Lagrangian Duality in 3D SLAM: Verification Techniques and Optimal Solutions

Original

Lagrangian Duality in 3D SLAM: Verification Techniques and Optimal Solutions / Carlone, Luca; Rosen, David M.; Calafiore, Giuseppe Carlo; Leonard, John J.; Dellaert, Frank. - STAMPA. - (2015). (Intervento presentato al convegno Int. Conf. on Intelligent Robots and Systems (IROS) tenutosi a Hamburg nel September 2015).

Availability:

This version is available at: 11583/2619534 since: 2015-10-08T13:17:06Z

Publisher:

Published

DOI:

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Lagrangian Duality in 3D SLAM: Verification Techniques and Optimal Solutions

Luca Carlone, David M. Rosen, Giuseppe Calafiore, John J. Leonard, Frank Dellaert

Abstract—State-of-the-art techniques for simultaneous localization and mapping (SLAM) employ iterative nonlinear optimization methods to compute an estimate for robot poses. While these techniques often work well in practice, they do not provide guarantees on the quality of the estimate. This paper shows that *Lagrangian duality* is a powerful tool to assess the quality of a given candidate solution. Our contribution is threefold. First, we discuss a revised formulation of the SLAM inference problem. We show that this formulation is probabilistically grounded and has the advantage of leading to an optimization problem with quadratic objective. The second contribution is the derivation of the corresponding *Lagrangian dual problem*. The SLAM dual problem is a (convex) *semidefinite program*, which can be solved reliably and globally by off-the-shelf solvers. The third contribution is to discuss the relation between the original SLAM problem and its dual. We show that from the dual problem, one can evaluate the quality (i.e., the suboptimality gap) of a candidate SLAM solution, and ultimately provide a certificate of optimality. Moreover, when the duality gap is zero, one can compute a guaranteed optimal SLAM solution from the dual problem, circumventing non-convex optimization. We present extensive (real and simulated) experiments supporting our claims and discuss practical relevance and open problems.

I. INTRODUCTION

Simultaneous localization and mapping (SLAM) is an enabling technology for many applications, including service and industrial robotics, autonomous driving, search and rescue, planetary exploration, and augmented reality.

The last decade has witnessed several groundbreaking results in SLAM, and state-of-the-art approaches are now transitioning from academic research to industrial applications. Standard techniques compute an estimate (e.g., for robot poses) by minimizing a nonlinear cost function, whose global minimum is the *maximum likelihood estimate* (or *maximum a posteriori* estimate in presence of priors). The optimization problem underlying SLAM is commonly solved using iterative nonlinear optimization methods, e.g. the Gauss-Newton method [1], [2], [3], the gradient method [4], [5], trust region methods [6], or ad-hoc approximations [7], [8].

Despite the success of state-of-the-art techniques, some practical and theoretical problems remain open. While iterative approaches are observed to work well in many problem instances, they cannot guarantee the correctness (global optimality) of the estimates that they compute. This is due to the fact that the optimization problem is non-convex, hence iterative

Fig. 1. This paper exploits Lagrangian duality to evaluate the quality of a SLAM solution. For instance, our verification techniques are able to certify the optimality of the solutions in the first row of the figure, while they reject as suboptimal the estimates shown in the bottom row.

techniques may be trapped in local minima¹ (Fig. 1), which correspond to wrong estimates. Recent work [9], [10] shows that iterative techniques fail to converge to a correct estimate even in fairly simple (real and simulated) 3D problems. Recent research efforts have addressed the issue of global convergence from several angles. Olson *et al.* [4], Grisetti *et al.* [5], Rosen *et al.* [6], and Tron *et al.* [11] study iterative techniques with larger basins of convergence. Carlone *et al.* [8], [12], and Rosen *et al.* [10] propose initialization techniques to bootstrap iterative optimization. Huang *et al.* [13], Wang *et al.* [14], Carlone [15], and Khosoussi *et al.* [16] investigate the factors influencing local convergence and the quality of the optimal SLAM solution. While these techniques provide remarkable insights into the problem, and working solutions to improve convergence, none of them can guarantee the recovery of a globally optimal solution to SLAM.

The motivation behind this work is that the transition of SLAM from research topic to industrial technology requires techniques with *guaranteed performance*. In autonomous vehicle applications, failure to produce a correct SLAM solution may put passengers' lives at risk. In other applications, SLAM failures can possibly cascade into path planning failures (if the plan is computed using a wrong map), and this may prevent the reliable operation of mobile robots.

Therefore, in this paper we address the following question: **Verification Problem:** *Given a candidate SLAM estimate (e.g., returned by a state-of-the-art iterative solver), is it possible to evaluate the quality of this estimate (e.g., its suboptimality gap), possibly certifying its optimality?*

To address this problem, we introduce a powerful tool,

¹We use the term “local minimum” to denote a stationary point of the cost which does not attain the optimal objective.

L. Carlone and F. Dellaert are with the Georgia Institute of Technology, Atlanta, GA, USA, luca.carlone@gatech.edu, frank@cc.gatech.edu.

D.M. Rosen and J.J. Leonard are with the Massachusetts Institute of Technology, Cambridge, MA, USA {dmrosen,jleonard}@mit.edu.

