$.

Output: A realization of $X \sim \text{Langevin}(M, \kappa)$.

```

1: function SAMPLEISOTROPICLANGEVIN( $M, \kappa$ )
2:   Sample a rotation angle  $\theta \sim \text{vM}(0, 2\kappa)$ .‡
3:   if  $d = 2$  then
4:     Set perturbation matrix  $P \leftarrow R(\theta)$ .
5:   else  $\triangleright d = 3$ 
6:     Sample an axis of rotation  $\hat{v} \sim \mathcal{U}(S^2)$ .
7:     Set perturbation matrix  $P \leftarrow \exp(\theta[\hat{v}]_\times)$ .
8:   end if
9:   return  $MP$ 
10: end function
```

it follows from (58)–(60) that

$$p(\theta; \mu, \lambda) \sim \frac{1}{\sqrt{2\pi\lambda^{-1}}} \exp\left(-\frac{(\theta - \mu)^2}{2\lambda^{-1}}\right), \quad \lambda \rightarrow \infty, \quad (61)$$

which we recognize as the density for the Gaussian distribution $\mathcal{N}(\mu, \lambda^{-1})$.

These observations lead to a particularly convenient generative description of the isotropic Langevin distribution in dimensions 2 and 3: namely, a realization of $X \sim \text{Langevin}(M, \kappa)$ is obtained by perturbing the mode M by a rotation through an angle $\theta \sim \text{vM}(0, 2\kappa)$ about a uniformly-distributed axis (Algorithm 4). Furthermore, it follows from (58) that the standard deviation of the angle θ of the relative rotation between M and X is given by:

$$\text{SD}[\theta] = \sqrt{\int_{-\pi}^{\pi} \theta^2 \frac{\exp(2\kappa \cos(\theta))}{2\pi I_0(2\kappa)} d\theta}, \quad (62)$$

which provides a convenient and intuitive measure of the dispersion of $\text{Langevin}(M, \kappa)$. The right-hand side of (62) can be efficiently evaluated to high precision via numerical quadrature for values of κ less than 150 (corresponding to an angular standard deviation of 3.31 degrees). For $\kappa > 150$, one can alternatively use the estimate of the angular standard deviation coming from the asymptotic Gaussian approximation (61) for the von Mises distribution $\text{vM}(0, 2\kappa)$:

$$\text{SD}[\theta] \sim \frac{1}{\sqrt{2\kappa}}, \quad \kappa \rightarrow \infty. \quad (63)$$

This approximation is accurate to within 1% for $\kappa > 12.87$ (corresponding to an angular standard deviation of 11.41 degrees) and to within .1% for $\kappa > 125.3$

[‡]See e.g. (Best and Fisher, 1979)

(corresponding to an angular standard deviation of 3.62 degrees).

B Reformulating the estimation problem

B.1 Deriving Problem 2 from Problem 1

In this section we show how to derive Problem 2 from Problem 1. Using the fact that $\text{vec}(v) = v$ for any $v \in \mathbb{R}^{d \times 1}$ and the fact that $\text{vec}(XY) = (Y^T \otimes I_k) \text{vec}(X)$ for all $X \in \mathbb{R}^{m \times k}$ and $Y \in \mathbb{R}^{k \times l}$ (Horn and Johnson, 1991, Lemma 4.3.1), we can write each summand of (11) in a vectorized form as:

$$\begin{aligned} \tau_{ij} \|t_j - t_i - R_i \tilde{t}_{ij}\|_2^2 &= \tau_{ij} \|t_j - t_i - (\tilde{t}_{ij}^T \otimes I_d) \text{vec}(R_i)\|_2^2 \\ &= \|(\sqrt{\tau_{ij}} I_d) t_j - (\sqrt{\tau_{ij}} I_d) t_i - \sqrt{\tau_{ij}} (\tilde{t}_{ij}^T \otimes I_d) \text{vec}(R_i)\|_2^2 \end{aligned} \quad (64)$$

and

$$\begin{aligned} \kappa_{ij} \|R_j - R_i \tilde{R}_{ij}\|_F^2 &= \kappa_{ij} \|\text{vec}(R_j) - \text{vec}(R_i \tilde{R}_{ij})\|_2^2 \\ &= \|\sqrt{\kappa_{ij}} (I_d \otimes I_d) \text{vec}(R_j) - \sqrt{\kappa_{ij}} (\tilde{R}_{ij}^T \otimes I_d) \text{vec}(R_i)\|_2^2. \end{aligned} \quad (65)$$

Letting $t \in \mathbb{R}^{dn}$ and $R \in \mathbb{R}^{d \times dn}$ denote the concatenations of the t_i and R_i as defined in (16), (64) and (65) imply that (11) can be rewritten in a vectorized form as:

$$p_{\text{MLE}}^* = \min_{\substack{t \in \mathbb{R}^{dn} \\ R \in \text{SO}(d)^n}} \left\| B \begin{pmatrix} t \\ \text{vec}(R) \end{pmatrix} \right\|_2^2 \quad (66)$$

where the coefficient matrix $B \in \mathbb{R}^{(d+d^2)m \times (d+d^2)n}$ (with $m = |\mathcal{E}|$) has the block decomposition:

$$B \triangleq \begin{pmatrix} B_1 & B_2 \\ 0 & B_3 \end{pmatrix} \quad (67)$$

and $B_1 \in \mathbb{R}^{dm \times dn}$, $B_2 \in \mathbb{R}^{dm \times d^2 n}$, and $B_3 \in \mathbb{R}^{d^2 m \times d^2 n}$ are block-structured matrices whose block rows and columns are indexed by the elements of \mathcal{E} and \mathcal{V} , respectively, and whose (e, k) -block elements are given by:

$$(B_1)_{ek} = \begin{cases} -\sqrt{\tau_{kj}} I_d, & e = (k, j) \in \mathcal{E}, \\ \sqrt{\tau_{ik}} I_d, & e = (i, k) \in \mathcal{E}, \\ 0_{d \times d}, & \text{otherwise,} \end{cases} \quad (68a)$$

$$(B_2)_{ek} = \begin{cases} -\sqrt{\tau_{kj}} (\tilde{t}_{kj}^T \otimes I_d), & e = (k, j) \in \mathcal{E}, \\ 0_{d \times d^2}, & \text{otherwise,} \end{cases} \quad (68b)$$

$$(B_3)_{ek} = \begin{cases} -\sqrt{\kappa_{kj}} (\tilde{R}_{kj}^T \otimes I_d), & e = (k, j) \in \mathcal{E}, \\ \sqrt{\kappa_{ik}} (I_d \otimes I_d), & e = (i, k) \in \mathcal{E}, \\ 0_{d \times d}, & \text{otherwise.} \end{cases} \quad (68c)$$

We can further expand the squared ℓ_2 -norm objective in (66) to obtain:

$$p_{\text{MLE}}^* = \min_{\substack{t \in \mathbb{R}^{dn} \\ R \in \text{SO}(d)^n}} \begin{pmatrix} t \\ \text{vec}(R) \end{pmatrix}^T B^T B \begin{pmatrix} t \\ \text{vec}(R) \end{pmatrix} \quad (69)$$

with

$$B^T B = \begin{pmatrix} B_1^T B_1 & B_1^T B_2 \\ B_2^T B_1 & B_2^T B_2 + B_3^T B_3 \end{pmatrix}. \quad (70)$$

Computing each of the constituent products in (70) blockwise using (68), we obtain:

$$B_1^T B_1 = L(W^T) \otimes I_d, \quad (71a)$$

$$B_1^T B_2 = \tilde{V} \otimes I_d, \quad (71b)$$

$$B_2^T B_2 = \tilde{\Sigma} \otimes I_d, \quad (71c)$$

$$B_3^T B_3 = L(\tilde{G}^p) \otimes I_d \quad (71d)$$

where $L(W^T)$, \tilde{V} , $\tilde{\Sigma}$, and $L(\tilde{G}^p)$ are as defined in equations (12a), (14), (15), and (13), respectively. This establishes that $B^T B = M \otimes I_d$, where M is the matrix defined in (17b), and consequently that Problem 1 is equivalent to Problem 2.

B.2 Deriving Problem 3 from Problem 2

In this subsection, we show how to analytically eliminate the translations t appearing in Problem 2 to obtain the simplified form of Problem 3. We make use of the following lemma (cf. (Boyd and Vandenberghe, 2004, Appendix A.5.5) and (Gallier, 2010, Proposition 4.2)):

Lemma 4. *Given $A \in \text{Sym}(p)$ and $b \in \mathbb{R}^p$, the function*

$$f(x) = x^T A x + 2b^T x \quad (72)$$

attains a minimum if and only if $A \succeq 0$ and $(I_d - A A^\dagger)b = 0$, in which case

$$\min_{x \in \mathbb{R}^p} f(x) = -b^T A^\dagger b \quad (73)$$

and

$$\underset{x \in \mathbb{R}^p}{\text{argmin}} f(x) = \left\{ -A^\dagger b + U \begin{pmatrix} 0 \\ z \end{pmatrix} \mid z \in \mathbb{R}^{p-r} \right\}, \quad (74)$$

where $A = U \Lambda U^T$ is an eigendecomposition of A and $r = \text{rank}(A)$.

Now $L(W^\tau) \otimes I_d \succeq 0$ since $L(W^\tau)$ is a (necessarily positive semidefinite) graph Laplacian, and $I_{dn} - (L(W^\tau) \otimes I_d)(L(W^\tau) \otimes I_d)^\dagger$ is the orthogonal projection operator onto $\ker((L(W^\tau) \otimes I_d)^\top) = \ker(L(W^\tau) \otimes I_d)$ (Meyer, 2000, eq. (5.13.12)). Using the relation $\text{vec}(AYC) = (C^\top \otimes A) \text{vec}(Y)$ (Horn and Johnson, 1991, Lemma 4.3.1), we find that $y \in \ker(L(W^\tau) \otimes I_d)$ if and only if $y = \text{vec}(Y)$ for some $Y \in \mathbb{R}^{d \times n}$ satisfying $YL(W^\tau) = 0$, or equivalently $L(W^\tau)Y^\top = 0$. Since G is connected, then $\ker(L(W^\tau)) = \text{span}\{\mathbb{1}_n\}$, and therefore we must have $Y^\top = \mathbb{1}_n c^\top \in \mathbb{R}^{n \times d}$ for some $c \in \mathbb{R}^d$. Altogether, this establishes that:

$$\ker(L(W^\tau) \otimes I_d) = \{\text{vec}(c\mathbb{1}_n^\top) \mid c \in \mathbb{R}^d\}. \quad (75)$$

Now let $b = (\tilde{V} \otimes I_d) \text{vec}(R)$; we claim that $b \perp y$ for all $y = \text{vec}(Y) \in \ker(L(W^\tau) \otimes I_d)$. To see this, observe that $b = \text{vec}(R\tilde{V}^\top)$ and therefore $\langle b, y \rangle_2 = \langle R\tilde{V}^\top, Y \rangle_F = \text{tr}(Y\tilde{V}R^\top)$. However, $\mathbb{1}_n^\top \tilde{V} = 0$ since the sum down each column of \tilde{V} is identically 0 by (14), and therefore $Y\tilde{V} = 0$ by (75). This establishes that $b \perp \ker(L(W^\tau) \otimes I_d)$ for any value of R .

Consequently, if we fix $R \in \text{SO}(d)^n$ and consider performing the optimization in (18) over the decision variables t only, we can apply Lemma 4 to compute the optimal value of the objective and a minimizing value t^* of t as functions of R . This in turn enables us to analytically eliminate t from (18), producing the equivalent optimization problem (76), with a corresponding optimal translational estimate t^* given by:

$$t^* = - \left(L(W^\tau)^\dagger \tilde{V} \otimes I_d \right) \text{vec}(R^*). \quad (77)$$

Rewriting (76) and (77) in a more compact matrixized form gives (19) and (20), respectively.

