SCORE: A Second-Order Conic Initialization for Range-Aided SLAM

Alan Papalia[1,2], Joseph Morales[1], Kevin J. Doherty[1,2], David M. Rosen[3,4], John J. Leonard[1]

***Abstract*— We present a novel initialization technique for the range-aided simultaneous localization and mapping (RA-SLAM) problem. In RA-SLAM we consider measurements of point-to-point distances in addition to measurements of rigid transformations to landmark or pose variables. Standard formulations of RA-SLAM approach the problem as non-convex optimization, which requires a good initialization to obtain quality results. The initialization technique proposed here relaxes the RA-SLAM problem to a convex problem which is then solved to determine an initialization for the original, non-convex problem. The relaxation is a second-order cone program (SOCP), which is derived from a quadratically constrained quadratic program (QCQP) formulation of the RA-SLAM problem. As a SOCP, the method is highly scalable. We name this relaxation Second-order COnic RElaxation for RA-SLAM (SCORE). To our knowledge, this work represents the first convex relaxation for RA-SLAM. We present real-world and simulated experiments which show SCORE initialization permits the efficient recovery of quality solutions for a variety of challenging single- and multi-robot RA-SLAM problems with thousands of poses and range measurements.**

I. Introduction

Range-aided simultaneous localization and mapping (RA-SLAM) is a key robotic task, with applications in planetary [1], subterranean [2], and sub-sea [3]–[5] environments. RA-SLAM combines range measurements (e.g., distances between acoustic beacons) with measurements of relative rigid transformations (e.g., odometry) to estimate a set of robot poses and landmark positions.

RA-SLAM differs from pose-graph simultaneous localization and mapping (SLAM) in that the sensing models of range measurements induce substantial difficulties for state-estimation. However, ranging is a valuable sensor modality and further developments in RA-SLAM could substantially advance the state of robot navigation.

The state-of-the-art formulation of SLAM problems is as *maximum a posteriori* (MAP) inference, which takes the form of an optimization problem. However, because robot orientations are a non-convex set, the MAP problem is non-convex. Additionally, in RA-SLAM, the measurement models of range sensing introduce further non-convexities. As a result, standard approaches to solving these MAP problems use local-search techniques and can only guarantee locally optimal solutions. Thus, good initializations to the MAP problem are key to obtaining quality RA-SLAM solutions.

[1] Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology (MIT), Cambridge, MA 02139, USA, {apapalia, jrales, kdoherty, jleonard}@mit.edu
[2] Department of Applied Ocean Physics & Engineering, Woods Hole Oceanographic Institution, Woods Hole, MA 02543, USA
[3] Department of Electrical & Computer Engineering, Northeastern University, Boston, MA 02115, USA, d.rosen@northeastern.edu
[4] Department of Math, Northeastern University, Boston, MA 02115, USA

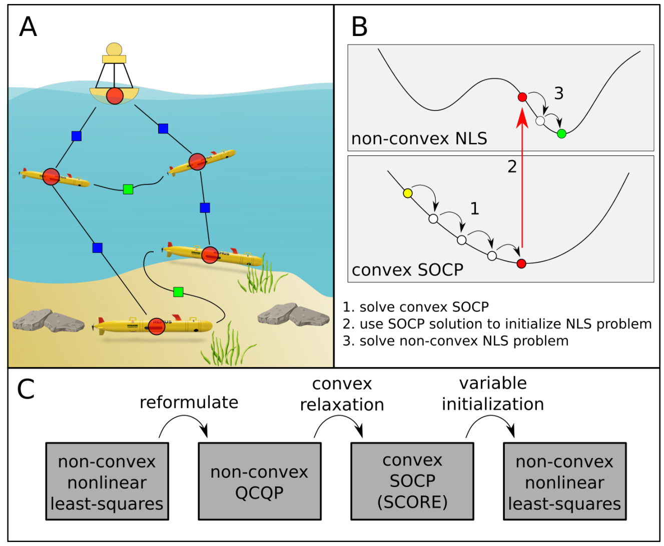

Fig. 1. Key concepts. (A) An example multi-robot range-aided SLAM problem. Red circles are problem variables (i.e., robot poses and beacon position), green squares are relative pose measurements (e.g., odometry), and blue squares are distance measurements (e.g., acoustic ranging). (B) A second-order cone program (SOCP) is derived as a convex relaxation of the original non-convex nonlinear least-squares (NLS) formulation of the range-aided SLAM problem. The solution to the SOCP is then used to obtain the initial estimate for the NLS problem, and the NLS problem is solved to refine the SOCP initialization. (C) Illustration of contributions. The range-aided SLAM problem is first represented as a non-convex NLS problem. This work shows how to reformulate the NLS problem into a non-convex quadratically constrained quadratic program (QCQP) and how to relax the QCQP into a convex SOCP. The SOCP can be solved to obtain an initial estimate for the original NLS problem.

This work approaches initialization for the non-convex MAP estimation problem through *convex relaxation* [6], the construction of a convex optimization problem which attempts to approximate the original problem. Convex problems can be efficiently solved to global optimality, and many convex relaxations have shown success as initialization strategies in related robotic problems [7]–[10].

This work presents SCORE as a novel methodology for initializing RA-SLAM problems. SCORE applies to 2D and 3D scenarios with arbitrary numbers of robots and landmarks. As SCORE is a SOCP, it is easily implemented in existing convex optimization libraries, and scales gracefully to large problems. We summarize our contributions as follows:

- A novel, convex-relaxation approach to initializing RA-SLAM problems.
- The first QCQP formulation of RA-SLAM, connecting RA-SLAM to a broader body of work [8]–[11].
- Our implementation of SCORE is open-sourced[1].

[1] https://github.com/MarineRoboticsGroup/score

II. Related Work

The current state-of-the-art formulation of RA-SLAM is through nonlinear least-squares (NLS) optimization based upon MAP inference [12]. This work specifically focuses on *initialization* for this MAP formulation. Though we focus on MAP estimation, there are many important previous formulations of RA-SLAM using: extended Kalman filters [3], [13]–[15], particle filters [16], [17], mixture models [18], and NLS optimization [1], [2], [19].

We first survey previous initialization strategies for RA-SLAM problems and then expand the scope to cover initialization strategies for general pose-graph SLAM problems. As our initialization approach is a convex relaxation of the RA-SLAM problem, we discuss related convex relaxations in robotic state-estimation. Finally, as the key challenge that differentiates RA-SLAM from pose-graph SLAM is the nature of range measurements, we discuss the closely related field of sensor network localization (SNL) and convex relaxations developed specifically for SNL problems.