G. Calafiore is with the Politecnico di Torino, Torino, Italy, giuseppe.calafiore@polito.it.

Lagrangian duality, borrowing the corresponding theory from the optimization community. Duality was first applied to 2D SLAM by Carlone and Dellaert [9]. In this work, we provide a nontrivial extension of [9] to 3D SLAM.

This paper contains three main contributions. The first contribution is a revised formulation of the SLAM inference problem. Our formulation is a probabilistically grounded maximum likelihood (ML) estimator, and has the advantage of leading to an optimization problem with quadratic objective, which facilitates the derivation of the dual problem. Our revised SLAM formulation is presented in Section II.

The second contribution is the derivation of the *Lagrangian dual problem*. A key step towards this goal consists of rewriting the 3D SLAM problem as a quadratic optimization problem with quadratic equality constraints. Intuitively, the constraints impose that the pose estimates are members of $SE(3)$, while the objective minimizes the mismatch w.r.t. the measurements. The dual SLAM problem, introduced in Section III, is a (convex) *semidefinite program* (SDP), and can be solved globally by off-the-shelf solvers.

The third contribution is to provide *verification techniques* to assess the quality of a given SLAM solution, leveraging the relation between the standard SLAM problem and its dual. We show that solving the dual problem allows us to bound the sub-optimality gap of a given candidate solution, hence we are able to quantify how far the candidate solution is from being optimal. Since current SDP solvers do not scale well to large problems, we also propose a second verification technique that does not require solving the SDP. As a by-product of our derivation, we show that, when the duality gap is zero, we can compute an optimal SLAM solution directly from the dual problem. Our verification techniques, presented in Section IV, can be seamlessly integrated in standard SLAM pipelines. Experimental evidence (Section V) confirms that these techniques enable the certification of globally optimal solutions in both real and simulated experiments. Extra results and visualizations are given in the supplemental material [17].

II. 3D POSE GRAPH OPTIMIZATION REVISITED

We consider the *pose graph optimization* (PGO) formulation of the SLAM problem. PGO computes the *maximum likelihood* estimate for n poses $\mathbf{x}_1, \dots, \mathbf{x}_n$, given m relative pose measurements $\bar{\mathbf{x}}_{ij}$ between pairs of poses i and j . In a 3D setup, both the unknown poses and the measurements are quantities in $SE(3) \doteq \{(\mathbf{R}, \mathbf{t}) : \mathbf{R} \in SO(3), \mathbf{t} \in \mathbb{R}^3\}$. We use the notation $\mathbf{x}_i = (\mathbf{R}_i, \mathbf{t}_i)$ and $\bar{\mathbf{x}}_{ij} = (\bar{\mathbf{R}}_{ij}, \bar{\mathbf{t}}_{ij})$ to make explicit the rotation and the translation of each pose. PGO can be visualized as a directed graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, in which we associate a node $i \in \mathcal{V} = \{1, \dots, n\}$ to each pose $\mathbf{x}_i$ and an edge $(i, j) \in \mathcal{E}$ to each relative measurement $\bar{\mathbf{x}}_{ij}$.

In this section we propose a revised PGO formulation. The key difference w.r.t. related work is the use of the *chordal distance* to quantify the rotation errors (more details in Section II-B). To lay the groundwork for this formulation, we begin with a generative model for our measurements, and then derive the corresponding ML estimator.

A. Generative Noise Model

We assume the following generative model for the relative pose measurements $(\bar{\mathbf{R}}_{ij}, \bar{\mathbf{t}}_{ij})^2$:

$$\begin{aligned} \bar{\mathbf{t}}_{ij} &= \mathbf{R}_i^\top (\mathbf{t}_j - \mathbf{t}_i) + \mathbf{t}_\epsilon & \mathbf{t}_\epsilon &\sim \text{gaussian}(\mathbf{0}_3, \omega_t^2 \mathbf{I}_3) \\ \bar{\mathbf{R}}_{ij} &= \mathbf{R}_i^\top \mathbf{R}_j \mathbf{R}_\epsilon & \mathbf{R}_\epsilon &\sim \text{vonMises}(\mathbf{I}_3, \omega_R^2) \end{aligned} \quad (1)$$

where “ $\text{gaussian}(\boldsymbol{\mu}, \boldsymbol{\Omega})$ ” denotes a Gaussian distribution with mean $\boldsymbol{\mu}$ and *information* matrix $\boldsymbol{\Omega}$, while “ $\text{vonMises}(\mathbf{S}, \kappa)$ ” denotes the isotropic von Mises-Fisher distribution, with mean $\mathbf{S} \in SO(3)$ and *concentration parameter* κ . The key difference w.r.t. to measurement models in other PGO formulations lies in the use of the von Mises-Fisher distribution as the model for the rotational measurements errors $\mathbf{R}_\epsilon \in SO(3)$.