B.3 Simplifying the translational data matrix

In this subsection we derive the simplified form of the translational data matrix \tilde{Q}^τ given in (23c) from the one originally presented in (19b). To begin, recall from Appendix B.1 that

$$\begin{aligned} \tilde{Q}^\tau \otimes I_d &= \left(\tilde{\Sigma} - \tilde{V}^\top L(W^\tau)^\dagger \tilde{V} \right) \otimes I_d \\ &= B_2^\top B_2 - B_2^\top B_1 (B_1^\top B_1)^\dagger B_1^\top B_2 \\ &= B_2^\top \left(I_{dm} - B_1 (B_1^\top B_1)^\dagger B_1^\top \right) B_2, \end{aligned} \quad (78)$$

where B_1 and B_2 are the matrices defined in (68a) and (68b), respectively. Using Ω and \tilde{T} as defined in (21) and (22), respectively, we may write B_1 and B_2 alternatively as:

$$B_1 = \Omega^{\frac{1}{2}} A^\top \otimes I_d, \quad B_2 = \Omega^{\frac{1}{2}} \tilde{T} \otimes I_d, \quad (79)$$

where $A \triangleq A(\vec{G})$ is the incidence matrix of \vec{G} . Substituting (79) into (78), we derive as shown in (80), or equivalently:

$$\tilde{Q}^\tau = \tilde{T}^\top \Omega^{\frac{1}{2}} \underbrace{\left(I_m - \Omega^{\frac{1}{2}} A^\top (A \Omega A^\top)^\dagger A \Omega^{\frac{1}{2}} \right)}_{\Pi} \Omega^{\frac{1}{2}} \tilde{T}. \quad (81)$$

Now let us develop the term Π appearing in (81):

$$\begin{aligned} \Pi &= I_m - \Omega^{\frac{1}{2}} A^\top (A \Omega A^\top)^\dagger A \Omega^{\frac{1}{2}} \\ &= I_m - \left(A \Omega^{\frac{1}{2}} \right)^\top \left(\left(A \Omega^{\frac{1}{2}} \right) \left(A \Omega^{\frac{1}{2}} \right)^\top \right)^\dagger \left(A \Omega^{\frac{1}{2}} \right) \\ &= I_m - \left(A \Omega^{\frac{1}{2}} \right)^\dagger \left(A \Omega^{\frac{1}{2}} \right) \end{aligned} \quad (82)$$

where we have used the fact that $X^\top (X X^\top)^\dagger = X^\dagger$ for any matrix X in passing from line 2 to line 3 above. We may now recognize the final line of (82) as the matrix of the orthogonal projection operator $\pi: \mathbb{R}^m \rightarrow \ker(A \Omega^{\frac{1}{2}})$ onto the kernel of the weighted incidence matrix $A \Omega^{\frac{1}{2}}$ (Meyer, 2000, eq. (5.13.12)). Equation (23c) thus follows from (81) and (82).

Finally, although Π is generically dense, we now show that it admits a computationally convenient decomposition in terms of sparse matrices and their inverses. By the Fundamental Theorem of Linear Algebra, $\ker(A \Omega^{\frac{1}{2}})^\perp = \text{image}(\Omega^{\frac{1}{2}} A^\top)$, and therefore every vector $v \in \mathbb{R}^m$ admits the orthogonal decomposition:

$$v = \pi(v) + c, \quad \pi(v) \in \ker(A \Omega^{\frac{1}{2}}), \quad c \in \text{image}(\Omega^{\frac{1}{2}} A^\top). \quad (83)$$

Now $\text{rank}(A) = n - 1$ since A is the incidence matrix of the weakly-connected directed graph \vec{G} ; it follows that $\text{image}(\Omega^{\frac{1}{2}} A^\top) = \text{image}(\Omega^{\frac{1}{2}} \underline{A}^\top)$, where \underline{A} is the reduced incidence matrix of \vec{G} formed by removing the final row of A . Furthermore, since c is the complement of $\pi(v)$ in the orthogonal decomposition (83), it is itself the orthogonal projection of v onto $\text{image}(\Omega^{\frac{1}{2}} A^\top) = \text{image}(\Omega^{\frac{1}{2}} \underline{A}^\top)$, and is therefore the value of the product realizing the minimum norm in:

$$\min_{w \in \mathbb{R}^{n-1}} \|v - \Omega^{\frac{1}{2}} \underline{A}^\top w\|_2. \quad (84)$$

Consequently, it follows from (83) and (84) that:

$$\begin{aligned} \pi(v) &= v - \Omega^{\frac{1}{2}} \underline{A}^\top w^*, \\ w^* &= \underset{w \in \mathbb{R}^{n-1}}{\text{argmin}} \|v - \Omega^{\frac{1}{2}} \underline{A}^\top w\|_2. \end{aligned} \quad (85)$$

Since \underline{A} is full-rank, we can solve for the minimizer w^* in (85) via the normal equations, obtaining:

$$w^* = (\underline{A} \Omega \underline{A}^\top)^{-1} \underline{A} \Omega^{\frac{1}{2}} v = L^{-\top} L^{-1} \underline{A} \Omega^{\frac{1}{2}} v, \quad (86)$$

$$p_{\text{MLE}}^* = \min_{R \in \text{SO}(d)} \text{vec}(R)^\top \left(\left(L(\tilde{G}^\rho) + \tilde{\Sigma} - \tilde{V}^\top L(W^\tau)^\dagger \tilde{V} \right) \otimes I_d \right) \text{vec}(R), \quad (76)$$

$$\begin{aligned} \tilde{Q}^\tau \otimes I_d &= B_2^\top (I_{dm} - B_1(B_1^\top B_1)^\dagger B_1^\top) B_2 \\ &= B_2^\top \left(I_{dm} - \left(\Omega^{\frac{1}{2}} A^\top \otimes I_d \right) \left(\left(\Omega^{\frac{1}{2}} A^\top \otimes I_d \right)^\top \left(\Omega^{\frac{1}{2}} A^\top \otimes I_d \right) \right)^\dagger \left(\Omega^{\frac{1}{2}} A^\top \otimes I_d \right)^\top \right) B_2 \\ &= B_2^\top \left(I_{dm} - \left(\Omega^{\frac{1}{2}} A^\top \otimes I_d \right) \left((A\Omega A^\top)^\dagger \otimes I_d \right) \left(\Omega^{\frac{1}{2}} A^\top \otimes I_d \right)^\top \right) B_2 \\ &= B_2^\top \left(I_{dm} - \left(\Omega^{\frac{1}{2}} A^\top (A\Omega A^\top)^\dagger A\Omega^{\frac{1}{2}} \right) \otimes I_d \right) B_2 \\ &= B_2^\top \left[\left(I_m - \Omega^{\frac{1}{2}} A^\top (A\Omega A^\top)^\dagger A\Omega^{\frac{1}{2}} \right) \otimes I_d \right] B_2 \\ &= \left(\tilde{T}^\top \Omega^{\frac{1}{2}} \left(I_m - \Omega^{\frac{1}{2}} A^\top (A\Omega A^\top)^\dagger A\Omega^{\frac{1}{2}} \right) \Omega^{\frac{1}{2}} \tilde{T} \right) \otimes I_d, \end{aligned} \quad (80)$$

where $A\Omega^{\frac{1}{2}} = LQ_1$ is a thin LQ decomposition of the weighted reduced incidence matrix $A\Omega^{\frac{1}{2}}$. Substituting (86) into (85), we obtain:

$$\begin{aligned} \pi(v) &= v - \Omega^{\frac{1}{2}} A^\top L^{-\top} L^{-1} A\Omega^{\frac{1}{2}} v \\ &= \left(I_m - \Omega^{\frac{1}{2}} A^\top L^{-\top} L^{-1} A\Omega^{\frac{1}{2}} \right) v. \end{aligned} \quad (87)$$

Since v was arbitrary, we conclude that the matrix in parentheses on the right-hand side of (87) is Π , which gives (38).

B.4 Deriving Problem 7 from Problem 6

This derivation is a straightforward application of the duality theory for semidefinite programs (Vandenberghe and Boyd, 1996). Letting $\lambda_{iuv} \in \mathbb{R}$ denote the (u, v) -element of the i th diagonal block of Λ for $i = 1, \dots, n$ and $1 \leq u \leq v \leq d$, we can rewrite Problem 6 in the primal standard form of (Vandenberghe and Boyd, 1996, eq. (1)) as:

$$\begin{aligned} -p_{\text{SDP}}^* &= \min_{\lambda \in \mathbb{R}^{\frac{d(d+1)}{2}n}} c^\top \lambda \\ \text{s.t. } &F(\lambda) \succeq 0, \end{aligned} \quad (88)$$

with the problem data in (88) given explicitly by:

$$\begin{aligned} F(\lambda) &\triangleq F_0 + \sum_{i=1}^n \sum_{1 \leq u \leq v \leq d} \lambda_{iuv} F_{iuv}, \\ F_0 &= Q, \\ F_{iuv} &= \begin{cases} -\text{Diag}(e_i) \otimes E_{uu}, & u = v, \\ -\text{Diag}(e_i) \otimes (E_{uv} + E_{vu}), & u \neq v, \end{cases} \\ c_{iuv} &= \begin{cases} -1, & u = v, \\ 0, & u \neq v. \end{cases} \end{aligned} \quad (89)$$

The standard dual problem for (88) is then (cf. equations (1) and (27) of (Vandenberghe and Boyd, 1996)):

$$\begin{aligned} -d_{\text{SDP}}^* &= \max_{Z \in \text{Sym}(dn)} -\text{tr}(F_0 Z) \\ \text{s.t. } &\text{tr}(F_{iuv} Z) = c_{iuv}, \\ &Z \succeq 0, \end{aligned} \quad (90)$$

for all $i \in \{1, \dots, n\}$ and $1 \leq u \leq v \leq d$. Comparing (90) with (89) reveals that the equality constraints are satisfied precisely when Z 's $(d \times d)$ -block-diagonal is comprised of identity matrices, which gives the form of Problem 7. Furthermore, since $\tilde{Q} - \Lambda \succ 0$ for any $\Lambda = -sI_{dn}$ with $s > \|\tilde{Q}\|_2$ is strictly feasible for Problem 6 (hence also (88)) and $I_{dn} \succ 0$ is strictly feasible for (90), Theorem 3.1 of (Vandenberghe and Boyd, 1996) implies that the optimal sets of (88) and (90) are nonempty, and that strong duality holds between them (so that the optimal value of Problem 7 is p_{SDP}^* , as claimed).