A. Initialization in RA-SLAM

Notable work in the single-robot RA-SLAM case used spectral graph clustering to initially estimate beacon locations [20]. Similarly, [21] used spectral decomposition to jointly estimate robot poses and beacon locations. However, these approaches [20], [21] do not account for sensor noise models and do not readily extend to multi-robot scenarios. Other initializations for multi-robot RA-SLAM used coordinated movement patterns to find the (single) relative transform between robots [22], [23].

Finally, a series of works developed increasingly sophisticated methods to compute relative transforms between two robots using range and odometry measurements [24]–[30]. While these approaches represent great progress, they are only capable of solving for a single inter-robot relative transform and rely on (noisy) odometric pose composition to initialize all other poses. In fact, [28]–[30] utilize convex relaxation to solve the relative transformation problem.

B. Initialization and Convex Relaxation in State-Estimation

In pose-graph SLAM many existing initializations solve approximations of the MAP problem. [7] leveraged the relationships between rotation and translation measurements in SLAM and explored the use of several rotation averaging algorithms [31]–[34] to obtain initial estimates. The same relationships were used to approximate the SLAM problem as two sequential linear estimation problems [35]. However, these approaches do not consider range measurements and thus do not extend to RA-SLAM problems.

With exception to [34], these initialization procedures represent convex relaxations of the pose-graph SLAM problem. Subsequent works in pose-graph SLAM developed convex relaxations which obtained exact solutions to the non-convex MAP problem [8], [11], [36]–[38]. Other works developed exact convex relaxations for the problems of extrinsic calibration [9], two-view relative pose estimation [39], [40], and spline-based trajectory estimation from range measurements with known beacon locations [41]. With exception to [41], these convex relaxations also do not extend to range measurements. However, [41] required known beacon locations, allowed only for range measurements, and only considered single-robot localization. SCORE allows for multi-robot RA-SLAM, and combines measurements of rigid transformations and ranges without necessitating any information *a priori*.

A number of convex relaxations exist for sensor network localization (SNL). As SNL centers around point-to-point distance measurements, these works are closely related to SCORE. Previous works developed and analyzed semidefinite program (SDP) [10], [42], [43] and SOCP [44]–[47] relaxations of the SNL problem. Whereas these relaxations only consider range measurements and can only estimate the Euclidean position of points, SCORE allows for rigid transformation measurements and estimation of poses. Additionally, as SCORE is a SOCP, it is substantially more scalable than these SDP relaxations.

C. Novelty of Our Approach

SCORE is similar to [31] in its relaxation of rotations and to [45] in its relaxation of range measurements. However, it differs from these works in that it jointly considers both relative pose and range measurements. Notably, SCORE generalizes the chordal-initialization of [31] and the SOCP relaxation of [45].

Additionally, there is much commonality to [28]–[30] in that the convex relaxation of SCORE considers both range and relative transformation measurements. Unlike these works, SCORE generalizes to arbitrary dimensions (e.g. 2- or 3-D) and multi-robot cases. Furthermore, SCORE differs in that the convex relaxation is derived from the full MAP problem, and thus utilizes all measurements and jointly solves for an initialization of all RA-SLAM variables.

III. Notation and Mathematical Preliminaries

Symbols: This work generalizes across dimensionalities (e.g., 2-D or 3-D). As such the mathematical presentation will assume to be working in d-dimensional space. We denote the d-dimensional identity matrix as I_d and the special orthogonal group as $\mathrm{SO(d)}$. We refer to the set of real values as $\mathbb{R}$ and the set of nonnegative reals as $\mathbb{R}_+$. Furthermore, we indicate vectors and scalars with lowercase symbols (e.g., t) and matrices with uppercase symbols (e.g., R). Noisy measurements are indicated with a tilde (e.g., $\tilde{R}$). The true (latent) value of a quantity will be underlined (e.g., $\underline{R}$)

Mathematical Operators: We denote the matrix trace as $\mathrm{tr}(\cdot)$. A diagonal matrix with diagonal elements $[a_1, \dots, a_n]$ is indicated as $\mathrm{diag}([a_1, \dots, a_n])$. The vector 2-norm is indicated as $\|\cdot\|_2$ and the Frobenius norm is indicated as $\|\cdot\|_F$. We note the following relationship between the Frobenius norm and trace operators: $\|A\|_F^2 = \mathrm{tr}(AA^\top) = \mathrm{tr}(A^\top A)$.

Probability Distributions: We indicate a Gaussian probability distribution over Euclidean space, with mean $\mu \in \mathbb{R}^d$ and variance $\Sigma \in \mathbb{R}^{d\times d}$ as $\mathcal{N}(\mu, \Sigma)$. Similarly, we denote the isotropic Langevin probability distribution[2] over matrix elements of the special orthogonal group, $\mathrm{SO(d)}$, with mode $R_\mu \in \mathbb{R}^{d\times d}$ and concentration parameter $\kappa \in \mathbb{R}$ as $\mathrm{Langevin}(R_\mu, \kappa)$ [48]. This is the distribution whose probability density function is:

$$\begin{aligned} &R^\epsilon \sim \text{Langevin}(R_\mu, \kappa) \\ &\mathrm{P}(R^\epsilon) = \tfrac{1}{c(\kappa)} \exp\{\kappa \operatorname{tr}(R_\mu^\top R^\epsilon)\} \end{aligned} \tag{1}$$

with normalizing constant $c(\kappa)$.

IV. Range-Aided SLAM Formulation

In this section we model the RA-SLAM problem and derive its corresponding MAP estimation problem.

A. Inference Over a Graph

We formulate RA-SLAM as inference over a directed graph, with nodes as variables (i.e. robot poses and landmark positions) and the edges representing measurements. Without loss of generality, we assume directionality in the edges. The edge-set of all measurements of relative rigid transformations is denoted $\mathcal{E}_p$, where $(i,j) \in \mathcal{E}_p$ represents a measured rigid transformation from node i to node j. Similarly, the edge-set of all range measurements is denoted $\mathcal{E}_r$ and each $(i,j) \in \mathcal{E}_r$ represents a measured distance from node i to node j.