The isotropic *von Mises-Fisher* (or *Langevin*) [18] distribution on $SO(n)$ with mean $\mathbf{S} \in SO(n)$ and concentration parameter $\kappa \geq 0$ can be written explicitly as:

$$\mathbb{P}(\mathbf{R}_\epsilon) = \frac{1}{c_n(\kappa)} \exp(\kappa \text{tr}(\mathbf{S}^\top \mathbf{R}_\epsilon)), \quad (2)$$

where $\text{tr}(\cdot)$ is the matrix trace and $c_n(\kappa)$ is a normalization term. Closed-form expressions for $c_n(\kappa)$ are given in [18]; these are inconsequential for our derivation. For $\kappa \rightarrow 0$, the distribution tends to the uniform distribution over $SO(n)$. For $\kappa \rightarrow \infty$, $\mathbf{R}_\epsilon = \mathbf{S}$ with probability one. Roughly speaking, one may think at κ in terms of *information* content.

We are now ready to introduce the maximum likelihood estimator for the poses, given the measurement model (1).

B. Maximum Likelihood Estimator

The ML estimate corresponds to the set of poses maximizing the likelihood of the measurements, or, equivalently, minimizing the negative log-likelihood:

$$f_{\text{ML}}^* = \min_{\{\mathbf{x}_i \in SE(3)\}} \sum_{(i,j) \in \mathcal{E}} -\log \mathcal{L}(\bar{\mathbf{R}}_{ij} | \mathbf{x}) - \log \mathcal{L}(\bar{\mathbf{t}}_{ij} | \mathbf{x}). \quad (3)$$

The negative log-likelihood of the Cartesian measurements can be easily computed from the Gaussian distribution:

$$-\log \mathcal{L}(\bar{\mathbf{t}}_{ij} | \mathbf{x}) = \omega_t^2 \|\mathbf{t}_j - \mathbf{t}_i - \mathbf{R}_i \bar{\mathbf{t}}_{ij}\|^2 + \text{const.} \quad (4)$$

Using (2), the negative log-likelihood for $\bar{\mathbf{R}}_{ij}$ is:

$$\begin{aligned} -\log \mathcal{L}(\bar{\mathbf{R}}_{ij} | \mathbf{x}) &= -\omega_R^2 \text{tr}(\mathbf{R}_i^\top \mathbf{R}_j \bar{\mathbf{R}}_{ij}) + \text{const.} \\ &= \frac{\omega_R^2}{2} \|\mathbf{R}_j - \mathbf{R}_i \bar{\mathbf{R}}_{ij}\|_F^2 + \text{const.} \end{aligned} \quad (5)$$

where $\|\cdot\|_F$ is the Frobenius matrix norm (sum of the squares of the entries), and we used $\|\mathbf{S} - \mathbf{R}\|_F^2 = \text{tr}((\mathbf{S} - \mathbf{R})(\mathbf{S} - \mathbf{R})^\top)$. The norm $\|\mathbf{S} - \mathbf{R}\|_F^2$ is usually referred to as the *chordal distance* between two rotations $\mathbf{S}$ and $\mathbf{R}$ [19].

Plugging (4) and (5) back into (3) we obtain our ML estimator:

$$\begin{aligned} f_{\text{ML}}^* &= \min_{\substack{\{\mathbf{t}_i \in \mathbb{R}^3\} \\ \{\mathbf{R}_i \in SO(3)\}}} \sum_{(i,j) \in \mathcal{E}} \omega_t^2 \|\mathbf{t}_j - \mathbf{t}_i - \mathbf{R}_i \bar{\mathbf{t}}_{ij}\|^2 \\ &\quad + \frac{\omega_R^2}{2} \|\mathbf{R}_j - \mathbf{R}_i \bar{\mathbf{R}}_{ij}\|_F^2 \end{aligned} \quad (\text{PGO}) \quad (6)$$

²To keep notation simple, we consider measurements with the same distribution. The extension to heterogeneous ω_R^2 and ω_t^2 is trivial.

The main difference between (6) and formulations in related work is the use of the chordal distance (related work instead uses the *geodesic distance* $\|\text{Log}(\bar{\mathbf{R}}_{ij}^\top \mathbf{R}_i^\top \mathbf{R}_j)\|^2$). References [9], [19] show that for small residual errors $\frac{1}{2} \|\mathbf{R}_j - \mathbf{R}_i \bar{\mathbf{R}}_{ij}\|_F^2 \approx \|\text{Log}(\bar{\mathbf{R}}_{ij}^\top \mathbf{R}_i^\top \mathbf{R}_j)\|^2$, making the formulations equivalent from a practical standpoint.

The advantage of the formulation (6) is that it has a quadratic objective function. This facilitates the derivation of the Lagrangian dual problem, as shown in the next section.

III. LAGRANGIAN DUALITY IN 3D PGO

The main goal of this paper is to provide tools to check if a candidate SLAM solution $\hat{\mathbf{x}}$ is globally optimal. If we knew the optimal cost f_{ML}^* this would be easy: calling $f_{\text{ML}}(\cdot)$ the objective function of (6), if $f_{\text{ML}}(\hat{\mathbf{x}}) = f_{\text{ML}}^*$ then $\hat{\mathbf{x}}$ is optimal. Unfortunately, f_{ML}^* is unknown. Our contribution is to show that we can compute close proxies of f_{ML}^* using duality theory. To make the derivation easier, we first rewrite the problem as a quadratic problem with equality constraints (Section III-A), and then derive the dual (Section III-B).