C Proof of Proposition 2

In this appendix we prove Proposition 2, following the general roadmap of the proof of a similar result for the special case of *angular synchronization* due to (Bandeira et al., 2016). At a high level, our approach is based upon exploiting the Lagrangian duality between Problems 5 and 7 to identify a matrix C (constructed from an optimal solution R^* of Problem 5) with the property that $C \succeq 0$ and $\text{rank}(C) = dn - d$ imply that $Z^* = R^{*\top} R^*$ is the unique optimal solution of Problem 7; we then show that these conditions can be assured by controlling the magnitude of the deviation $\Delta Q \triangleq \tilde{Q} - Q$ of the observed data matrix \tilde{Q} from its exact latent

value Q . Specifically, our proof proceeds according to the following chain of reasoning:

1. We begin by deriving the first-order necessary optimality conditions for the extrinsic formulation of Problem 5; these take the form $(\tilde{Q} - \Lambda^*)R^{*\top} = 0$, where $R^* \in O(d)^n$ is a minimizer of Problem 5 and $\Lambda^* = \text{SymBlockDiag}_d(\tilde{Q}R^{*\top}R^*)$ is a symmetric block-diagonal matrix of Lagrange multipliers corresponding to the orthogonality constraints in (25).
2. Exploiting the Lagrangian duality between Problems 5 and 7 together with nondegeneracy results for semidefinite programs (Alizadeh et al., 1997), we identify a matrix $C \triangleq \tilde{Q} - \Lambda^*$ with the property that $C \succeq 0$ and $\text{rank}(C) = dn - d$ imply that $Z^* = R^{*\top}R^*$ is the unique optimal solution of Problem 7 (Theorem 7).
3. We next observe that for the idealized case in which the measurements \tilde{x}_{ij} of the relative transforms are *noiseless*, the true latent rotations R comprise a minimizer of Problem 5 with corresponding certificate $C = Q \succeq L(\underline{Q}^\rho)$. We then show (by means of similarity) that the spectrum of $L(\underline{Q}^\rho)$ consists of d copies of the spectrum of the rotational weight graph W^ρ ; in particular, $L(\underline{Q}^\rho)$ has d eigenvalues equal to 0, and the remaining $d(n-1)$ are lower-bounded by the algebraic connectivity $\lambda_2(L(W^\rho)) > 0$ of W^ρ . Consequently $C \succeq 0$, and $\text{rank}(C) = dn - d$ since $R \in \ker \tilde{Q} = \ker C$. It follows that Problem 7 is *always* exact in the absence of measurement noise (Theorem 10).
4. In the presence of noise, the minimizer $R^* \in O(d)^n$ of Problem 5 will generally not coincide with the true latent rotations $R \in SO(d)^n$. Nevertheless, we can still derive an upper bound for the error in the estimate R^* in terms of the magnitude of the deviation $\Delta Q \triangleq \tilde{Q} - Q$ of the data matrix \tilde{Q} from the true latent value Q (Theorem 12).
5. The first-order necessary optimality condition for Problem 5 (point 1 above) can alternatively be read as $CR^{*\top} = 0$, which shows that d eigenvalues of C are always 0. Since in general the eigenvalues of a matrix X are continuous functions of X , it follows from points 3 and 4 above and the definition of C that the other $d(n-1)$ eigenvalues of C can be controlled into remaining nonnegative by controlling the norm of ΔQ . In light of point 2, this establishes the existence of a constant $\beta_1 > 0$ such that, if $\|\tilde{Q} - Q\|_2 < \beta_1$, $Z^* = R^{*\top}R^*$ is the unique solution of Problem 7.

6. Finally, we observe that since $O(d)$ is the disjoint union of two components separated by a distance $\sqrt{2}$ in the Frobenius norm, and the true latent rotations R_i all lie in the $+1$ component, it follows from point 4 that there exists a constant $\beta_2 > 0$ such that, if $\|\tilde{Q} - Q\|_2 < \beta_2$, a minimizer $R^* \in O(d)^n$ of Problem 5 must in fact lie in $SO(d)^n$, and is therefore also a minimizer of Problem 4.
7. Taking $\beta \triangleq \min\{\beta_1, \beta_2\}$, Proposition 2 follows from points 5 and 6 and Theorem 1.

The remainder of this appendix is devoted to rigorously establishing each of claims 1–7 above.

C.1 Gauge symmetry and invariant metrics for Problems 4 and 5

A critical element of the proof of Proposition 2 is the derivation of an upper bound for the estimation error of a minimizer R^* of Problem 5 as a function of ΔQ (point 4). However, we have observed previously that Problems 1–5 always admit *multiple* (in fact, infinitely many) solutions for dimensions $d \geq 2$ due to gauge symmetry.¹⁹ In consequence, it may not be immediately clear how we should quantify the estimation error of a *particular* point estimate R^* obtained as a minimizer of Problem 4 or Problem 5, since R^* is an arbitrary representative of an infinite set of *distinct* but *equivalent* minimizers that are related by gauge transforms. To address this complication, in this subsection we study the gauge symmetries of Problems 4 and 5, and then develop a pair of *gauge-invariant* metrics suitable for quantifying estimation error in these problems in a consistent, “symmetry aware” manner.

Recall from Section 3.2 (cf. footnote 6 on page 29) that solutions of Problem 1 are determined only up to a global gauge symmetry (corresponding to the diagonal left-action of $SE(d)$ on $SE(d)^n$). Similarly, it is straightforward to verify that if R^* is any minimizer of (23) (resp. (24)), then $G \bullet R^*$ also minimizes (23) (resp. (24)) for any choice of $G \in SO(d)$ (resp. $G \in O(d)$), where \bullet is the diagonal left-action of $O(d)$ on $O(d)^n$:

$$G \bullet R \triangleq (GR_1, \dots, GR_n). \quad (91)$$

It follows that the sets of minimizers of Problems 4 and 5 are partitioned into *orbits* of the form:

$$\mathcal{S}(R) \triangleq \{G \bullet R \mid G \in SO(d)\} \subset SO(d)^n, \quad (92a)$$

$$\mathcal{O}(R) \triangleq \{G \bullet R \mid G \in O(d)\} \subset O(d)^n, \quad (92b)$$

each of which comprise a set of point estimates that are *equivalent* to R for the estimation problems (23) and (24), respectively. Consequently, when quantifying the dissimilarity between a pair of feasible points X, Y for Problem 4 or 5, we are interested in measuring the distance between the *orbits* determined by these two points, *not* the distance between the *specific representatives* X and Y themselves (which may in fact represent *equivalent* solutions, but differ in their assignments of coordinates by a global gauge symmetry). We therefore introduce the following *orbit distances*:

$$d_S(X, Y) \triangleq \min_{G \in \text{SO}(d)} \|X - G \bullet Y\|_F, \quad X, Y \in \text{SO}(d)^n, \quad (93a)$$

$$d_O(X, Y) \triangleq \min_{G \in \text{O}(d)} \|X - G \bullet Y\|_F, \quad X, Y \in \text{O}(d)^n; \quad (93b)$$

these functions report the Frobenius norm distance between the two closest representatives of the orbits (92) in $\text{SO}(d)^n$ and $\text{O}(d)^n$ determined by X and Y , respectively. Using the orthogonal invariance of the Frobenius norm, it is straightforward to verify that these functions satisfy:

$$d_S(X, Y) = d_S(G_1 \bullet X, G_2 \bullet Y) \quad \forall X, Y \in \text{SO}(d)^n, G_1, G_2 \in \text{SO}(d), \quad (94a)$$

$$d_O(X, Y) = d_O(G_1 \bullet X, G_2 \bullet Y) \quad \forall X, Y \in \text{O}(d)^n, G_1, G_2 \in \text{O}(d), \quad (94b)$$

i.e., they define notions of dissimilarity between feasible points of Problems 4 and 5, respectively, that are *invariant* with respect to the action of the gauge symmetries for these problems, and so provide a consistent, gauge-invariant means of quantifying the estimation error.²⁰

The following result enables us to compute these distances easily in practice:

Theorem 5. Computing the orbit distance. *Given $X, Y \in \text{O}(d)^n$, let*

$$XY^\top = U\Sigma V^\top \quad (95)$$

be a singular value decomposition of XY^\top with $\Sigma = \text{Diag}(\sigma_1, \dots, \sigma_d)$ and $\sigma_1 \geq \dots \geq \sigma_d \geq 0$. Then the orbit distance $d_O(X, Y)$ between $\mathcal{O}(X)$ and $\mathcal{O}(Y)$ in $\text{O}(d)^n$ is:

$$d_O(X, Y) = \sqrt{2dn - 2\|XY^\top\|_*}, \quad (96)$$

and a minimizer $G_O^ \in \text{O}(d)$ realizing this optimal value in (93b) is given by:*

$$G_O^* = UV^\top. \quad (97)$$

If additionally $X, Y \in \text{SO}(d)$, then the orbit distance $d_S(X, Y)$ between $\mathcal{S}(X)$ and $\mathcal{S}(Y)$ in $\text{SO}(d)^n$ is given by:

$$d_S(X, Y) = \sqrt{2dn - 2\text{tr}(\Xi\Sigma)}, \quad (98)$$

where Ξ is the matrix:

$$\Xi = \text{Diag}(1, \dots, 1, \det(UV^\top)) \in \mathbb{R}^{d \times d}, \quad (99)$$

and a minimizer $G_S^ \in \text{SO}(d)$ realizing this optimal value in (93a) is given by:*

$$G_S^* = U\Xi V^\top. \quad (100)$$

Proof. Observe that

$$\begin{aligned} \|X - G \bullet Y\|_F^2 &= \sum_{i=1}^n \|X_i - GY_i\|_F^2 \\ &= 2dn - 2 \sum_{i=1}^n \langle X_i, GY_i \rangle_F \\ &= 2dn - 2 \left\langle G, \sum_{i=1}^n X_i Y_i^\top \right\rangle_F \\ &= 2dn - 2 \langle G, XY^\top \rangle_F. \end{aligned} \quad (101)$$

Consequently, a minimizer $G_O^* \in \text{O}(d)$ attaining the optimal value in (93b) is determined by:

$$G_O^* \in \underset{G \in \text{O}(d)}{\text{argmax}} \langle G, XY^\top \rangle_F. \quad (102)$$

But now we may recognize (102) as an instance of the orthogonal Procrustes problem, with maximizer G_O^* given by (97) (Hanson and Norris, 1981; Umeyama, 1991). Substituting (95) and (97) into (101) and simplifying the resulting expression using the orthogonal invariance of the Frobenius inner product then shows that the orbit distance $d_O(X, Y)$ is:

$$d_O(X, Y) = \sqrt{2dn - 2\text{tr}(\Sigma)} = \sqrt{2dn - 2\|XY^\top\|_*}, \quad (103)$$

which is (96). If additionally $X, Y \in \text{SO}(d)$, then (101) implies that a minimizer $G_S^* \in \text{SO}(d)$ attaining the optimal value in (93a) is determined by:

$$G_S^* \in \underset{G \in \text{SO}(d)}{\text{argmax}} \langle G, XY^\top \rangle_F. \quad (104)$$

Equation (104) is an instance of the special orthogonal Procrustes problem, with a maximizer $G_S^* \in \text{SO}(d)$ given by (100) (Hanson and Norris, 1981; Umeyama, 1991). Substituting (95) and (100) into (101) and once again simplifying the resulting expression using the orthogonal invariance of the Frobenius inner product then shows that the orbit distance $d_S(X, Y)$ is given by (98).

C.2 A sufficient condition for exact recovery in Problem 5

In this subsection we address points 1 and 2 in our roadmap, deriving sufficient conditions to ensure the recovery of a minimizer $R^* \in O(d)^n$ of Problem 5 by means of solving the dual semidefinite relaxation Problem 7. Our approach is based upon exploiting the Lagrangian duality between Problem 5 and Problems 6 and 7 to construct a matrix C whose positive semidefiniteness serves as a *certificate of optimality* for $Z^* = R^{*\top} R^*$ as a solution of Problem 7.