B. Measurement Noise Models

For relative rigid transformation measurements we follow the formulation of [11] in the noise model we assume over our measurements. We denote the true translation and orientation of node i as $\underline{t}_i \in \mathbb{R}^d$ and $\underline{R}_i \in \mathrm{SO(d)}$, with the true relative rotation from node i to node j as $\underline{R}_{ij} \triangleq \underline{R}_i^\top \underline{R}_j$. Similarly, we denote the true relative translations as $\underline{t}_{ij} \triangleq \underline{R}_i^\top (\underline{t}_j - \underline{t}_i)$. We indicate noisy measurements with a tilde (e.g., $\tilde{R}$):

$$\begin{aligned} \tilde{t}_{ij} &= \underline{t}_{ij} + t_{ij}^\epsilon, \quad t_{ij}^\epsilon \sim \mathcal{N}\left(0, \tau_{ij}^{-1} I_d\right), \\ \tilde{R}_{ij} &= \underline{R}_{ij} R_{ij}^\epsilon, \quad R_{ij}^\epsilon \sim \text{Langevin}\left(I_d, \kappa_{ij}\right). \end{aligned} \tag{2}$$

We assume the following generative Gaussian model for range measurements:

$$\tilde{d}_{ij} = \|\underline{t}_i - \underline{t}_j\|_2 + d_{ij}^\epsilon, \quad d_{ij}^\epsilon \sim \mathcal{N}\left(0, \sigma_{ij}^2\right). \tag{3}$$

C. MAP Estimation for RA-SLAM

We write the MAP problem corresponding to the sensor models of Section IV-B. The MAP problem is as follows, where $t_i \in \mathbb{R}^d$ and $R_i \in \mathrm{SO(d)}$ denote the variables corresponding to the estimated translation and rotation of node i. Bolded symbols indicate groupings (e.g., $\mathbf{t} \triangleq \{t_i \ \forall\ i = 1, \ldots, n\}$):

$$\max_{\mathbf{t}, \mathbf{R}} \ \mathrm{P}(\tilde{\mathbf{t}}, \tilde{\mathbf{R}}, \tilde{\mathbf{d}} \mid \mathbf{t}, \mathbf{R}). \tag{4}$$

This general MAP problem of Eq. (4) is equivalently solved by minimizing its negative log-likelihood [12]. Given the models of Section IV-B, the MAP estimate of the RA-SLAM problem can be expressed as the following NLS problem:

$$\begin{aligned} \min_{\substack{t_i \in \mathbb{R}^d \\ R_i \in \mathrm{SO(d)}}} \quad & \sum_{(i,j)\in\mathcal{E}_p} \kappa_{ij} \|R_j - R_i \tilde{R}_{ij}\|_F^2 \\ &+ \sum_{(i,j)\in\mathcal{E}_p} \tau_{ij} \left\|t_j - t_i - R_i \tilde{t}_{ij}\right\|_2^2 \\ &+ \sum_{(i,j)\in\mathcal{E}_r} \frac{1}{\sigma_{ij}^2} (\|t_i - t_j\|_2 - \tilde{d}_{ij})^2 . \end{aligned} \tag{5}$$

where $\|\cdot\|_F$ denotes the matrix Frobenius norm:

In Eq. (5) the first two summands correspond to relative transformation measurements whereas the third summand corresponds to range measurements. Derivation of the cost terms in the first two summands can be found in [11]. The cost terms in the third summand follow immediately from the negative log-likelihood of the model in Eq. (3).

Notice that there are two distinct sources of non-convexity in Eq. (5). The rotation variables are members of the special orthogonal group, i.e., $R_i \in \mathrm{SO(d)}$, which is a non-convex set. Additionally, the cost terms due to range measurements are non-convex due to the $\|t_i - t_j\|_2$ component.

V. SCORE: Second Order Conic Relaxation

In this section we reformulate the MAP problem of Eq. (5) as a QCQP. We then show how to construct SCORE by relaxing this QCQP to a SOCP.

A. RA-SLAM as a QCQP

We perform two changes to convert the non-convex NLS problem of Eq. (5) problem to the QCQP of Eq. (6). First, we rewrite the $\mathrm{SO(d)}$ constraint as the orthonormality constraint and remove the associated determinant constraint. Secondly, we introduce the auxiliary variables, $d_{ij} \in \mathbb{R}_+$ to rewrite the range-measurement cost terms as quadratics. In this formulation the cost terms are quadratic and convex with respect to the problem variables and all constraints are quadratic equality constraints. Thus, Eq. (6) is a QCQP.

Previous work [8], [36], [37] demonstrated that $\det(R_i) = 1$ can be written as a set of quadratic equalities and that the contribution of this determinant constraint is negligible, justifying its removal. With exception to the determinant constraint, this QCQP is an exact reformulation of Eq. (5):

$$\begin{aligned} \min_{\substack{t_i \in \mathbb{R}^d \\ R_i \in \mathbb{R}^{d\times d} \\ d_{ij} \in \mathbb{R}_+}} \quad & \sum_{(i,j)\in\mathcal{E}_p} \kappa_{ij} \|R_j - R_i \tilde{R}_{ij}\|_F^2 \\ &+ \sum_{(i,j)\in\mathcal{E}_p} \tau_{ij} \left\|t_j - t_i - R_i \tilde{t}_{ij}\right\|_2^2 \\ &+ \sum_{(i,j)\in\mathcal{E}_r} \frac{1}{\sigma_{ij}^2} (d_{ij} - \tilde{d}_{ij})^2 \\ \text{subject to} \quad & R_i R_i^\top = I_d, \ i = 1, \ldots, n \\ & d_{ij}^2 = \|t_i - t_j\|_2^2, \ \forall (i,j) \in \mathcal{E}_r . \end{aligned} \tag{6}$$

[2]This work assumes isotropic Langevin noise for simplicity of presentation, but our approach generalizes to anisotropic Langevin noise.

B. Relaxing Equation (6) to a SOCP

We will now introduce the primary contribution of our paper, a Second-order COnic RElaxation for RA-SLAM (SCORE). As previously mentioned, SCORE naturally arises as a relaxation of the QCQP in Eq. (6). We emphasize that Eq. (6) is not solved in our approach, as it is not computationally practical to do so. Eq. (6) is important as an intermediate representation to derive SCORE.