A. Quadratic Problem with Quadratic Equality Constrains

In this section, we rewrite (6) in order to (i) have vector variables (the rotations $\mathbf{R}_i$ are matrices), and (ii) formulate the constraints $\mathbf{R}_i \in \text{SO}(3)$ as quadratic equality constraints.

We define $\mathbf{r}_i \in \mathbb{R}^9$ as the vectorized version of $\mathbf{R}_i$: $\mathbf{r}_i \doteq [\mathbf{R}_i^{(1)} \ \mathbf{R}_i^{(2)} \ \mathbf{R}_i^{(3)}]^\top$, where $\mathbf{R}_i^{(k)}$ is the k th row of $\mathbf{R}_i$. We use the shorthand $\mathbf{r}_i = \text{rows}(\mathbf{R}_i)$ to obtain the vector representation $\mathbf{r}_i$ of a 3×3 matrix $\mathbf{R}_i$. Using this parametrization, each summand in the objective in (6) becomes (using that $\|\mathbf{R}\|_F = \|\mathbf{R}^\top\|_F$ in the first expression):

$$\begin{aligned} & \omega_t^2 \|\mathbf{t}_j - \mathbf{t}_i - \mathbf{R}_i \bar{\mathbf{t}}_{ij}\|^2 + \frac{\omega_R^2}{2} \|\mathbf{R}_j^\top - \bar{\mathbf{R}}_{ij}^\top \mathbf{R}_i^\top\|_F^2 \\ = & \omega_t^2 \|\mathbf{t}_j - \mathbf{t}_i - \mathbf{T}_{ij} \mathbf{r}_i\|^2 + \frac{\omega_R^2}{2} \|\mathbf{r}_j - \mathbf{Q}_{ij} \mathbf{r}_i\|^2 \end{aligned} \quad (7)$$

where $\mathbf{T}_{ij} \doteq \mathbf{I}_3 \otimes \bar{\mathbf{t}}_{ij}^\top \in \mathbb{R}^{3 \times 9}$, $\mathbf{Q}_{ij} \doteq \mathbf{I}_3 \otimes \bar{\mathbf{R}}_{ij}^\top \in \mathbb{R}^{9 \times 9}$, and $\otimes$ is the Kronecker product.

We cannot choose arbitrary vectors $\mathbf{r}_i \in \mathbb{R}^9$, but have to limit ourself to choices of $\mathbf{r}_i$ that produce meaningful rows of a rotation matrix $\mathbf{R}_i \in \text{SO}(3)$. The rotation group $\text{SO}(3)$ is defined as $\text{SO}(3) \doteq \{\mathbf{R} \in \mathbb{R}^{3 \times 3} : \mathbf{R}^\top \mathbf{R} = \mathbf{I}_3, \det(\mathbf{R}) = 1\}$, which, written in terms of the rows of $\mathbf{R}_i$, becomes:

$$\begin{aligned} \mathbf{R}_i^\top \mathbf{R}_i = \mathbf{I}_3 & \Leftrightarrow (\mathbf{R}_i^{(u)})^\top \mathbf{R}_i^{(v)} = \begin{cases} 1 & \text{if } u = v, \\ 0 & \text{if } u \neq v, \end{cases} \quad u, v = 1, 2, 3 \\ \det(\mathbf{R}_i) = 1 & \Leftrightarrow \mathbf{R}_i^{(1)} \times \mathbf{R}_i^{(2)} = \mathbf{R}_i^{(3)} \end{aligned} \quad (8)$$

where $\times$ is the cross product. In other words, the rows of a rotation matrix have to be orthonormal, and have to satisfy the right-hand rule. To derive the dual problem we relax the second condition ($\det(\mathbf{R}) = 1$), which amounts to performing estimation in $\text{O}(3)$ rather than $\text{SO}(3)$ (i.e., resulting matrices can have determinant $\det(\mathbf{R}) = \pm 1$). Then in Proposition 4 we show how to reconcile our verification techniques to work directly on the original PGO problem (6).

Using (7) and (8) and relaxing the determinant constraints, we rewrite the PGO problem (6) as:

$$\begin{aligned} f^* = & \min_{\{\mathbf{r}_i, \mathbf{t}_i\}_{(i,j) \in \mathcal{E}}} \sum \omega_t^2 \|\mathbf{t}_j - \mathbf{t}_i - \mathbf{T}_{ij} \mathbf{r}_i\|^2 + \frac{\omega_R^2}{2} \|\mathbf{r}_j - \mathbf{Q}_{ij} \mathbf{r}_i\|^2 \\ \text{subject to} & \left. \begin{aligned} \mathbf{r}_i^\top \mathbf{E}_{uv} \mathbf{r}_i &= 1, & u = v \\ \mathbf{r}_i^\top \mathbf{E}_{uv} \mathbf{r}_i &= 0, & u \neq v \end{aligned} \right\} \quad u, v = 1, 2, 3 \\ & \quad i = 1, \dots, n \end{aligned} \quad (9)$$

where $\mathbf{E}_{uv}$ is a 9×9 selection matrix composed of 3×3 blocks that are zero everywhere except the 3×3 block in position (u, v) , which is the identity matrix. The matrices $\mathbf{E}_{uv}$ are built such that $\mathbf{r}_i^\top \mathbf{E}_{uv} \mathbf{r}_i = (\mathbf{R}_i^{(u)})^\top \mathbf{R}_i^{(v)}$, hence the constraints in (9) correspond to the orthonormality constraints in (8).