We begin by deriving the first-order necessary optimality conditions for (25):

Lemma 6. First-order necessary optimality conditions for Problem 5. *If $R^* \in O(d)^n$ is a minimizer of Problem 5, then there exists a matrix $\Lambda^* \in \text{SBD}(d, n)$ such that*

$$(\tilde{Q} - \Lambda^*)R^{*\top} = 0. \quad (105)$$

Furthermore, Λ^* can be computed in closed form according to:

$$\Lambda^* = \text{SymBlockDiag}_d(\tilde{Q}R^{*\top}R^*). \quad (106)$$

Proof. If we consider (24) as an *unconstrained* minimization of the objective $F(R) \triangleq \text{tr}(\tilde{Q}R^\top R)$ over the Riemannian manifold $O(d)^n$, then the first-order necessary optimality condition is simply that the Riemannian gradient $\text{grad} F$ must vanish at the minimizer R^* :

$$\text{grad} F(R^*) = 0. \quad (107)$$

Furthermore, if we consider $O(d)^n$ as an embedded submanifold of $\mathbb{R}^{d \times dn}$, then this embedding induces a simple relation between the Riemannian gradient $\text{grad} F$ of F viewed as a function restricted to $O(d)^n \subset \mathbb{R}^{d \times dn}$ and ∇F , the gradient of F considered as a function on the ambient Euclidean space $\mathbb{R}^{d \times dn}$. Specifically, we have:

$$\text{grad} F(R) = \text{Proj}_R \nabla F(R), \quad (108)$$

where $\text{Proj}_R: T_R(\mathbb{R}^{d \times dn}) \rightarrow T_R(O(d)^n)$ is the orthogonal projection operator onto the tangent space of $O(d)^n$ at R (Absil et al., 2008, eq. (3.37)); this latter is given explicitly by:

$$\begin{aligned} \text{Proj}_R: T_R(\mathbb{R}^{d \times dn}) &\rightarrow T_R(O(d)^n) \\ \text{Proj}_R(X) &= X - R \text{SymBlockDiag}_d(R^\top X). \end{aligned} \quad (109)$$

Straightforward differentiation shows that the Euclidean gradient is $\nabla F(R) = 2R\tilde{Q}$, and consequently (107)–(109) imply that:

$$\begin{aligned} 0 &= \text{Proj}_{R^*} \nabla F(R^*) \\ &= 2R^*\tilde{Q} - 2R^* \text{SymBlockDiag}_d(R^{*\top}R^*\tilde{Q}). \end{aligned} \quad (110)$$

Dividing both sides of (110) by 2 and taking the transpose produces (105), using the definition of Λ^* given in (106).

Despite its simplicity, it turns out that Lemma 6 is actually already enough to enable the derivation of sufficient conditions to ensure the exactness of the semidefinite relaxation Problem 7 with respect to Problem 5. Comparing (105) with the *extrinsic* formulation (25) of Problem 5, we may recognize Λ^* as nothing more than a matrix of Lagrange multipliers corresponding to the orthogonality constraints $R_i^\top R_i = I_d$. Consequently, in the case that exactness holds between Problems 5 and 7 (i.e., that $Z^* = R^{*\top}R^*$ is a minimizer of Problem 7), strong duality obtains *a fortiori* between Problems 5 and 6, and therefore Λ^* also comprises an optimal solution for the Lagrangian relaxation (31) (cf. e.g. (Boyd and Vandenberghe, 2004, Sec. 5.5.3)). It follows that $\tilde{Q} - \Lambda^* \succeq 0$ (since Λ^* is *a fortiori* feasible for Problem 6 if it is optimal), and $(\tilde{Q} - \Lambda^*)Z^* = 0$ (from the definition of Z^* and (105)). But now observe that these last two conditions are precisely the first-order necessary and sufficient conditions for Z^* to be an optimal solution of Problem 7 (cf. equation (33) of (Vandenberghe and Boyd, 1996)); furthermore, they provide a *closed-form expression* for a KKT certificate for Z^* (namely, $\tilde{Q} - \Lambda^*$) in terms of a minimizer R^* of Problem 5 using (106). The utility of this expression is that, although it was originally derived under the *assumption* of exactness, it can also be exploited to derive a *sufficient condition* for same, as shown in the next theorem.

Theorem 7. A sufficient condition for exact recovery in Problem 5. *Let $R^* \in O(d)^n$ be a minimizer of Problem 5 with corresponding Lagrange multiplier matrix $\Lambda^* \in \text{SBD}(d, n)$ as in Lemma 6, and define*

$$C \triangleq \tilde{Q} - \Lambda^*. \quad (111)$$

If $C \succeq 0$, then $Z^ = R^{*\top}R^*$ is a minimizer of Problem 7. If additionally $\text{rank}(C) = dn - d$, then Z^* is the unique minimizer of Problem 7.*

Proof. Since $C = \tilde{Q} - \Lambda^* \succeq 0$ by hypothesis, and equation (105) of Lemma 6 implies that

$$(\tilde{Q} - \Lambda^*)Z^* = (\tilde{Q} - \Lambda^*)(R^{*\top}R^*) = 0, \quad (112)$$

Λ^* and Z^* satisfy the necessary and sufficient conditions characterizing primal-dual pairs of optimal solutions for the strictly feasible primal-dual pair of semidefinite programs (31) and (32) (cf. Appendix B.4 and Theorem 3.1 in (Vandenberghe and Boyd, 1996)). In other words, $C \succeq 0$ *certifies* the optimality of Z^* as a solution of Problem 7.

Now suppose further that $\text{rank}(C) = dn - d$; we establish that Z^* is the *unique* solution of Problem 7 using nondegeneracy results from (Alizadeh et al., 1997). Specifically, we observe that the equivalent formulations (90) and (88) of Problems 7 and 6 derived in Appendix B.4 match the forms of the primal and dual semidefinite programs given in equations (2) and (3) of (Alizadeh et al., 1997), respectively. Consequently, Theorem 10 of (Alizadeh et al., 1997) guarantees that Problem 7 has a unique solution provided that we can exhibit a dual nondegenerate solution of Problem 6. Since we have already identified Λ^* as a solution of Problem 6 (via (112)), it suffices to show that Λ^* is dual nondegenerate. To that end, let

$$\tilde{Q} - \Lambda^* = (U \quad V) \text{Diag}(0, 0, 0, \sigma_{d+1}, \dots, \sigma_{dn}) (U \quad V)^T \quad (113)$$

be an eigendecomposition of $\tilde{Q} - \Lambda^*$ as in equation (16) of (Alizadeh et al., 1997) (with $\sigma_k > 0$ for $k \geq d+1$, $U \in \mathbb{R}^{dn \times d}$, and $V \in \mathbb{R}^{dn \times (dn-d)}$). Theorem 9 of (Alizadeh et al., 1997) states that Λ^* is dual nondegenerate if and only if

$$\{U^T \Xi U \mid \Xi \in \text{SBD}(d, n)\} = \text{Sym}(d). \quad (114)$$

Now the matrix U appearing in (113) can be characterized as a matrix whose columns form an orthonormal basis for the d -dimensional null space of $\tilde{Q} - \Lambda^*$. But equation (105) shows that the columns of $R^{*\top}$ span this same subspace, and are pairwise orthogonal since $R^{*\top}$ is composed of orthogonal blocks. Consequently, without loss of generality we may take $U = \frac{1}{\sqrt{n}} R^{*\top}$ in (113). Now we can write the left-hand side of (114) more explicitly as:

$$\begin{aligned} & \{U^T \Xi U \mid \Xi \in \text{SBD}(d, n)\} \\ &= \left\{ \frac{1}{n} \sum_{i=1}^n R_i^* \Xi_i R_i^{*\top} \mid \Xi \in \text{SBD}(d, n) \right\}, \quad (115) \end{aligned}$$

and it is immediate from (115) that given any $S \in \text{Sym}(d)$, we have $S = U^T \Xi U$ for $\Xi = \text{Diag}(nR_1^{*\top} S R_1^*, 0, \dots, 0)$. This shows that (114) holds, so Λ^* is a dual nondegenerate solution of Problem 6, and we conclude that Z^* is indeed the *unique* minimizer of Problem 7, as claimed.

In short, Theorem 7 enables us to reduce the question of Problem 7's exactness to the problem of verifying the positive semidefiniteness of a certain matrix C that can be constructed from an optimal solution of Problem 5. The remainder of this appendix is devoted to establishing conditions that are sufficient to guarantee that this latter (much simpler) criterion is satisfied.

C.3 The noiseless case

As our first application of Theorem 7, in this subsection we prove that the semidefinite relaxation Problem 7 is always exact in the (highly idealized) case that the measurements \tilde{x}_{ij} in (10) are noiseless. In addition to providing a baseline sanity check for the feasibility of our overall strategy, our analysis of this idealized case will also turn out to admit a straightforward generalization that suffices to prove Proposition 2.

To that end, let $L(\mathcal{G}^\rho)$ and Q^τ denote the rotational and translational data matrices of the form appearing in Problem 4 constructed using the true (latent) relative transforms \underline{x}_{ij} appearing in (10). The following pair of lemmas characterize several important properties of these matrices:

Lemma 8. Exact rotational connection Laplacians. Let $L(\mathcal{G}^\rho)$ be the rotational connection Laplacian constructed using the true (latent) relative rotations $\underline{R}_{ij} \triangleq \underline{R}_i^{-1} \underline{R}_j$ in (10), W^ρ the corresponding rotational weight graph, and $R \in \text{SO}(d)^n$ the matrix of true (latent) rotational states. Then the following hold:

- (i) $L(W^\rho) \otimes I_d = SL(\mathcal{G}^\rho) S^{-1}$ for $S = \text{Diag}(\underline{R}_1, \dots, \underline{R}_n) \in \mathbb{R}^{dn \times dn}$;
- (ii) $\lambda_{d+1}(L(\mathcal{G}^\rho)) = \lambda_2(L(W^\rho))$;
- (iii) $\ker(L(\mathcal{G}^\rho)) = \{R^\top v \mid v \in \mathbb{R}^d\}$.

Proof. A direct computation using the definitions of S in claim (i) and of $L(\mathcal{G}^\rho)$ in (13) shows that the (i, j) -block of the product $SL(\mathcal{G}^\rho) S^{-1}$ is given by:

$$(SL(\mathcal{G}^\rho) S^{-1})_{ij} = \begin{cases} d_i^\rho I_d, & i = j, \\ -\kappa_{ij} I_d, & \{i, j\} \in \mathcal{E}, \\ 0_{d \times d}, & \{i, j\} \notin \mathcal{E}, \end{cases} \quad (116)$$

which we recognize as the $(d \times d)$ -block description of $L(W^\rho) \otimes I_d$; this proves claim (i). For claim (ii), we observe that $L(W^\rho) \otimes I_d$ and $L(\mathcal{G}^\rho)$ have the same spectrum (since claim (i) shows that they are similar), and the spectrum of $L(W^\rho) \otimes I_d$ consists of d copies of the spectrum of $L(W^\rho)$ (this follows from the fact that the spectra of $A \in \mathbb{R}^{d_1 \times d_1}$, $B \in \mathbb{R}^{d_2 \times d_2}$ and $A \otimes B$ are related by $\lambda(A \otimes B) = \{\lambda_i(A) \lambda_j(B) \mid i \in [d_1], j \in [d_2]\}$, see e.g. (Horn and Johnson, 1991, Theorem 4.2.12)). For claim (iii), another direct computation using definition (13) shows that $L(\mathcal{G}^\rho) R^\top = 0$, and therefore that $\text{image}(R^\top) = \{R^\top v \mid v \in \mathbb{R}^d\} \subseteq \ker(L(\mathcal{G}^\rho))$; furthermore, $\dim(\text{image}(R^\top)) = d$ since $\text{rank}(R^\top) = d$ (as it has d orthonormal columns). On the other hand, claim (ii) shows that $\lambda_{d+1}(L(\mathcal{G}^\rho)) = \lambda_2(L(W^\rho)) > 0$ since G is connected, and therefore $\dim(\ker(L(\mathcal{G}^\rho))) \leq d$; consequently, $\text{image}(R^\top)$ is all of $\ker(L(\mathcal{G}^\rho))$.

Lemma 9. Orthogonal projections of exact measurements. Let $\underline{T} \in \mathbb{R}^{m \times dn}$ denote the data matrix of the form (22) constructed using the true (latent) values of the translational measurements \underline{t}_{ij} in (10) and $\underline{R} \in \text{SO}(d)^n$ the matrix of true (latent) rotational states. Then $\Omega^{\frac{1}{2}} \underline{T} \underline{R}^T \in \ker \Pi$.