We relax the orthonormality constraints by removing them entirely, eliminating that specific source of non-convexity. To eliminate the non-convexity from the quadratic equality constraint on our auxiliary variables (d_{ij}) we first recall that $d_{ij} \geq 0$, so we can drop the squares on each side of the inequality without modifying the feasible set. From here, we can relax the equality constraint to the convex inequality constraint $d_{ij} \geq \|t_i - t_j\|_2$. Importantly, this specific constraint relating the auxiliary variables and position variables is a (convex) second-order cone constraint [6], [49]. This relaxed problem takes the following form, where $\mathrm{F}(\mathbf{R}, \mathbf{t}, \mathbf{d})$ is the same cost function as in Eq. (6).

$$\begin{aligned} \min_{\mathbf{t},\mathbf{R},\mathbf{d}} \quad & \mathrm{F}(\mathbf{R}, \mathbf{t}, \mathbf{d}) \\ \text{subject to} \quad & d_{ij} \geq \|t_i - t_j\|_2, \ \forall (i,j) \in \mathcal{E}_r \\ & R_0 = I_d \end{aligned} \tag{7}$$

With this previously mentioned second-order cone constraint on d_{ij}, it follows that Eq. (7) is a SOCP, i.e., it is the minimization of quadratic cost functions over the intersection of an affine subspace with second-order cones. This is a known form of SOCPs [49], and thus existing solvers (e.g., [50]) can efficiently solve Eq. (7) to global optimality. Similarly to [31], the first rotation (R_0) is constrained to prevent a trivial solution of all zeros.

VI. Using SCORE to Initialize Local-Optimization

Though SCORE can be solved to global optimality, in general the SCORE solution will differ from that of the original problem due to the relaxed constraints. This is not a serious limitation, as the intent is that the SCORE solution will be close to the solution of the original problem. If true, the SCORE solution can be projected to the feasible set of the original problem, and thereby serve as a robust and principled initialization procedure.

Projecting the SCORE solution to the original problem's feasible set only involves the estimated rotations. As SCORE relaxes the $\mathrm{SO(d)}$ constraints, the estimated values, R_i, will generally not be valid rotation matrices. As in [31], we project the approximate rotation estimates to the nearest rotation matrix via singular value decomposition, as follows:

$$R_i^* = \underset{R_i}{\operatorname{argmin}} \quad \|R_i - \hat{R}_{i,approx}\|_F^2 \tag{8}$$

$$U, S, V^\top = \mathrm{SVD}(\hat{R}) \tag{9}$$

$$R_i^* = U \ \ \mathrm{diag}([1, 1, \det(UV^\top)]) \ \ V^\top \tag{10}$$

Projecting all approximate rotations to $\mathrm{SO(d)}$ produces a valid set of pose and landmark estimates. We then use this projected solution to initialize the MAP problem in Eq. (5) and solve the problem with an iterative, local-search based NLS solver (e.g., [51]) to refine the SCORE initial estimate.

VII. Experiments

In this section we present an experimental evaluation of SCORE's performance on several RA-SLAM problems, consisting of (1) three real-world trials involving an autonomous underwater vehicle (AUV) and acoustic beacons, (2) simulated single-robot experiments with a single static beacon, and (3) simulated multi-robot scenarios with inter-robot ranging. These experiments were chosen to respectively demonstrate (1) SCORE's utility in real-world settings, (2) the challenge of RA-SLAM in even seemingly simple cases (single-robot and single-beacon), and (3) SCORE's performance in the challenging, but useful, scenario in which the relative transformation between robots must be estimated from odometry and range measurements.

A. Experimental Evaluation

We compare SCORE to other initialization techniques using the MAP estimate obtained from the given initialization. Our SCORE implementation follows the details of Section VI, using Drake [52] and the Gurobi convex solver [50].

We emphasize that the initialization strategies compared against represent either idealistic scenarios (ground truth values) or the existing state-of-the-art (odometric approaches) for the problems presented in these experiments. Whereas sophisticated initialization techniques exist for pose-graph SLAM problems, RA-SLAM lacks such techniques.

Quantitative analysis was in comparison to ground truth values and was based on absolute pose error (APE) [53].

B. Real-World AUV Experiments

The AUV data was collected using four acoustic beacons and a Bluefin Odyssey III off the coast of Italy. As GPS was not available for most of the collection, we obtained ground truth AUV positioning via MAP estimation, using GPS for the initial pose and combining odometry and acoustic ranging with known beacon locations to localize the remaining trajectory. Ground truth beacon positions were obtained by surveyed acoustic trilateration prior to the experiment. We obtained odometry measurements by estimating speed from Doppler velocity log (DVL) measurements and combining the speed estimate with heading from the INS magnetometer. We extracted range measurements via two-way time-of-flight ranging with the acoustic beacons.

1) Single-Robot Odometry Initialization (Odom): The Odom initialization set the first pose to be the identity matrix and used composition along the (noisy) odometry chain to obtain the remaining poses. The static beacon initialization uses a random sample drawn from the environment, as done in standard techniques [3], [22].

2) Evaluation of Single-Robot Experiments: As seen in Fig. 4, the SCORE initialized solution closely matched the ground truth trajectory. The APE results in Fig. 5 support this argument, showing consistently lesser error for the SCORE initialized solution in comparison to the Odom solution.

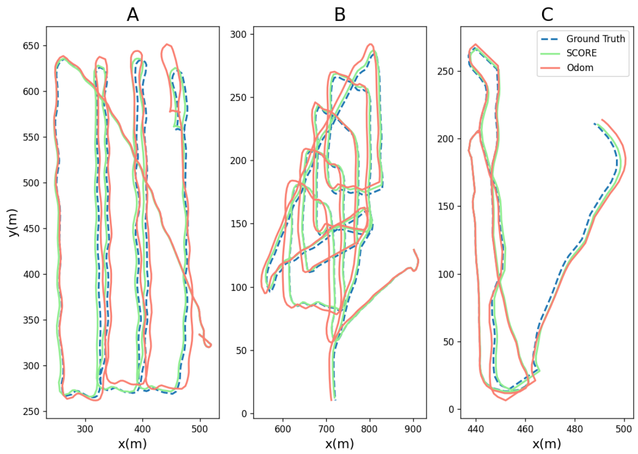

Fig. 2. The estimated trajectories from 3 separate autonomous underwater vehicle (AUV) trials. The dashed lines indicate the ground truth trajectories of the robot, green lines indicate the estimated trajectory using SCORE initialization, and orange lines indicate the estimated trajectory using odometric initialization.

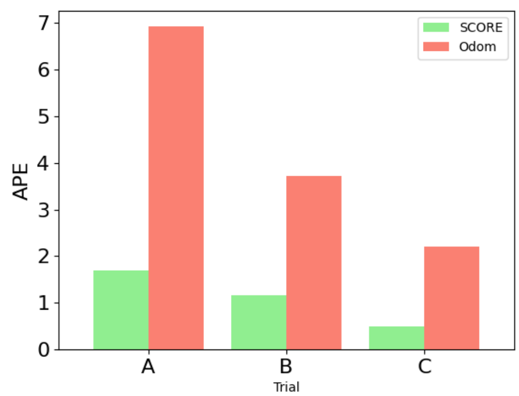

Fig. 3. The absolute pose error (APE) of the estimated trajectories in the AUV trials for both SCORE and Odom initialization strategies. Green indicates SCORE initialization, and orange indicates odometric initialization.