In order to write (9) is a more compact matrix notation, we define the vector $\check{\mathbf{x}} = [\mathbf{t}_1^\top, \dots, \mathbf{t}_n^\top, \mathbf{r}_1^\top, \dots, \mathbf{r}_n^\top]^\top \in \mathbb{R}^{12n}$. Using this notation, (9) becomes:

$$\begin{aligned} f^* = & \min_{\check{\mathbf{x}}} \|\check{\mathbf{A}} \check{\mathbf{x}}\|^2 \\ \text{subject to} & \left. \begin{aligned} \check{\mathbf{x}}^\top \check{\mathbf{E}}_{iuv} \check{\mathbf{x}} &= 1, & u = v \\ \check{\mathbf{x}}^\top \check{\mathbf{E}}_{iuv} \check{\mathbf{x}} &= 0, & u \neq v \end{aligned} \right\} \quad u, v = 1, 2, 3 \\ & \quad i = 1, \dots, n \end{aligned} \quad (10)$$

where the matrices $\check{\mathbf{A}}$ and $\check{\mathbf{E}}_{iuv}$ are obtained by stacking the coefficient matrices in (9), with suitable zero blocks for padding; ω_t^2 and ω_R^2 are included in the definition of $\check{\mathbf{A}}$.

Finally, since absolute poses are not observable from relative measurements, we fix a pose to be our reference frame. Without loss of generality we fix the pose of the first node to the identity pose ($\mathbf{t}_1 = \mathbf{0}_3$ and $\mathbf{R}_1 = \mathbf{I}_3$, or, equivalently $\mathbf{r}_1 = \text{rows}(\mathbf{I}_3)$). This process is usually called *anchoring*. Fixing the first pose modifies (10) as follows:

$$\begin{aligned} f^* = & \min_{\mathbf{x}} \|\mathbf{A} \mathbf{x} - \mathbf{b}\|^2 \\ \text{subject to} & \left. \begin{aligned} \mathbf{x}^\top \mathbf{E}_{iuv} \mathbf{x} &= 1, & u = v \\ \mathbf{x}^\top \mathbf{E}_{iuv} \mathbf{x} &= 0, & u \neq v \end{aligned} \right\} \quad u, v = 1, 2, 3 \\ & \quad i = 1, \dots, n-1 \end{aligned} \quad (11)$$

where $\mathbf{x} \in \mathbb{R}^{12(n-1)}$ is obtained by removing the first pose from $\check{\mathbf{x}}$, $\mathbf{A}$ is obtained by removing from $\check{\mathbf{A}}$ the columns corresponding to the first pose, and $\mathbf{b}$ is the known right-hand-side arising from anchoring; $\mathbf{E}_{iuv}$ are the same as $\check{\mathbf{E}}_{iuv}$ but without the rows and columns corresponding to the first pose.

We conclude this section by transforming (11) into an equivalent problem with *homogeneous* objective (i.e., without constant terms in the squared cost). For this purpose, we note that solving (11) is the same as solving:

$$\begin{aligned} f^* = & \min_{\mathbf{x}, y} \|\mathbf{A} \mathbf{x} - \mathbf{b} y\|^2 \quad (\text{primal problem}) \\ \text{subject to} & \left. \begin{aligned} \mathbf{x}^\top \mathbf{E}_{iuv} \mathbf{x} &= 1, & u = v \\ \mathbf{x}^\top \mathbf{E}_{iuv} \mathbf{x} &= 0, & u \neq v \end{aligned} \right\} \quad u, v = 1, 2, 3 \\ & \quad i = 1, \dots, n-1 \\ & \quad y^2 = 1 \end{aligned} \quad (12)$$

meaning that the two problems have the same optimal objective, and the corresponding solutions can be mapped to each other. Intuitively, if the solution of (12) is $[\mathbf{x}_H^* \ 1]$, then $\mathbf{x}^* = \mathbf{x}_H^*$ is also optimal for (11), while if the solution is $[\mathbf{x}_H^* \ -1]$, then $\mathbf{x}^* = -\mathbf{x}_H^*$ will be optimal for (11). The inclusion of the *slack* variable y is often referred to as *homogenization*. We refer to (12) as the *primal* problem.

B. The dual problem

In this section we apply Lagrangian duality to the primal problem (12), borrowing the corresponding theory from the optimization community [20], [21]. We begin by recalling basic properties and notions about duality theory and then we tailor these concepts to our SLAM problem.