Proof. It follows from (2) and the definitions of \underline{t}_{ij} in (10) and \underline{T} in (22) that the product $\underline{T} \underline{R}^T \in \mathbb{R}^{m \times d}$ is a $(1 \times d)$ -block structured matrix with rows indexed by the edges $(i, j) \in \vec{\mathcal{E}}$ and whose (i, j) -th row is given by:

$$(\underline{T} \underline{R}^T)_{(i,j)} = -\underline{t}_{ij}^T \underline{R}_i^T = -(R_i \underline{t}_{ij})^T = \underline{t}_i^T - \underline{t}_j^T. \quad (117)$$

Now observe that the quantities $\underline{t}_i^T - \underline{t}_j^T$ associated with each edge $(i, j) \in \vec{\mathcal{E}}$ in the product $\underline{T} \underline{R}^T$ are realizable as differences of values $\underline{t}_i, \underline{t}_j$ assigned to the endpoints of (i, j) , i.e., the columns of $\underline{T} \underline{R}^T$ are realizable as *potential differences* associated with the *potential function* assigning \underline{t}_i^T to vertex $i \in \mathcal{V}$ for all i (Biggs, 1997). Formally, we have from (117) that

$$\underline{T} \underline{R}^T = A(\vec{G})^T \begin{pmatrix} -\underline{t}_1^T \\ \vdots \\ -\underline{t}_n^T \end{pmatrix}, \quad (118)$$

so that the columns of $\underline{T} \underline{R}^T$ lie in $\text{image}(A(\vec{G})^T)$. It follows that $\Omega^{\frac{1}{2}} \underline{T} \underline{R}^T \in \text{image}(\Omega^{\frac{1}{2}} A(\vec{G})^T)$. But $\text{image}(\Omega^{\frac{1}{2}} A(\vec{G})^T) \perp \ker(A(\vec{G})\Omega^{\frac{1}{2}})$ by the Fundamental Theorem of Linear Algebra, and therefore $\Omega^{\frac{1}{2}} \underline{T} \underline{R}^T$ lies in the kernel of the orthogonal projector Π , as claimed.

With the aid of Lemmas 8 and 9, it is now straightforward to show that Problem 7 is always exact in the noiseless case:

Theorem 10. Problem 7 is exact in the noiseless case. Let \underline{Q} be the data matrix of the form (23b) constructed using the true (latent) relative transforms \underline{x}_{ij} in (10). Then $\underline{Z}^* = \underline{R}^T \underline{R}$ is the unique solution of the instance of Problem 7 parameterized by \underline{Q} .

Proof. Since $L(\underline{G}^\rho) \succeq 0$ by Lemma 8(i) and $\underline{Q}^\tau \succeq 0$ (immediate from the definition (23c)), $\underline{Q} = L(\underline{G}^\rho) + \underline{Q}^\tau \succeq 0$ as well, and therefore the optimal value of Problem 5 satisfies $p_O^* \geq 0$. Furthermore, $\underline{R}^T \in \ker(L(\underline{G}^\rho)), \ker(\underline{Q}^\tau)$ by Lemmas 8(iii) and 9, respectively, so $\underline{R}^T \in \ker(\underline{Q})$ as well. This implies that $\text{tr}(\underline{Q} \underline{R}^T \underline{R}) = 0$, and we conclude that \underline{R} is an optimal solution of the noiseless version of Problem 5 (since it is a feasible point that attains the lower bound of 0 for Problem 5's optimal value p_O^*). This also implies $\text{rank}(\underline{Q}) = dn - d$, since $\underline{Q} \succeq L(\underline{G}^\rho)$ (and $L(\underline{G}^\rho)$ has $dn - d$ positive eigenvalues by Lemma 8(i))

and $\text{image}(\underline{R}^T) \subseteq \ker(\underline{Q})$ with $\dim(\text{image}(\underline{R}^T)) = d$. Finally, a straightforward computation using equations (106) and (111) shows that the candidate certificate matrix corresponding to the optimal solution \underline{R} is $\underline{C} = \underline{Q}$. The claim then follows from an application of Theorem 7.

In addition to providing a useful sanity check on the feasibility of our overall strategy by showing that it will at least succeed under ideal conditions, the proof of Theorem 10 also points the way towards a proof of the more general Proposition 2, as we now describe. Observe that in the noiseless case, the certificate matrix $\underline{C} = \underline{Q} = L(\underline{G}^\rho) + \underline{Q}^\tau$ corresponding to the optimal solution \underline{R} has a spectrum consisting of d copies of 0 and $dn - d$ strictly positive eigenvalues that are lower-bounded by $\lambda_2(L(\underline{W}^\rho)) > 0$. Now in the more general (noisy) case, both the data matrix $\tilde{\underline{Q}}$ and the minimizer \underline{R}^* of Problem 5 will vary as a function of the noise added to the measurements \tilde{x}_{ij} in (10), and in consequence so will the matrix \underline{C} . However the first-order condition (105) appearing in Lemma 6 can alternatively be read as $\underline{C} \underline{R}^{*\top} = 0$, which guarantees that \underline{C} always has at least d eigenvalues fixed to 0; furthermore, in general the eigenvalues of a matrix \underline{X} are continuous functions of \underline{X} , and equations (106) and (111) show that \underline{C} is a continuous function of $\tilde{\underline{Q}}$ and \underline{R}^* . Consequently, if we can bound the magnitude of the estimation error $d_O(\underline{R}, \underline{R}^*)$ for a minimizer \underline{R}^* of Problem 5 as a function of the magnitude of the noise $\Delta \underline{Q} = \tilde{\underline{Q}} - \underline{Q}$ corrupting the data matrix $\tilde{\underline{Q}}$, then by controlling $\Delta \underline{Q}$ we can in turn ensure (via continuity) that the eigenvalues of the matrix \underline{C} constructed at the minimizer \underline{R}^* remain nonnegative, and hence Problem 7 will remain exact.

C.4 An upper bound for the estimation error in Problem 5

In this subsection we derive an upper bound on the estimation error $d_O(\underline{R}, \underline{R}^*)$ of a minimizer \underline{R}^* of Problem 5 as a function of the noise $\Delta \underline{Q} \triangleq \tilde{\underline{Q}} - \underline{Q}$ corrupting the data matrix $\tilde{\underline{Q}}$. To simplify the derivation, in the sequel we will assume (without loss of generality) that \underline{R}^* is an element of its orbit (92b) attaining the orbit distance $d_O(\underline{R}, \underline{R}^*)$ defined in (93b).

To begin, the optimality of \underline{R}^* implies that:

$$\begin{aligned} \text{tr}(\tilde{\underline{Q}} \underline{R}^T \underline{R}) &= \text{tr}(\Delta \underline{Q} \underline{R}^T \underline{R}) + \text{tr}(\underline{Q} \underline{R}^T \underline{R}) \\ &\geq \text{tr}(\Delta \underline{Q} \underline{R}^{*\top} \underline{R}^*) + \text{tr}(\underline{Q} \underline{R}^{*\top} \underline{R}^*) \quad (119) \\ &= \text{tr}(\tilde{\underline{Q}} \underline{R}^{*\top} \underline{R}^*). \end{aligned}$$

Now $\text{tr}(QR^\top R) = 0$ since we showed in the previous subsection that $\text{image}(R^\top) = \ker(Q)$, and the identity $\text{tr}(\Delta Q R^\top R) = \text{vec}(R)^\top (\Delta Q \otimes I_d) \text{vec}(R)$ together with the submultiplicativity of the spectral norm shows that

$$\begin{aligned} |\text{tr}(\Delta Q R^\top R)| &\leq \|\Delta Q \otimes I_d\|_2 \|\text{vec}(R)\|_2^2 \\ &= \|\Delta Q\|_2 \|R\|_F^2 = dn \|\Delta Q\|_2 \end{aligned} \quad (120)$$

(and similarly for $|\text{tr}(\Delta Q R^{*\top} R^*)|$); consequently, (119) in turn implies:

$$2dn \|\Delta Q\|_2 \geq \text{tr}(QR^{*\top} R^*). \quad (121)$$

We will now lower-bound the right-hand side of (121) as a function of the estimation error $d_{\mathcal{O}}(\underline{R}, R^*)$, thereby enabling us to upper-bound this error by controlling $\|\Delta Q\|_2$. To do so, we make use of the following:

Lemma 11. Fix $R \in \mathcal{O}(d)^n \subset \mathbb{R}^{d \times dn}$, and let $M = \{WR \mid W \in \mathbb{R}^{d \times d}\} \subset \mathbb{R}^{d \times dn}$ denote the subspace of matrices whose rows are contained in $\text{image}(R^\top)$. Then

$$\begin{aligned} \text{Proj}_V: \mathbb{R}^{dn} &\rightarrow \text{image}(R^\top) \\ \text{Proj}_V(x) &= \frac{1}{n} R^\top R x \end{aligned} \quad (122)$$

is the orthogonal projection operator onto $\text{image}(R^\top)$ with respect to the usual ℓ_2 inner product on \mathbb{R}^{dn} , and the mapping

$$\begin{aligned} \text{Proj}_M: \mathbb{R}^{d \times dn} &\rightarrow M \\ \text{Proj}_M(X) &= \frac{1}{n} X R^\top R \end{aligned} \quad (123)$$

that applies Proj_V to each row of X is the orthogonal projection operator onto M with respect to the Frobenius inner product on $\mathbb{R}^{d \times dn}$.

Proof. If $x \in \text{image}(R^\top)$, then $x = R^\top v$ for some $v \in \mathbb{R}^d$, and

$$\text{Proj}_V(x) = \frac{1}{n} R^\top R (R^\top v) = R^\top v = x, \quad (124)$$

since $RR^\top = nI_d$ as $R \in \mathcal{O}(d)^n$ by hypothesis; this shows that Proj_V is a projection onto $\text{image}(R^\top)$. To show that Proj_V is orthogonal projection with respect to the ℓ_2 inner product on \mathbb{R}^{dn} , it suffices to show that $\text{image}(\text{Proj}_V) \perp \ker(\text{Proj}_V)$. To that end, let

$x, y \in \mathbb{R}^{dn}$, and observe that

$$\begin{aligned} &\langle \text{Proj}_V(x), y - \text{Proj}_V(y) \rangle \\ &= \left\langle \frac{1}{n} R^\top R x, y - \frac{1}{n} R^\top R y \right\rangle_2 \\ &= \frac{1}{n} \langle R^\top R x, y \rangle_2 - \frac{1}{n^2} \langle R^\top R x, R^\top R y \rangle_2 \\ &= \frac{1}{n} x^\top R^\top R y - \frac{1}{n^2} x^\top R^\top R R^\top R y \\ &= 0. \end{aligned} \quad (125)$$

Next, let $X \in \mathbb{R}^{d \times dn}$ and observe that $\text{Proj}_M(X) = \frac{1}{n} X R^\top R$ is the matrix obtained by applying the projection Proj_V to each row of X ; this immediately implies that Proj_M is itself a projection onto M . Furthermore, given $X, Y \in \mathbb{R}^{d \times dn}$, we observe that

$$\begin{aligned} &\langle \text{Proj}_M(X), Y - \text{Proj}_M(Y) \rangle_F \\ &= \langle \text{Proj}_M(X)^\top, Y^\top - \text{Proj}_M(Y)^\top \rangle_F \\ &= \langle \text{vec}(\text{Proj}_M(X)^\top), \text{vec}(Y^\top - \text{Proj}_M(Y)^\top) \rangle_2 \\ &= 0, \end{aligned} \quad (126)$$

since we have already established that Proj_M acts row-wise by Proj_V , which is orthogonal projection with respect to the ℓ_2 inner product.