C. Generation of Simulations

We simulated RA-SLAM problems in which robots moved along a grid and measurements used the noise models of Section IV-B. In all simulated trials the initial position of any robots or beacons were randomly placed within the simulated environment. At each time step odometry measurements were recorded for each robot and each robot had a 50% chance of generating a range measurement. If generated, the remaining endpoint of the range measurement was chosen by uniform sampling from the set of all other robots and beacons present.

To evaluate the effect of different sensor noise levels, we tested low- and high-noise cases for both the single-robot and multi-robot scenarios. Each set of noise parameters comes from values found in previous real-world experiments. All odometry measurements were 1-meter movements, so a standard deviation of 0.01 meters in translation measurements corresponds to 1% of the distance traveled.

The low-noise setting uses $\sigma_{ij}^2 = (0.01)^2$, $\tau_{ij}^{-1} = (0.01)^2$, and $\kappa_{ij}^{-1} = 2 \times (0.002)^2$. This corresponds to standard deviations of 0.01 meters, 0.01 meters, and 0.002 radians (0.41 degrees) for distance, translation, and rotation measurements, respectively. The high-noise setting uses $\sigma_{ij}^2 = (0.5)^2$, $\tau_{ij}^{-1} = (0.05)^2$, and $\kappa_{ij}^{-1} = 2 \times (0.02)^2$.

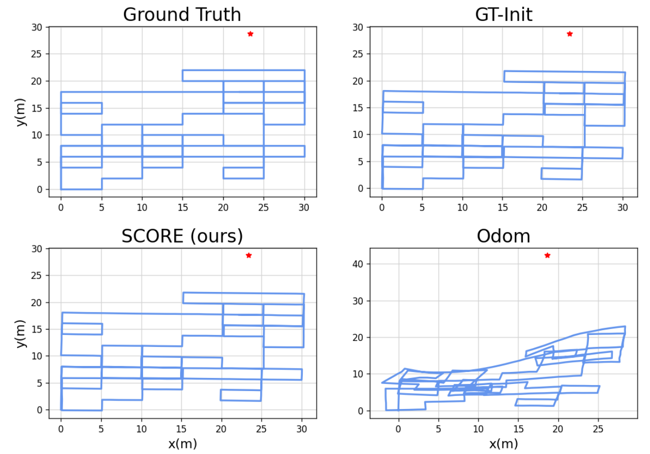

Fig. 4. The final estimated trajectories and beacon positions from a single-robot trial with low-noise parameters as defined in Section VII-C. Blue lines indicate trajectory and the red dot is the beacon position. From top-left to bottom-right: (Top-Left) The ground-truth trajectory of the robot, (Top-Right) Estimated trajectory using ground-truth initialization, (Bottom-Left) Estimated trajectory using SCORE initialization, (Bottom-Right) Estimated trajectory using odometry chain initialization (Odom).

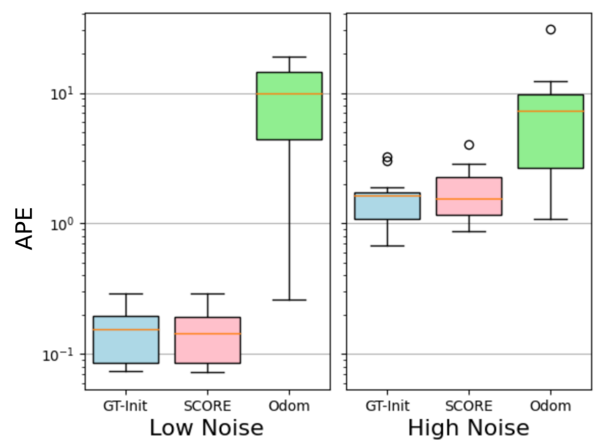

Fig. 5. Box and whisker plots of the absolute pose error (APE) of the estimated trajectory of a single robot for each initialization strategy. Each subplot represents results from 20 simulations. The left subplot refers to low noise experiments and the right subplot concerns high noise experiments.

D. Simulated Single-Robot Experiments

In this section we evaluate the single-robot experiments, which considered single robot and single static ranging beacon. For both the high and low noise cases we generated 20 trials, with each trial consisting of 400 poses. For each trial we compare the results of three different initialization strategies: SCORE, simple odometry composition, and GT-Init (in which the ground-truth is used as initialization).

1) Evaluation of Single-Robot Experiments: As seen in Fig. 4, the SCORE initialized solution appeared to exactly match the GT-Init initialized solution. This is supported by

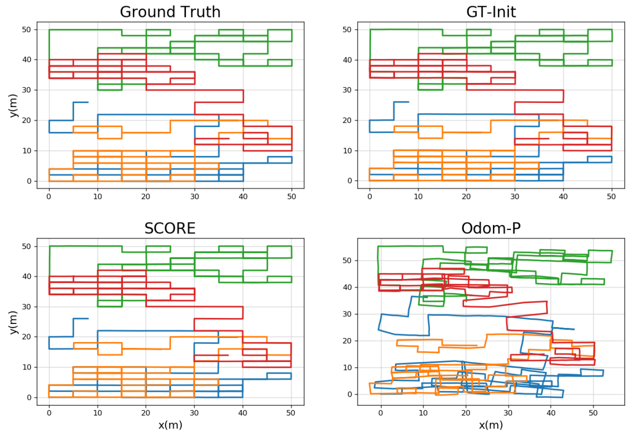

Fig. 6. A series of estimated trajectories from a multi-robot trial using the low-noise parameters defined in Section VII-C. From top-left to bottom-right: (Top-Left) Estimated trajectory from GT-Init initialization, (Bottom-Left) Estimated trajectory from SCORE initialization, (Bottom-Right) Estimated trajectory from Odom-P initialization.