The key insight of duality is that for every constrained optimization problem of the form:

$$\begin{aligned} f^* = \quad & \min_{\mathbf{x}} \quad f(\mathbf{x}) \\ \text{subject to} \quad & c_i(\mathbf{x}) = 0 \quad \forall i \in \mathcal{C} \end{aligned} \quad (13)$$

(where $\mathcal{C}$ is a set indexing the constraints $c_i(\mathbf{x})$), there is an associated *unconstrained* optimization problem:

$$d^* = \max_{\boldsymbol{\lambda}} \left(\overbrace{\inf_{\mathbf{x}} f(\mathbf{x}) + \sum_{i \in \mathcal{C}} \lambda_i c_i(\mathbf{x})}^{d(\boldsymbol{\lambda})} \right) \quad (14)$$

called the *dual problem*. The scalar variables λ_i appearing in (14) are called *Lagrange multipliers* or *dual variables*, and the function $d(\boldsymbol{\lambda})$ is called the *dual function*; $\boldsymbol{\lambda}$ is a vector stacking all dual variables. With reference to the dual problem (14), problem (13) is referred to as the *primal* problem. Intuitively, the minimization (“inf”) in $d(\boldsymbol{\lambda})$ can be understood as a relaxation of the original problem (13) in which the constraints are transformed into penalty terms in the objective, whose “importance” is controlled by $\boldsymbol{\lambda}$; hence, the maximization (w.r.t., $\boldsymbol{\lambda}$) tries to make this relaxation as tight as possible.

The dual problem (14) has two important properties. First, since the dual function $d(\boldsymbol{\lambda})$ is the pointwise infimum of a family of affine functions of $\boldsymbol{\lambda}$, it is *always* concave, and therefore the dual *maximization* problem (14) is a convex program [20, Sec. 5.2]. Its convexity guarantees that the dual problem can always be solved *globally optimally* using *local* search techniques. Second, given any feasible $\mathbf{x}$ for (13) (for which $c_i(\mathbf{x}) = 0$), the definition of $d(\boldsymbol{\lambda})$ in (14) shows that $d(\boldsymbol{\lambda}) \leq f(\mathbf{x})$ for any choice of $\boldsymbol{\lambda}$. In particular, this must also hold at the optima $\mathbf{x}^*$ and $\boldsymbol{\lambda}^*$ in (13) and (14), so that:

$$d^* \leq f^*. \quad (15)$$

The inequality (15) is referred to as *weak (Lagrangian) duality*, and it enables us to lower-bound the optimal value f^* of the (possibly very difficult, nonconvex) primal problem (13) using the optimal value d^* of the (convex) dual problem (14). For some problems, the inequality (15) is tight (i.e. $d^* = f^*$), for which we say that *strong (Lagrangian) duality* holds. The quantity $f^* - d^* \geq 0$ is called the *duality gap*.

Using weak duality, it is easy to show that, given a primal feasible point $\hat{\mathbf{x}}$ and a dual point $\hat{\boldsymbol{\lambda}}$, the following chain of inequality holds:

$$d(\hat{\boldsymbol{\lambda}}) \leq d(\boldsymbol{\lambda}^*) \doteq d^* \leq f^* \doteq f(\mathbf{x}^*) \leq f(\hat{\mathbf{x}}) \quad (16)$$

where the first inequality stems from the fact that $\boldsymbol{\lambda}^*$ attains the maximum over all $\boldsymbol{\lambda}$, and the last follows from the optimality of f^* (which is the global minimum among all feasible $\mathbf{x}$).

Therefore, in this work we exploit a simple idea: given a candidate solution $\hat{\mathbf{x}}$, if we are able to find a $\hat{\boldsymbol{\lambda}}$, for which

$d(\hat{\boldsymbol{\lambda}}) = f(\hat{\mathbf{x}})$, then the chain of inequalities (16) becomes tight ($d(\hat{\boldsymbol{\lambda}}) = d^* = f^* = f(\hat{\mathbf{x}})$), which implies that $\hat{\mathbf{x}}$ is an optimal solution. Equation (16) thus provides a means of *certifying* the global optimality of a candidate solution $\hat{\mathbf{x}}$ for (13) and enables our derivation of *algorithmic* approaches for certifying the correctness of SLAM solutions.

We are now ready to apply duality to our primal problem (12). From (12) and (14), the dual function is:

$$\begin{aligned} d(\boldsymbol{\lambda}) = \quad & \inf_{\mathbf{x}, \mathbf{y}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\mathbf{y}\|^2 + \sum_{i=1}^{n-1} \left[\sum_{u=1,2,3} \lambda_{iuv} (1 - \mathbf{x}^\top \mathbf{E}_{iuv} \mathbf{x}) \right. \\ & \left. + \sum_{\substack{u,v=1,2,3 \\ u \neq v}} \lambda_{iuv} (-\mathbf{x}^\top \mathbf{E}_{iuv} \mathbf{x}) \right] + \lambda_y (1 - \mathbf{y}^2), \end{aligned} \quad (17)$$

where $\boldsymbol{\lambda}$ is the vector of Lagrange multipliers λ_{iuv} and λ_y , associated with the orthonormality and homogeneity constraints in (12), respectively. We observe that the quadratic terms in (17) can be written more compactly as:

$$\begin{aligned} & \|\mathbf{A}\mathbf{x} - \mathbf{b}\mathbf{y}\|^2 - \mathbf{x}^\top \left[\sum_{i=1}^{n-1} \sum_{u,v=1}^3 \lambda_{iuv} \mathbf{E}_{iuv} \right] \mathbf{x} - \lambda_y \mathbf{y}^2 \\ = \quad & \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}^\top \begin{bmatrix} \mathbf{H}(\boldsymbol{\lambda}) & -\mathbf{A}^\top \mathbf{b} \\ -\mathbf{b}^\top \mathbf{A} & \mathbf{b}^\top \mathbf{b} - \lambda_y \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}, \end{aligned} \quad (18)$$

where

$$\mathbf{H}(\boldsymbol{\lambda}) \doteq \mathbf{A}^\top \mathbf{A} - \sum_{i=1}^{n-1} \sum_{u,v=1}^3 \lambda_{iuv} \mathbf{E}_{iuv}. \quad (19)$$