Since $\ker(Q) = \text{image}(R^\top)$ and $\dim(\text{image}(R^\top)) = d$, it follows from Lemma 11 that

$$\text{tr}(QR^{*\top} R^*) \geq \lambda_{d+1}(Q) \|P\|_F^2, \quad (127)$$

where

$$\begin{aligned} R^* &= K + P, \\ K &= \text{Proj}_M(R^*) = \frac{1}{n} R^* R^\top R, \\ P &= R^* - \text{Proj}_M(R^*) = R^* - \frac{1}{n} R^* R^\top R \end{aligned} \quad (128)$$

is an orthogonal decomposition of R^* with respect to the Frobenius inner product on $\mathbb{R}^{d \times dn}$, and the rows of P are contained in $\text{image}(R^\top)^\perp = \ker(Q)^\perp$. Using (128), we compute:

$$\begin{aligned} \|K\|_F^2 &= \frac{1}{n^2} \text{tr}(R^\top R R^{*\top} R^* R^\top R) \\ &= \frac{1}{n} \text{tr}(R R^{*\top} R^* R^\top) = \frac{1}{n} \|R R^{*\top}\|_F^2 \end{aligned} \quad (129)$$

where we have used the cyclic property of the trace and the fact that $R R^\top = nI_d$. Since (128) is an orthogonal decomposition, it follows that

$$\|P\|_F^2 = \|R^*\|_F^2 - \|K\|_F^2 = dn - \frac{1}{n} \|R R^{*\top}\|_F^2. \quad (130)$$

We may therefore lower-bound $\|P\|_F^2$ by upper-bounding $\|\underline{R}R^{*\top}\|_F^2$ as functions of $d_{\mathcal{O}}(\underline{R}, R^*)$. To that end, recall that R^* is by hypothesis a representative of its orbit (92b) that attains the orbit distance (93b); Theorem 5 then implies that

$$d_{\mathcal{O}}(\underline{R}, R^*)^2 = \|\underline{R} - R^*\|_F^2 = 2dn - 2 \sum_{i=1}^d \sigma_i, \quad (131)$$

where

$$\underline{R}R^{*\top} = U \text{Diag}(\sigma_1, \dots, \sigma_d) V^{\top} \quad (132)$$

is a singular value decomposition of $\underline{R}R^{*\top}$. It follows from (132) and the orthogonal invariance of the Frobenius inner product that

$$\|\underline{R}R^{*\top}\|_F^2 = \|\text{Diag}(\sigma_1, \dots, \sigma_d)\|_F^2 = \sum_{i=1}^d \sigma_i^2, \quad (133)$$

and therefore (131) and (133) imply that we may obtain an upper bound ϵ^2 for $\|\underline{R}R^{*\top}\|_F^2$ in terms of $\delta^2 = d_{\mathcal{O}}(\underline{R}, R^*)^2$ as the optimal value of:

$$\begin{aligned} \epsilon^2 &= \max_{\sigma_i \geq 0} \sum_{i=1}^d \sigma_i^2 \\ \text{s.t. } 2dn - 2 \sum_{i=1}^d \sigma_i &= \delta^2. \end{aligned} \quad (134)$$

The first-order necessary optimality condition for (134) is

$$2\sigma_i = -2\lambda \quad (135)$$

for all $i \in [d]$, where $\lambda \in \mathbb{R}$ is a Lagrange multiplier, and therefore $\sigma_1 = \dots = \sigma_d = \sigma$ for some $\sigma \in \mathbb{R}$. Solving the constraint in (134) for σ shows that

$$\sigma = n - \frac{\delta^2}{2d}, \quad (136)$$

and therefore the optimal value of the objective in (134) is

$$\epsilon^2 = d \left(n - \frac{\delta^2}{2d} \right)^2. \quad (137)$$

Recalling the original definitions of ϵ^2 and δ^2 , we conclude from (137) and (130) that

$$\begin{aligned} \|P\|_F^2 &\geq dn - \frac{d}{n} \left(n - \frac{d_{\mathcal{O}}(\underline{R}, R^*)^2}{2d} \right)^2 \\ &= d_{\mathcal{O}}(\underline{R}, R^*)^2 - \frac{d_{\mathcal{O}}(\underline{R}, R^*)^4}{4dn}. \end{aligned} \quad (138)$$

Applying the inequality $d_{\mathcal{O}}(\underline{R}, R^*)^2 \leq 2dn$ (which follows immediately from the nonnegativity of the

nuclear norm in (96)), we may in turn lower-bound the right-hand side of (138) as:

$$\begin{aligned} d_{\mathcal{O}}(\underline{R}, R^*)^2 - \frac{d_{\mathcal{O}}(\underline{R}, R^*)^4}{4dn} \\ &= \left(1 - \frac{d_{\mathcal{O}}(\underline{R}, R^*)^2}{4dn} \right) d_{\mathcal{O}}(\underline{R}, R^*)^2 \\ &\geq \frac{1}{2} d_{\mathcal{O}}(\underline{R}, R^*)^2. \end{aligned} \quad (139)$$

Finally, combining inequalities (121), (127), (138), and (139), we obtain the following:

Theorem 12. An upper bound for the estimation error in Problem 5. *Let \underline{Q} be the data matrix of the form (23b) constructed using the true (latent) relative transforms $\underline{x}_{ij} = (\underline{t}_{ij}, \underline{R}_{ij})$ in (10), $\underline{R} \in \text{SO}(d)^n$ the matrix composed of the true (latent) rotational states, and $R^* \in \text{O}(d)^n$ an estimate of \underline{R} obtained as a minimizer of Problem 5. Then the estimation error $d_{\mathcal{O}}(\underline{R}, R^*)$ admits the following upper bound:*

$$\sqrt{\frac{4dn \|\tilde{\underline{Q}} - \underline{Q}\|_2}{\lambda_{d+1}(\underline{Q})}} \geq d_{\mathcal{O}}(\underline{R}, R^*). \quad (140)$$

C.5 Finishing the proof

Finally, we complete the proof of Proposition 2 with the aid of Theorems 7 and 12.

Proof of Proposition 2. Let $\underline{R} \in \text{SO}(d)^n$ be the matrix of true (latent) rotations, $R^* \in \text{O}(d)^n$ an estimate of \underline{R} obtained as a minimizer of Problem 5, and assume without loss of generality that R^* is an element of its orbit (92b) attaining the orbit distance $d_{\mathcal{O}}(\underline{R}, R^*)$ defined in (93b). Set $\Delta \underline{Q} \triangleq \tilde{\underline{Q}} - \underline{Q}$ and $\Delta R \triangleq R^* - \underline{R}$, and consider the following decomposition of the certificate matrix C defined in (111) and (106):

where we have used the fact that $\text{image}(\underline{R}^{\top}) = \ker(\underline{Q})$ in passing from lines 2 to 3 above (cf. Lemmas 8 and 9). Observe that the term labeled ΔC in (141) depends continuously upon $\Delta \underline{Q}$ and ΔR , with $\Delta C = 0$ for $(\Delta \underline{Q}, \Delta R) = (0, 0)$; furthermore, $\Delta \underline{Q} \rightarrow 0$ implies $\Delta R \rightarrow 0$ by Theorem 12. It therefore follows from continuity that there exists some $\beta_1 > 0$ such that $\|\Delta C\|_2 < \lambda_{d+1}(\underline{Q})$ for all $\|\Delta \underline{Q}\|_2 < \beta_1$. Moreover, if $\|\Delta C\|_2 < \lambda_{d+1}(\underline{Q})$, it follows from (141) that

$$\lambda_i(C) \geq \lambda_i(\underline{Q}) - \|\Delta C\|_2 > \lambda_i(\underline{Q}) - \lambda_{d+1}(\underline{Q}), \quad (142)$$

and therefore $\lambda_i(C) > 0$ for $i \geq d+1$; i.e., C has at least $dn - d$ strictly positive eigenvalues. Furthermore, Lemma 6 shows that $CR^{*\top} = 0$, which implies that $\ker(C) \supseteq \text{image}(R^{*\top})$; since $\dim(\text{image}(R^{*\top})) =$

$$\begin{aligned}
C &= \tilde{Q} - \text{SymBlockDiag}_d(\tilde{Q}R^{*\top}R^*) \\
&= (\underline{Q} + \Delta Q) - \text{SymBlockDiag}_d((\underline{Q} + \Delta Q)(R + \Delta R)^\top(R + \Delta R)) \\
&= \underline{Q} + \Delta Q - \text{SymBlockDiag}_d\left(\begin{aligned} &\underline{Q}R^\top R + \Delta QR^\top R + \underline{Q}R^\top \Delta R + \underline{Q}\Delta R^\top R \\ &+ \Delta QR^\top \Delta R + \Delta Q\Delta R^\top R + \underline{Q}\Delta R^\top \Delta R + \Delta Q\Delta R^\top \Delta R \end{aligned}\right) \\
&= \underline{Q} + \Delta Q - \underbrace{\text{SymBlockDiag}_d\left(\begin{aligned} &\Delta QR^\top R + \underline{Q}\Delta R^\top R + \Delta QR^\top \Delta R \\ &+ \Delta Q\Delta R^\top R + \underline{Q}\Delta R^\top \Delta R + \Delta Q\Delta R^\top \Delta R \end{aligned}\right)}_{\Delta C},
\end{aligned} \tag{141}$$

d , this in turn implies that C has at least d eigenvalues equal to 0. Since this exhausts C 's dn eigenvalues, we conclude that $C \succeq 0$ and $\text{rank}(C) = dn - d$, and consequently Theorem 7 guarantees that $Z^* = R^{*\top}R^*$ is the unique minimizer of Problem 7.

Now suppose further that $\|\Delta Q\|_2 < \beta_2$ with $\beta_2 \triangleq \lambda_{d+1}(\underline{Q})/2dn$. Then Theorem 12 implies $d_{\mathcal{O}}(R, R^*) = \|R - R^*\|_F < \sqrt{2}$, and therefore in particular that $\|R_i - R_i^*\|_F < \sqrt{2}$ for all $i \in [n]$. But the $+1$ and -1 components of $\text{O}(d)$ are separated by a distance $\sqrt{2}$ under the Frobenius norm, so $R_i \in \text{SO}(d)$ and $\|R_i - R_i^*\|_F < \sqrt{2}$ for all $i \in [n]$ together imply that $R^* \in \text{SO}(d)^n$, and therefore that R^* is in fact an optimal solution of Problem 4 as well.

Proposition 2 then follows from the preceding paragraphs by taking $\beta \triangleq \min\{\beta_1, \beta_2\} > 0$.

Notes

1. The nomenclature originates with the prototypical example of this class of problems: synchronization of clocks over a communications network (Karp et al., 2003; Giridhar and Kumar, 2006) (corresponding to synchronization over the additive group \mathbb{R}).
2. This reprojection operation is often referred to as *rounding*.
3. Note that our definitions of directed and undirected graphs exclude loops and parallel edges. While all of our results can be straightforwardly generalized to admit parallel edges (and indeed our experimental implementation of SE-Sync supports them), we have adopted this restriction in order to simplify the following presentation.
4. If G is not connected, then the problem of estimating the unknown states x_1, \dots, x_n decomposes into a set of independent estimation problems that are in one-to-one correspondence with the connected components of G ; thus, the general case is always reducible to the case of connected graphs.
5. We use a directed graph to model the measurements \tilde{x}_{ij} sampled from (10) because the distribution of the noise corrupting the latent values x_{ij} is not invariant under $\text{SE}(d)$'s group inverse operation, as can be seen by composing (10) with (2b). Consequently, we must keep track of which state x_i was the “base frame” for each measurement.
6. Note that a minimizer of Problem 1 is a maximum-likelihood estimate (rather than *the* maximum-likelihood estimate) because Problem 1 always has multiple (in fact, infinitely many) solutions: since the objective function in (11) is constructed from *relative* measurements of the form $x_i^{-1}x_j$, if $x^* = (x_1^*, \dots, x_n^*) \in \text{SE}(d)^n$ minimizes (11), then $g \bullet x^* \triangleq (g \cdot x_1^*, \dots, g \cdot x_n^*)$ also minimizes (11) for all $g \in \text{SE}(d)$. Consequently, the solution set of Problem 1 is organized into orbits of the diagonal action \bullet of $\text{SE}(d)$ on $\text{SE}(d)^n$. This gauge symmetry simply corresponds to the fact that *relative* measurements $x_i^{-1}x_j$ provide no information about the *absolute* values of the states x_i .
7. There is also some empirical evidence that the relaxation from $\text{SO}(d)$ to $\text{O}(d)$ is not the limiting factor in the exactness of our approach. In our prior work (Tron et al., 2015), we observed that in the specific case $d = 3$ it is possible to replace the (cubic) determinantal constraint in (1b) with an equivalent *quadratic* constraint by using the cross-product operation on the columns of each R_i to enforce the correct orientation; this leads to an equivalent formulation of Problem 1 as a quadratically-constrained quadratic program that can be relaxed directly to a semidefinite program (Luo et al., 2010) *without* the intermediate relaxation through $\text{O}(d)$. We found no significant difference between the sharpness of the relaxation incorporating the determinantal constraint and the relaxation without (Problem 7).
8. Note that Problem 7 could conceivably have multiple solutions, only *some* of which belong to the class specified in Theorem 1; in that case, it is possible that a numerical optimization method might converge to a minimizer of Problem 7 that does *not* correspond to a solution of Problem 1.