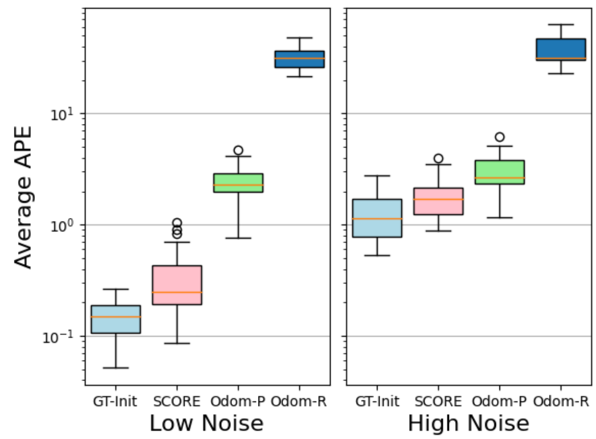

Fig. 7. Box and whisker plots of the average absolute pose error (APE) of the estimated trajectories of four robots for each initialization strategy. Each subplot represents results from 20 simulations. The left subplot refers to low noise experiments and the right subplot concerns high noise experiments.

the quantitative results in Fig. 5, which suggest that SCORE and GT-Init initialization effectively resulted in the same APE for the trials run. In contrast to SCORE and GT-Init, the Odom initialization obtained notably worse solutions (i.e. substantially higher APE and qualitatively incorrect). In Fig. 5 it can be seen that the Odom initialization resulted in an increase in APE of up to two orders of magnitude.

E. Multi-Robot Experiments

To test a scenario in which initialization is difficult, we simulated scenarios with four robots and no ranging beacons. For both the high and low noise cases we generated 20 trials with 400 poses per robot (1600 poses per trial). We compare the results of four different initialization strategies: SCORE, ground truth (GT-Init), odometry with perfect first pose (Odom-P), and odometry with random first pose (Odom-R). We describe these approaches in detail in Section VII-E.1.

1) Multi-Robot Initialization: GT-Init is an idealistic case in which the ground truth state is known *a priori* and all variables are initialized to their true values. While RA-SLAM is not needed if the true state is known, GT-Init is meant to serve as an upper bound for reasonable quality of an initialization technique.

In comparison to the GT-Init initialization, the odometry based initializations represent more realistic scenarios in which either the initial pose of each robot is known (Odom-P) or no prior information exists (Odom-R). Odom-P and Odom-R only differ in choosing the starting pose for each robot. After determining the starting pose for a given robot both Odom-P and Odom-R compose odometry measurements to initialize the remaining poses for all robots.

In Odom-P we assume a perfectly known starting pose of each robot and compose the odometry chain to initialize the other pose variables. Such a situation may occur in marine robotics, where AUVs begin with GPS before diving underwater. In contrast to Odom-P, for Odom-R we assume no prior information and choose the starting pose for each robot randomly. We note that Odom-R is a standard approach to initialization for multi-robot RA-SLAM [22], [23] and is arguably the fairest comparison for SCORE as both require no prior information.

2) Evaluation of Multi-Robot Experiments: As seen in the example multi-robot trial displayed in Fig. 6, SCORE provides reliable estimates of the trajectories of the four robots. In this trial the solution obtained from the SCORE initialization qualitatively matches the true trajectories. In addition, the refined SCORE initialization appears to match the trajectories estimated using the GT-Init initialization and provide a substantially better estimate than that provided by the Odom-P and Odom-R initializations.

These qualitative observations from Fig. 6 are supported by the quantitative analysis of the APE in Fig. 7. We observe that for the high-noise multi-robot experiments the SCORE initialization strategy resulted in similar APE to the GT-Init initialization and had notably reduced APE in comparison to Odom-P and Odom-R. For the low-noise multi-robot experiments SCORE resulted in higher APE than the GT-Init initialization approach but lesser APE than either of the odometric initialization strategies.

VIII. Discussion

In summary, the results shown in Section VII demonstrated that SCORE works in real-world settings and outperforms state-of-the-art initialization techniques for a wide range of RA-SLAM problems. We emphasize the difficulty of the multi-robot problems presented and that SCORE is capable of providing quality initializations for multi-robot RA-SLAM problems with only odometry and inter-robot range measurements.

Furthermore, in evaluating SCORE's performance on these multi-robot experiments we emphasize that Odom-R is the only other initialization technique which assumes the same amount of prior information as SCORE, i.e., no information beyond the robot's measurements. In comparing SCORE to Odom-R we observe that the SCORE initialization typically obtains an average APE roughly two orders of magnitude less

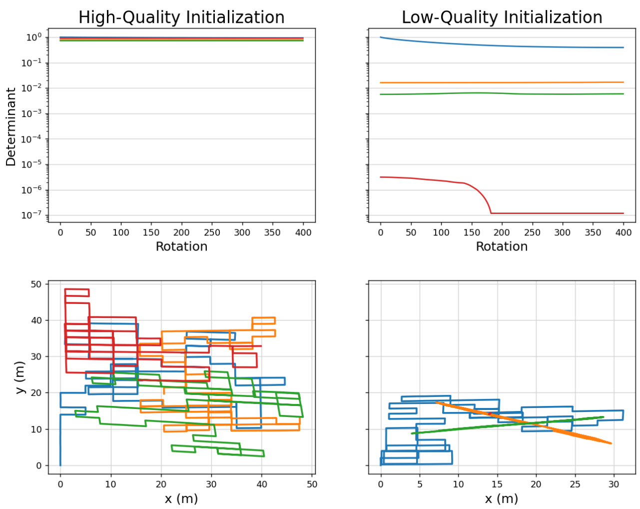

Fig. 8. Examples of low- and high-quality SCORE solutions (before refinement). The top row shows the determinants of each of the rotation matrices for four robots. Each colored curve represents a different robot's trajectory. The bottom row visualizes the corresponding SCORE initialization for each determinant plot. Visualized are: (Top-Left) The determinants for a high-quality initialization, (Bottom-Left) corresponding high-quality SCORE initialization, (Top-Right) the determinants for a low-quality initialization. (Bottom-Right) corresponding low-quality SCORE initialization. Observe how several trajectories with small determinants ($\det(R) \leq 0.05$) were 'flattened' into linear segments in the low-quality initialization.

than that obtained by Odom-R initialization. In fact, SCORE appears to obtain more accurate trajectory estimates than those generated by Odom-P, which assumes knowledge of each robot's first pose. These observations suggest that SCORE is a more effective initialization strategy than general odometry-based techniques (the current state-of-the-art initialization for these problems) even when each robot's first pose is available.

We note that several instances demonstrated substantial improvement in the estimate by refining the SCORE initialization. In particular, we observed the translation estimates were compressed (i.e., the estimated distance traveled was notably less than the true value of 1 meter). Despite this compression in the initialized translations, the final MAP estimate appeared high quality. We hypothesize that good solutions are obtained from these compressed translations because the SCORE initialization returns good initial estimates of the robot orientations and beacon locations. This would agree with known results in pose-graph SLAM [7] which, show that the quality of MAP solutions depends more on good orientation initializations than good translation initializations.