Calling $\mathbf{M}(\boldsymbol{\lambda})$ the matrix in (18), the dual function (17) can thus be written as:

$$d(\boldsymbol{\lambda}) = \inf_{\mathbf{x}, \mathbf{y}} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}^\top \mathbf{M}(\boldsymbol{\lambda}) \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} + \sum_{\substack{i=1, \dots, n-1 \\ u=1,2,3}} \lambda_{iuv} + \lambda_y. \quad (20)$$

Now in the dual problem (14) we try to maximize $d(\boldsymbol{\lambda})$; however, from (20) we see that $d(\boldsymbol{\lambda}) = -\infty$ if $\mathbf{M}(\boldsymbol{\lambda})$ has a negative eigenvalue (by letting $[\mathbf{x}, \mathbf{y}]$ lie in the corresponding eigenspace). Consequently, we can safely restrict our search to the vectors $\boldsymbol{\lambda}$ that preserve positive semi-definiteness of $\mathbf{M}(\boldsymbol{\lambda})$ [20, Sec. 5.1.5]. Moreover:

$$\mathbf{M}(\boldsymbol{\lambda}) \succeq 0 \quad \Rightarrow \quad \inf_{\mathbf{x}, \mathbf{y}} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}^\top \mathbf{M}(\boldsymbol{\lambda}) \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} = 0, \quad (21)$$

as the minimization over $[\mathbf{x}, \mathbf{y}]$ in the homogenized problem (21) is unconstrained. The dual problem (14) thus becomes:

$$\begin{aligned} d^* = \quad & \max_{\boldsymbol{\lambda}} \quad \sum_{\substack{i=1, \dots, n-1 \\ u=1,2,3}} \lambda_{iuv} + \lambda_y \quad (\text{dual problem}) \\ \text{subject to} \quad & \mathbf{M}(\boldsymbol{\lambda}) \succeq 0. \end{aligned} \quad (22)$$

The dual SLAM problem turns out to be a *semidefinite program* (SDP), for which specialized solvers exist [22]. In the following section we discuss the relations between the primal and the dual problem, and elucidate on its practical use.

IV. RELATION BETWEEN THE PRIMAL AND THE DUAL PROBLEM AND PRACTICAL USE

In this section we present two powerful applications of the dual problem (22). Section IV-A deals with the case in which one is given a candidate PGO solution, and wants to evaluate its quality, possibly certifying its optimality. Section IV-B shows that in particular cases (when the duality gap is zero) one can obtain an optimal solution of the primal problem from the solution λ^* of the dual.

In both sections, we use the following property.

Lemma 1 (Primal optimal solution and zero duality gap):

If the duality gap is zero ($d^* = f^*$), then any primal optimal solution $[x^* \ 1]$ of (12) is in the null space of the matrix $M(\lambda^*)$, where λ^* is the solution of the dual problem (22).

Proof: When the duality gap is zero, any minimizer of the primal problem is also a minimizer for the infimum in the dual function (14) [20, Sec. 5.5.5]. Consider a primal optimal solution $[x^* \ 1]$. We already observed that any such minimizer annihilates the quadratic term in (20), and therefore it holds that $[x^* \ 1]^T M(\lambda^*) [x^* \ 1] = 0$, which implies that $[x^* \ 1]$ is in the null space of $M(\lambda^*)$, proving the claim. ■

We give an alternative proof, which does not require prior knowledge on duality, in the supplemental material [17].

A. Verification

In this section we consider the case in which we are given a candidate solution $\hat{x}$ for the primal problem, and we want to evaluate the quality of this solution. For brevity, we denote with $f(\hat{x})$ the objective evaluated at $\hat{x}$, for both (11) and its homogeneous form (12) (for the latter we imply $y = 1$).

We begin with the following proposition, whose proof easily follows from (16) and the discussion in Section III-B.

Proposition 2 (Verification of Primal Optimal Objective):

Given a candidate solution $\hat{x}$ for the primal problem (12), if $f(\hat{x}) = d^*$, then the duality gap is zero and $\hat{x}$ is an optimal solution of (12). Moreover, even if the duality gap is nonzero, $f(\hat{x}) - d^* \geq f(\hat{x}) - f^*$, meaning that $f(\hat{x}) - d^*$ is an upper-bound for the sub-optimality gap of $\hat{x}$.