9. Ideally, one would like to have both (i) an explicit (i.e. closed-form) expression that lower-bounds the magnitude β of the admissible deviation of the data matrix \tilde{Q} from its exact value Q (as measured in some suitable norm) and (ii) a concentration inequality (Tropp, 2015) (or several) that upper-bounds the probability $p(\|\tilde{Q} - Q\| > \delta)$ of large deviations; together, these would enable the derivation of a lower bound on the probability that a given realization of Problem 4 sampled from the generative model (10) admits an exact semidefinite relaxation (32). While it is possible (with a bit more effort) to derive such lower bounds on β using straightforward (although somewhat tedious) quantitative refinements of the continuity argument given in Appendix C, to date the sharpest concentration inequalities that we have been able to derive appear to be significantly suboptimal, and therefore lead to estimates for the probability of exactness that are grossly conservative versus what we observe empirically (cf. also the discussion in Remark 4.6 and Sec. 5 of (Bandeira et al., 2016)). Consequently, we have chosen to state Proposition 2 as a simple existence result for β in order to simplify its presentation and proof, while still providing some rigorous justification for our convex relaxation approach.
10. We remark that as a practical matter, we have already shown in our previous work (Carlone et al., 2015a) (and do so again here in Section 6) that Problem 7 in fact remains exact with high probability when the measurements \tilde{x}_{ij} in (10) are corrupted with noise up to an order of magnitude greater than what is encountered in typical robotics and computer vision applications; consequently, we leave the derivation of sharper concentration inequalities and explicit lower bounds on the probability of exactness to future research.
11. This is the transpose of a QR decomposition (Golub and Loan, 1996, Sec. 5.2).
12. That is, a point satisfying $\text{grad } F(Y) = 0$ and $\text{Hess } F(Y) \succeq 0$ (cf. (40)–(43)).
13. Note that since *every* $Y \in \text{St}(d, r)^n$ is row rank-deficient for $r > dn$, this procedure is guaranteed to recover an optimal solution after searching at most $dn + 1$ levels of the hierarchy (37).
14. We point out that equations (41), (42), and (43) correspond to equations (7), (8), and (9) in (Boumal, 2015), with the caveat that (Boumal, 2015)’s definition of Y is the *transpose* of ours. Our notation follows the more common convention (cf. e.g. (Edelman et al., 1998)) that elements of a Stiefel manifold are matrices with orthonormal *columns*, rather than *rows*.
15. The requirement of high precision here is not superfluous: because Proposition 3 requires the identification of *rank-deficient* second-order critical points, whatever local search technique we apply to Problem 9 must be capable of numerically approximating a critical point precisely enough that its rank can be correctly determined.
16. This is the *special orthogonal Procrustes* problem, which admits a simple closed-form solution based upon the singular value decomposition (Hanson and Norris, 1981; Umeyama, 1991).
17. Interestingly, the Gauss-Newton method with chordal initialization (GN + chord) appears to obtain remarkably good solutions in the high-rotational-noise regime, although here we cannot certify their correctness, as our solution verification methods (Carlone et al., 2015a) depend upon the same semidefinite relaxation (32) as does SE-Sync.
18. Hence the nomenclature.
19. These are nonlinear programs in which the feasible set is a *semialgebraic set* (the set of real solutions of a system of polynomial (in)equalities) and the objective is a (rational) polynomial function.
20. Recall that $\text{SO}(1) = \{+1\}$, so for $d = 1$ Problems 3 and 4 admit only the (trivial) solution $R^* = (1, \dots, 1)$. Similarly, $\text{O}(1) = \{\pm 1\}$, so for $d = 1$ Problem 5 admits pairs of solutions related by multiplication by -1 .
21. We remark that while the formulation of the distance functions presented in (93) may at first appear somewhat *ad hoc*, one can justify this choice rigorously using the language of Riemannian geometry (Boothby, 2003; Kobayashi and Nomizu, 1996). Since the Frobenius norm distance is orthogonally invariant, the diagonal left-actions (91) of $\text{SO}(d)$ on $\text{SO}(d)^n$ and $\text{O}(d)$ on $\text{O}(d)^n$ are *isometries*, and are trivially free and proper; consequently, the quotient spaces $\mathcal{M}_S \triangleq \text{SO}(d)^n / \text{SO}(d)$ and $\mathcal{M}_O \triangleq \text{O}(d)^n / \text{O}(d)$ obtained by *identifying* the elements of the orbits (92) are manifolds, and the projections $\pi_S: \text{SO}(d)^n \rightarrow \mathcal{M}_S$ and $\pi_O: \text{O}(d)^n \rightarrow \mathcal{M}_O$ are submersions. Furthermore, it is straightforward to verify that the restrictions of the derivative maps $d(\pi_S)_R: T_R(\text{SO}(d)^n) \rightarrow T_{[R]}(\mathcal{M}_S)$ and $d(\pi_O)_R: T_R(\text{O}(d)^n) \rightarrow T_{[R]}(\mathcal{M}_O)$ to the horizontal spaces $H_R(\text{SO}(d)^n) \triangleq \ker(d(\pi_S)_R)^\perp$ and $H_R(\text{O}(d)^n) \triangleq \ker(d(\pi_O)_R)^\perp$ are linear isometries onto $T_{[R]}(\mathcal{M}_S)$ and $T_{[R]}(\mathcal{M}_O)$, respectively, and therefore induce well-defined Riemannian metrics on the quotient spaces \mathcal{M}_S and \mathcal{M}_O from the Riemannian metrics on the total spaces $\text{SO}(d)^n$ and $\text{O}(d)^n$, with corresponding distance functions $d_{\mathcal{M}_S}(\cdot, \cdot)$, $d_{\mathcal{M}_O}(\cdot, \cdot)$, respectively. The functions $d_S(\cdot, \cdot)$ and $d_O(\cdot, \cdot)$ defined in (93) are then simply the functions that report the distances between the images of X and Y after projecting them to

these quotient spaces: $d_S(X, Y) = d_{\mathcal{M}_S}(\pi_S(X), \pi_S(Y))$ and $d_O(X, Y) = d_{\mathcal{M}_O}(\pi_O(X), \pi_O(Y))$. Consequently, these are in fact the canonical distance functions for comparing points in $SO(d)^n$ and $O(d)^n$ while accounting for the gauge symmetry (91).

References

- Åström K and Murray R (2008) *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton, NJ, USA: Princeton University Press.
- Absil PA, Baker C and Gallivan K (2007) Trust-region methods on Riemannian manifolds. *Found. Comput. Math.* 7(3): 303–330.
- Absil PA, Mahony R and Sepulchre R (2008) *Optimization Algorithms on Matrix Manifolds*. Princeton University Press.
- Alizadeh F, Haeberly JPA and Overton M (1997) Complementarity and nondegeneracy in semidefinite programming. *Math. Program.* 77(1): 111–128.
- Arrigoni F, Rossi B and Fusiello A (2016) Spectral synchronization of multiple views in SE(3). *SIAM J. Imaging Sci.* 9(4): 1963–1990.
- Bandeira A (2016) A note on probably certifiably correct algorithms. *Comptes Rendus Mathématique* 354(3): 329–333.
- Bandeira A, Boumal N and Singer A (2016) Tightness of the maximum likelihood semidefinite relaxation for angular synchronization. *Math. Program.* .
- Bandeira A, Singer A and Spielman D (2013) A Cheeger inequality for the graph connection Laplacian. *SIAM J. Matrix Anal. Appl.* 34(4): 1611–1630.
- Best D and Fisher N (1979) Efficient simulation of the von Mises distribution. *Appl. Statist.* 28(2): 152–157.
- Biggs N (1997) Algebraic potential theory on graphs. *Bull. London Math. Soc.* 29: 641–682.
- Boothby W (2003) *An Introduction to Differentiable Manifolds and Riemannian Geometry*. 2nd edition. London: Academic Press.
- Boumal N (2015) A Riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints. ArXiv preprint: arXiv:1506.00575v2.
- Boumal N, Absil PA and Cartis C (2016a) Global rates of convergence for nonconvex optimization on manifolds. ArXiv preprint: arXiv:1605.08101v1.
- Boumal N, Mishra B, Absil PA and Sepulchre R (2014a) Manopt, a MATLAB toolbox for optimization on manifolds. *Journal of Machine Learning Research* 15(1): 1455–1459.
- Boumal N, Singer A, Absil PA and Blondel V (2014b) Cramér-Rao bounds for synchronization of rotations. *Information and Inference* 3: 1–39.
- Boumal N, Voroninski V and Bandeira A (2016b) The non-convex Burer-Monteiro approach works on smooth semidefinite programs. ArXiv preprint arXiv:1606.04970v1.
- Boyd S and Vandenberghe L (2004) *Convex Optimization*. Cambridge University Press.
- Bugarin F, Henrion D and Lasserre J (2016) Minimizing the sum of many rational functions. *Math. Prog. Comp.* 8: 83–111.
- Burer S and Monteiro R (2003) A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Program.* 95: 329–357.
- Burer S and Monteiro R (2005) Local minima and convergence in low-rank semidefinite programming. *Math. Program.* 103: 427–444.
- Burschka D, Corso J, Dewan M, Lau W, Li M, Lin H, Marayong P, Ramey N, Hager G, Hoffinan B, Larkin D and Hasser C (2005) Navigating inner space: 3-D assistance for minimally invasive surgery. *Robotics and Autonomous Systems* 52(1): 5–26.
- Candès E, Li X, Ma Y and Wright J (2011) Robust principal component analysis? *Journal of the ACM* 58(3): Article 11.
- Carlone L, Rosen D, Calafiore G, Leonard J and Dellaert F (2015a) Lagrangian duality in 3D SLAM: Verification techniques and optimal solutions. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. Hamburg, Germany.
- Carlone L, Tron R, Daniilidis K and Dellaert F (2015b) Initialization techniques for 3D SLAM: A survey on rotation estimation and its use in pose graph optimization. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. Seattle, WA, pp. 4597–4605.
- Chandrasekaran V, Recht B, Parrilo P and Willsky A (2012) The convex geometry of linear inverse problems. *Found. Comput. Math.* 12: 805–849.