A. When Does SCORE Return Poor Initializations?

Though SCORE generally performed well as an initialization technique, we observed several multi-robot RA-SLAM instances where SCORE produced poor initializations. These poor initializations were due to the cost function tending towards a trivial solution of all zeros (i.e., $R_i = 0_{d\times d}, t_i = d_i = 0_d$). While the constraint $R_0 = I_d$ should prevent this, in practice certain scenarios demonstrate the remaining variables rapidly decaying to zero. This phenomenon is unsurprising, as it is found in similar SLAM relaxations [7], [54] and the zeros elements indeed minimizes the problem cost. Importantly, we note that these failures can be readily detected by observing the determinants of the estimated rotations, which will quickly decay to less than $5\text{e}{-2}$ (see Fig. 8). As a result, failures can be identified and alternative initializations used in this case.

We find that the experiments with failed (poor) initializations resemble the conditions identified by [44] for the SNL SOCP derived. Namely, the SOCP relaxation returns good solutions when the robot trajectories had well-constrained poses (via range measurements) within the convex hull of robot poses connected via odometry to R_0 (i.e., the trajectory of the robot with constrained first rotation). Though at present we lack formal analysis, empirical results suggest the analysis and connections to rigidity theory in [44] are closely linked to SCORE.

IX. Conclusions

We presented SCORE, a novel SOCP relaxation of the RA-SLAM problem. Through simulated and real-world experiments we demonstrated SCORE outperforms state-of-the-art initialization techniques and can obtain high-quality initializations for RA-SLAM problems. The initializations provided by SCORE resulted in up to two orders of magnitude reduction in average pose error when compared to existing odometry-based initializations. Notably, we showed that SCORE can determine good initializations for the challenging multi-robot RA-SLAM problem when using only odometry and inter-robot range measurements. Along the way, we derived a QCQP formulation of the RA-SLAM problem, which is of independent scientific interest and connects RA-SLAM to a broader body of work. Finally, we discussed empirically-derived hypotheses as to the conditions under which SCORE can be expected to return a good initialization and drew connections to existing theoretical analysis in SNL.

References

[1] E. R. Boroson, R. Hewitt, N. Ayanian, and J.-p. D. Croix, "Inter-Robot Range Measurements in Pose Graph Optimization," *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4806–4813, 2020.

[2] N. Funabiki, B. Morrell, J. Nash, and A. A. Agha-Mohammadi, "Range-Aided Pose-Graph-Based SLAM: Applications of Deployable Ranging Beacons for Unknown Environment Exploration," *IEEE Robotics and Automation Letters*, vol. 6, no. 1, pp. 48–55, 2021.

[3] P. Newman and J. Leonard, "Pure range-aided sub-sea SLAM," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, vol. 2, Sept. 2003, pp. 1921–1926.

[4] A. Bahr and J. Leonard, "Cooperative Localization for Autonomous Underwater Vehicles," in *Intl. Sym. on Experimental Robotics (ISER)*, Rio de Janeiro, Brasil, July 2006.

[5] A. Bahr, J. Leonard, and A. Martinoli, "Dynamic positioning of beacon vehicles for coop- erative underwater navigation," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Algarve, Portugal, 2012.

[6] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

[7] L. Carlone, R. Tron, K. Daniilidis, and F. Dellaert, "Initialization techniques for 3D SLAM: A survey on rotation estimation and its use in pose graph optimization," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2015-June, no. June, pp. 4597–4604, 2015.

[8] L. Carlone, D. M. Rosen, G. Calafiore, J. J. Leonard, and F. Dellaert, "Lagrangian duality in 3D SLAM: Verification techniques and optimal solutions," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Hamburg, Germany, sep 2015.

[9] M. Giamou, Z. Ma, V. Peretroukhin, and J. Kelly, "Certifiably Globally Optimal Extrinsic Calibration From Per-Sensor Egomotion," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 367–374, 2019.

[10] A. M. C. So and Y. Ye, "Theory of semidefinite programming for sensor network localization," *Mathematical Programming*, vol. 109, no. 2-3, pp. 367–384, 2007.

[11] D. M. Rosen, L. Carlone, A. S. Bandeira, and J. J. Leonard, "SE-sync: A certifiably correct algorithm for synchronization over the special euclidean group," *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 95–125, 2019.

[12] F. Dellaert, M. Kaess, *et al.*, "Factor graphs for robot perception," *Foundations and Trends® in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.

[13] E. Menegatti, A. Zanella, S. Zilli, F. Zorzi, and E. Pagello, "Range-only slam with a mobile robot and a wireless sensor networks," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 8–14, 2009.

[14] J. Djugash, S. Singh, and B. P. Grocholsky, "Modeling mobile robot motion with polar representations," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Oct. 2009.

[15] J. Djugash and S. Singh, "A robust method of localization and mapping using only range," *Springer Tracts in Advanced Robotics*, vol. 54, pp. 341–351, 2009.

[16] J. González, J. L. Blanco, C. Galindo, A. Ortiz-de Galisteo, J. A. Fernández-Madrigal, F. A. Moreno, and J. L. Martínez, "Mobile robot localization based on Ultra-Wide-Band ranging: A particle filter approach," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 496–507, 2009.

[17] J. L. Blanco, J. González, and J. A. Fernández-Madrigal, "A pure probabilistic approach to range-only SLAM," *Proceedings - IEEE International Conference on Robotics and Automation*, pp. 1436–1441, 2008.

[18] J. L. Blanco, J. A. Fernández-Madrigal, and J. González, "Efficient probabilistic range-only SLAM," *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, pp. 1017–1022, 2008.

[19] F. Herranz, A. Llamazares, E. Molinos, and M. Ocana, "A comparison of slam algorithms with range only sensors," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014, pp. 4606–4611.

[20] E. Olson, J. Leonard, and S. Teller, "Robust range-aided beacon localization," *IEEE Journal of Oceanic Engineering*, vol. 31, no. 4, pp. 949–958, Oct. 2006.

[21] B. Boots and G. J. Gordon, "A Spectral learning approach to range-only SLAM," *30th International Conference on Machine Learning, ICML 2013*, vol. 28, no. PART 1, pp. 19–26, 2013.

[22] K. Guo, Z. Qiu, W. Meng, L. Xie, and R. Teo, "Ultra-wideband based cooperative relative localization algorithm and experiments for multiple unmanned aerial vehicles in gps denied environments," *International Journal of Micro Air Vehicles*, vol. 9, no. 3, pp. 169–186, 2017.