- Chiuso A, Picci G and Soatto S (2008) Wide-sense estimation on the special orthogonal group. *Communications in Information and Systems* 8(3): 185–200.
- Cucuringu M, Lipman Y and Singer A (2012) Sensor network localization by eigenvector synchronization over the special Euclidean group. *ACM Trans. on Sensor Networks* 8(3): 19:1–19:42.
- Dellaert F, Carlson J, Ila V, Ni K and Thorpe C (2010) Subgraph-preconditioned conjugate gradients for large scale SLAM. In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. Taipei, Taiwan, pp. 2566–2571.
- Dellaert F and Kaess M (2006) Square Root SAM: Simultaneous localization and mapping via square root information smoothing. *Intl. J. of Robotics Research* 25(12): 1181–1203.
- Dembo R and Steihaug T (1983) Truncated-Newton algorithms for large-scale unconstrained optimization. *Math. Program.* 26: 190–212.
- Edelman A, Arias T and Smith S (1998) The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.* 20(2): 303–353.
- Fallon M, Kuindersma S, Karumanchi S, Antone M, Schneider T, Dai H, D’Arpino C, Deits R, DiCicco M, Fourie D, Koolen T, Marion P, Posa M, Valenzuela A, Yu KT, Shah J, Iagnemma K, Tedrake R and Teller S (2015) An architecture for online affordance-based perception and whole-body planning. *Intl. J. of Robotics Research* 32(2): 229–254.
- Ferguson T (1996) *A Course in Large Sample Theory*. Boca Raton, FL: Chapman & Hall/CRC.
- Fisher R (1953) Dispersion on a sphere. In: *Proc. Roy. Soc.*, volume 217. pp. 295–305.
- Gallier J (2010) The Schur complement and symmetric positive semidefinite (and definite) matrices. Unpublished note, available online: <http://www.cis.upenn.edu/~jean/schur-comp.pdf>.
- Giridhar A and Kumar P (2006) Distributed clock synchronization over wireless networks: Algorithms and analysis. In: *IEEE Conf. on Decision and Control*. San Diego, CA, pp. 4915–4920.
- Golub G and Loan CV (1996) *Matrix Computations*. 3rd edition. Baltimore, MD: Johns Hopkins University Press.
- Grisetti G, Kümmerle R, Stachniss C and Burgard W (2010) A tutorial on graph-based SLAM. *IEEE Intelligent Transportation Systems Magazine* 2(4): 31–43.
- Grisetti G, Stachniss C and Burgard W (2009) Nonlinear constraint network optimization for efficient map learning. *IEEE Trans. on Intelligent Transportation Systems* 10(3): 428–439.
- Guillemin V and Pollack A (1974) *Differential Topology*. Prentice-Hall.
- Hanson R and Norris M (1981) Analysis of measurements based on the singular value decomposition. *SIAM J. Sci. Stat. Comput.* 2(3): 363–373.
- Hartley R, Trunf J, Dai Y and Li H (2013) Rotation averaging. *Intl. J. of Computer Vision* 103(3): 267–305.
- Horn R and Johnson C (1991) *Topics in Matrix Analysis*. Cambridge University Press.
- Huang S, Lai Y, Frese U and Dissanayake G (2010) How far is SLAM from a linear least squares problem? In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. Taipei, Taiwan, pp. 3011–3016.
- Huang S, Wang H, Frese U and Dissanayake G (2012) On the number of local minima to the point feature based SLAM problem. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. St. Paul, MN, pp. 2074–2079.
- Johnson M, Shrewsbury B, Bertrand S, Wu T, Duran D, Floyd M, Abeles P, Stephen D, Mertins N, Lesman A, Carff J, Rifenburgh W, Kaveti P, Straatman W, Smith J, Griffioen M, Layton B, de Boer T, Koolen T, Neuhaus P and Pratt J (2015) Team IHMC’s lessons learned from the DARPA Robotics Challenge Trials. *Intl. J. of Robotics Research* 32(2): 192–208.
- Kaess M, Johannsson H, Roberts R, Ila V, Leonard J and Dellaert F (2012) iSAM2: Incremental smoothing and mapping using the Bayes tree. *Intl. J. of Robotics Research* 31(2): 216–235.
- Karp R, Elson J, Estrin D and Shenker S (2003) Optimal and global time synchronization in sensor networks. Technical report, Center for Embedded Network Sensing, University of California.
- Kobayashi S and Nomizu K (1996) *Foundations of Differential Geometry*, volume 1. Wiley Classics.

- Konolige K (2010) Sparse sparse bundle adjustment. In: *Proc. British Machine Vision Conf. (BMVC)*. pp. 1–11.
- Kümmerle R, Grisetti G, Strasdat H, Konolige K and Burgard W (2011) g2o: A general framework for graph optimization. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. Shanghai, China, pp. 3607–3613.
- Lasserre J (2001) Global optimization and the problem of moments. *SIAM J. Optim.* 11(3): 796–817.
- Lasserre J (2006) Convergent SDP-relaxations in polynomial optimization with sparsity. *SIAM J. Optim.* 17(3): 822–843.
- Laurent M (2009) Sums of squares, moment matrices and optimization over polynomials. In: Putinar M and Sullivant S (eds.) *Emerging Applications of Algebraic Geometry, The IMA Volumes in Mathematics and Its Applications*, volume 149. New York: Springer, pp. 157–270.
- Leonard J, How J, Teller S, Berger M, Campbell S, Fiore G, Fletcher L, Frazzoli E, Huang A, Karaman S, Koch O, Kuwata Y, Moore D, Olson E, Peters S, Teo J, Truax R, Walter M, Barrett D, Epstein A, Maheloni K, Moyer K, Jones T, Buckley R, Antone M, Galejs R, Krishnamurthy S and Williams J (2008) A perception-driven autonomous urban vehicle. *Intl. J. of Robotics Research* 25(10): 727–774.
- Lourakis M and Argyros A (2009) SBA: A software package for generic sparse bundle adjustment. *ACM Trans. Math. Softw.* 36(1): 1–30.
- Luo ZQ, Ma WK, So A, Ye Y and Zhang S (2010) Semidefinite relaxation of quadratic optimization problems. *IEEE Signal Processing Magazine* 27(3): 20–34.
- Martinec D and Pajdla T (2007) Robust rotation and translation estimation in multiview reconstruction. In: *IEEE Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*. Minneapolis, MN, pp. 1–8.
- Meyer C (2000) *Matrix Analysis and Applied Linear Algebra*. Philadelphia, PA: The Society for Industrial and Applied Mathematics (SIAM).
- Nash S (2000) A survey of truncated-Newton methods. *Journal of Computational and Applied Mathematics* 124: 45–59.
- Nie J and Demmel J (2008) Sparse SOS relaxations for minimizing functions that are summations of small polynomials. *SIAM J. Optim.* 19(4): 1534–1558.
- Nie J, Demmel J and Sturmfels B (2006) Minimizing polynomials via sum of squares over the gradient ideal. *Math. Program.* 106: 587–606.
- Nocedal J and Wright S (2006) *Numerical Optimization*. 2nd edition. New York: Springer Science+Business Media.
- Olson E, Leonard J and Teller S (2006) Fast iterative alignment of pose graphs with poor initial estimates. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. Orlando, FL, pp. 2262–2269.
- Olver F, Olde Daalhuis A, Lozier D, Schneider B, Boisvert R, Clark C, Miller B and Saunders B (eds.) (2016) *NIST Digital Library of Mathematical Functions*, Release 1.0.13. Available online: <http://dlmf.nist.gov/>.
- Özyeşil O, Singer A and Basri R (2015) Stable camera motion estimation using convex programming. *SIAM J. Imaging Sci.* 8(2): 1220–1262.
- Parrilo P (2003) Semidefinite programming relaxations for semialgebraic problems. *Math. Program.* 96(2): 293–320.
- Peters J, Borra D, Paden B and Bullo F (2015) Sensor network localization on the group of three-dimensional displacements. *SIAM J. Control Optim.* 53(6): 3534–3561.
- Pratt G and Manzo J (2013) The DARPA Robotics Challenge. *IEEE Robotics & Automation Magazine* 20(2): 10–12.
- Rosen D, Carlone L, Bandeira A and Leonard J (2016) A certifiably correct algorithm for synchronization over the special Euclidean group. In: *Intl. Workshop on the Algorithmic Foundations of Robotics (WAFR)*. San Francisco, CA.
- Rosen D, DuHadway C and Leonard J (2015) A convex relaxation for approximate global optimization in simultaneous localization and mapping. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. Seattle, WA, pp. 5822–5829.
- Rosen D, Kaess M and Leonard J (2014) RISE: An incremental trust-region method for robust online sparse least-squares estimation. *IEEE Trans. on Robotics* 30(5): 1091–1108.

- Russel S and Norvig P (2010) *Artificial Intelligence: A Modern Approach*. 3rd edition. Upper Saddle River, NJ, USA: Prentice Hall.
- Singer A (2011) Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis* 30(1): 20–36.
- Singer A and Wu HT (2012) Vector diffusion maps and the connection Laplacian. *Comm. Pure Appl. Math.* 65: 1067–1144.
- Stachniss C, Leonard J and Thrun S (2016) Simultaneous localization and mapping. In: *Springer Handbook of Robotics*. Springer International Publishing, pp. 1153–1176.
- Steihaug T (1983) The conjugate gradient method and trust regions in large scale optimization. *SIAM J. Numer. Anal.* 20(3): 626–637.
- Stengel R (1994) *Optimal Control and Estimation*. Dover Publications.
- Taylor R (2006) A perspective on medical robotics. *Proc. of the IEEE* 94(9): 1652–1664.
- Taylor R, Menciassi A, Fichtinger G and Dario P (2008) Medical robotics and computer-integrated surgery. In: *Springer Handbook of Robotics*. Springer, pp. 1199–1222.
- Thrun S, Montemerlo M, Dahlkamp H, Stavens D, Aron A, Diebel J, Fong P, Gale J, Halpenny M, Hoffmann G, Lau K, Oakley C, Palatucci M, Pratt V, Stang P, Strohband S, Dupont C, Jendrossek LE, Koelen C, Markey C, Rummel C, van Niekerk J, Jensen E, Alessandrini P, Bradski G, Davies B, Ettinger S, Kaehler A, Nefian A and Mahoney P (2006) Stanley: The robot that won the DARPA Grand Challenge. *Intl. J. of Robotics Research* 23(9): 661–692.
- Todd M (2001) Semidefinite optimization. *Acta Numerica* 10: 515–560.
- Tron R, Rosen D and Carlone L (2015) On the inclusion of determinant constraints in Lagrangian duality for 3D SLAM. Presented at Robotics: Science and Systems (RSS) in the workshop “The Problem of Mobile Sensors”.
- Tron R, Zhou X and Daniilidis K (2016) A survey on rotation optimization in structure from motion. In: *IEEE Intl. Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops*. Las Vegas, NV, pp. 77–85.
- Tropp J (2015) An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning* 8(1–2): 1–230.
- Umeyama S (1991) Least-squares estimation of transformation parameters between two point patterns. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 13(4): 376–380.
- Urmson C, Anhalt J, Bagnell D, Baker C, Bittner R, Clark M, Dolan J, Duggins D, Galatali T, Geyer C, Gittleman M, Harbaugh S, Hebert M, Howard T, Kolski S, Kelly A, Likhachev M, McNaughton M, Miller N, Peterson K, Pilnick B, Rajkumar R, Rybski P, Salesky B, Seo YW, Singh S, Snider J, Stentz A, Whittaker W, Wolkowicki Z, Ziglar J, Bae H, Brown T, Demitrish D, Litkouhi B, Nickolaou J, Sadekar V, Zhang W, Struble J, Taylor M, Darms M and Ferguson D (2008) Autonomous driving in urban environments: Boss and the Urban Challenge. *Intl. J. of Robotics Research* 25(8): 425–466.
- Vandenberghe L and Boyd S (1996) Semidefinite programming. *SIAM Review* 38(1): 49–95.
- Waki H, Kim S, Kojima M and Muramatsu M (2006) Sums of squares and semidefinite program relaxations for polynomial optimization problems with structured sparsity. *SIAM J. Optim.* 17(1): 218–242.
- Wang H, Hu G, Huang S and Dissanayake G (2012) On the structure of nonlinearities in pose graph SLAM. In: *Robotics: Science and Systems (RSS)*. Sydney, Australia.
- Wang L and Singer A (2013) Exact and stable recovery of rotations for robust synchronization. *Information and Inference* .
- Warner F (1983) *Foundations of Differentiable Manifolds and Lie Groups*. Springer Graduate Texts in Mathematics. New York: Springer.
- Zhou Z, Li X, Wright J, Candès E and Ma Y (2010) Stable principal component pursuit. In: *IEEE Intl. Symp. on Information Theory (ISIT)*. Austin, TX, pp. 1518–1522.