[23] S. Li, M. Coppola, C. De Wagter, and G. C. H. E. de Croon, "An autonomous swarm of micro flying robots with range-based relative localization," 2020. [Online]. Available: http://arxiv.org/abs/2003.05853

[24] X. S. Zhou and S. I. Roumeliotis, "Robot-to-robot relative pose estimation from range measurements," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1379–1393, 2008.

[25] N. Trawny, X. S. Zhou, K. X. Zhou, and S. I. Roumeliotis, "3d relative pose estimation from distance-only measurments," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2007, pp. 1071—-1078.

[26] N. Trawny, X. S. Zhou, and S. I. Roumeliotis, "3d relative pose estimation from six distances," in *Robotics: Science and Systems (RSS)*, vol. 5, 2010, pp. 233–240.

[27] N. Trawny, X. S. Zhou, K. Zhou, and S. I. Roumeliotis, "Interrobot transformations in 3-d," *IEEE Trans. Robotics*, vol. 26, no. 2, pp. 226–243, 2010.

[28] B. Jiang, B. D. Anderson, and H. Hmam, "3-D Relative Localization of Mobile Systems Using Distance-Only Measurements via Semidefinite Optimization," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 56, no. 3, pp. 1903–1916, 2020.

[29] M. Li, G. Liang, H. Luo, H. Qian, and T. L. Lam, "Robot-to-robot relative pose estimation based on semidefinite relaxation optimization," *IEEE International Conference on Intelligent Robots and Systems*, pp. 4491–4498, 2020.

[30] M. Li, T. L. Lam, and Z. Sun, "3-D Inter-Robot Relative Localization via Semidefinite Optimization," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 081–10 088, 2022.

[31] D. Martinec and T. Pajdla, "Robust rotation and translation estimation in multiview reconstruction," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2007.

[32] V. M. Govindu, "Combining two-view constraints for motion estimation," *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 2, pp. 218–225, 2001.

[33] J. Fredriksson and C. Olsson, "Simultaneous multiple rotation averaging using lagrangian duality," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 7726 LNCS, no. PART 3, pp. 245–258, 2013.

[34] R. Tron and R. Vidal, "Distributed 3-D localization of camera sensor networks from 2-D image measurements," *IEEE Transactions on Automatic Control*, vol. 59, no. 12, pp. 3325–3340, 2014.

[35] L. Carlone, R. Aragues, J. A. Castellanos, and B. Bona, "A fast and accurate approximation for planar pose graph optimization," *International Journal of Robotics Research*, vol. 33, no. 7, pp. 965–987, 2014.

[36] J. Briales and J. Gonzalez-Jimenez, "Cartan-sync: Fast and global se(d)-synchronization," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2127–2134, 2017.

[37] R. Tron, D. M. Rosen, and L. Carlone, "On the inclusion of determinant constraints in lagrangian duality for 3d slam," *Proc. Robotics: Science and Systems (RSS), Workshop "The Problem of Mobile Sensors"*, vol. 2, no. 1, 2015.

[38] T. Fan, H. Wang, M. Rubenstein, and T. Murphey, "CPL-SLAM: Efficient and Certifiably Correct Planar Graph-Based SLAM Using the Complex Number Representation," *IEEE Transactions on Robotics*, vol. 36, pp. 1719—-1737, 2020. [Online]. Available: http://arxiv.org/abs/2007.06708

[39] J. Briales, L. Kneip, and J. Gonzalez-Jimenez, "A Certifiably Globally Optimal Solution to the Non-minimal Relative Pose Problem," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2018, pp. 145–154.

[40] M. Garcia-Salguero, J. Briales, and J. Gonzalez-Jimenez, "Certifiable relative pose estimation," *Image and Vision Computing*, vol. 109, 2021.

[41] M. Pacholska, F. Dümbgen, and A. Scholefield, "Relax and Recover: Guaranteed Range-Only Continuous Localization," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2248–2255, 2020.

[42] P. Biswas, T. C. Liang, K. C. Toh, Y. Ye, and T. C. Wang, "Semidefinite programming approaches for sensor network localization with noisy distance measurements," *IEEE Transactions on Automation Science and Engineering*, vol. 3, no. 4, pp. 360–371, 2006.

[43] D. Shamsi, N. Taheri, Z. Zhu, and Y. Ye, "Conditions for correct sensor network localization using sdp relaxation," in *Discrete geometry and optimization*. Springer, 2013, pp. 279–301.

[44] P. Tseng, "Second-order cone programming relaxation of sensor network localization," *SIAM Journal on Optimization*, vol. 18, no. 1, pp. 156–185, 2007.

[45] G. Naddafzadeh-Shirazi, M. B. Shenouda, and L. Lampe, "Second order cone programming for sensor network localization with anchor position uncertainty," *IEEE Transactions on Wireless Communications*, vol. 13, no. 2, pp. 749–763, 2014.

[46] L. Doherty, K. S. Pister, and L. El Ghaoui, "Convex position estimation in wireless sensor networks," *Proceedings - IEEE INFOCOM*, vol. 3, pp. 1655–1663, 2001.

[47] S. Srirangarajan, A. Tewfik, and Z. Q. Luo, "Distributed sensor network localization using SOCP relaxation," *IEEE Transactions on Wireless Communications*, vol. 7, no. 12, pp. 4886–4895, 2008.

[48] A. Chiuso, G. Picci, and S. Soatto, "Wide-sense estimation on the special orthogonal group," *Communications in Information and Systems*, vol. 8, no. 3, pp. 185–200, 2008.

[49] F. Alizadeh and D. Goldfarb, "Second-order cone programming," *Mathematical Programming, Series B*, vol. 95, no. 1, pp. 3–51, 2003.

[50] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2021. [Online]. Available: https://www.gurobi.com

[51] F. Dellaert, "Factor graphs and gtsam: A hands-on introduction," Georgia Institute of Technology, Tech. Rep., 2012.

[52] R. Tedrake and the Drake Development Team, "Drake: Model-based design and verification for robotics," 2019. [Online]. Available: https://drake.mit.edu

[53] M. Grupp, "evo: Python package for the evaluation of odometry and slam." https://github.com/MichaelGrupp/evo, 2017.

[54] D. Rosen, C. DuHadway, and J. J. Leonard, "A convex relaxation for approximate global optimization in simultaneous localization and mapping," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, May 2015.