Proposition 2 ensures that the candidate $\hat{x}$ is optimal when $f(\hat{x}) = d^*$. Moreover, even in the case in which we get $f(\hat{x}) > d^*$, the quantity $f(\hat{x}) - d^*$ can be used as an indicator of how far $\hat{x}$ is from the global optimum.

While Proposition 2 already provides means of verifying a candidate solution, it requires solving the dual problem, to compute d^* . The following proposition provides a technique to verify the optimality of $\hat{x}$ without solving the SDP.

Proposition 3 (Verification of Primal Optimal Solution):

Given a candidate solution $\hat{x}$ for the primal problem (12), if the solution $\hat{\lambda}$ of the linear system

$$M(\hat{\lambda}) \begin{bmatrix} \hat{x} \\ 1 \end{bmatrix} = 0 \quad (\text{to be solved w.r.t. } \hat{\lambda}) \quad (23)$$

is such that $M(\hat{\lambda}) \succeq 0$ and $d(\hat{\lambda}) = f(\hat{x})$, then the duality gap is zero and $\hat{x}$ is a primal optimal solution.

Proof: From Lemma 1, we know that when the duality gap is zero, it must hold $M(\lambda^*)[x^* \ 1] = 0$. Therefore,

in Proposition 3 we solve the linear system (23), trying to obtain λ^* . When $M(\hat{\lambda}) \succeq 0$, the solution $\hat{\lambda}$ of (23) is such that $d(\hat{\lambda}) \leq d^*$ (recall that d^* is the maximum over λ). Therefore, it holds that (i) $d(\hat{\lambda}) \leq d^* \leq f^* \leq f(\hat{x})$ (by weak duality and optimality of f^*). However, if $d(\hat{\lambda}) = f(\hat{x})$, the chain of inequalities (i) becomes tight, $d(\hat{\lambda}) = d^* = f^* = f(\hat{x})$, implying that $\hat{x}$ attains the optimal objective f^* . ■

In practice, iterative SLAM solvers optimize (6), rather than the primal problem (12). A natural question is then how to use the results in Propositions 2-3 (which relate f^* and d^*) to verify the solution of (6), whose optimal value is $f_{\text{ML}}^* \geq f^*$. This extension is given by the following proposition, which essentially states that Propositions 2-3 can be applied directly to check the solution of (6).

Proposition 4 (Verification techniques for PGO): The following statements hold true:

- (V1) Given a candidate solution $\hat{x}$ for the PGO problem (6), if $f_{\text{ML}}(\hat{x}) = d^*$, then $\hat{x}$ is an optimal solution of (6). Moreover, $f_{\text{ML}}(\hat{x}) - d^* \geq f_{\text{ML}}(\hat{x}) - f_{\text{ML}}^*$, i.e., $f_{\text{ML}}(\hat{x}) - d^*$ is an upper-bound for the sub-optimality gap of $\hat{x}$.
- (V2) Given a candidate solution $\hat{x}$ for the PGO problem (6), if the solution $\hat{\lambda}$ of the linear system (23) is such that $M(\hat{\lambda}) \succeq 0$ and $d(\hat{\lambda}) = f_{\text{ML}}(\hat{x})$, then the duality gap is zero and $\hat{x}$ is an optimal solution of (6).

Proof: The first claim can be proven by observing that the following chain of inequalities holds (i) $d^* \leq f^* \leq f_{\text{ML}}^* \leq f_{\text{ML}}(\hat{x})$, hence $f_{\text{ML}}(\hat{x}) = d^*$ implies that $d^* = f^* = f_{\text{ML}}^* = f_{\text{ML}}(\hat{x})$, which implies that $\hat{x}$ is optimal. The inequality $f_{\text{ML}}(\hat{x}) - d^* \geq f_{\text{ML}}(\hat{x}) - f_{\text{ML}}^*$ easily follows from (i). The second claim can be proven in the same way, noting that (ii) $d(\hat{\lambda}) \leq d^* \leq f^* \leq f_{\text{ML}}^* \leq f_{\text{ML}}(\hat{x})$: if we are able to compute a $\hat{\lambda}$ that is dual feasible ($M(\hat{\lambda}) \succeq 0$) and such that $d(\hat{\lambda}) = f_{\text{ML}}(\hat{x})$, then it must hold $d(\hat{\lambda}) = d^* = f^* = f_{\text{ML}}^* = f_{\text{ML}}(\hat{x})$, which implies that $f_{\text{ML}}(\hat{x})$ is optimal. ■

B. Primal Optimal Solutions

In this section, rather than verifying the quality of a given candidate solution, we show how to use duality to compute a primal optimal solution directly. We focus on the particular case in which the duality gap is zero. This case is of interest as we observe that the duality gap is often zero in practice.

From Lemma 1 we know that an optimal solution must be in the null space of $M(\lambda^*)$. This motivates the following proposition, which provides a way to compute a primal optimal solution directly from the solution λ^* of (22).

Proposition 5: If the duality gap is zero and λ^* is an optimal solution of (22), then an optimal solution x^* of (12) can be computed by solving the following linear system:

$$\begin{bmatrix} H(\lambda^*) \\ b^T A \end{bmatrix} x^* = \begin{bmatrix} A^T b \\ \lambda_y^* - b^T b \end{bmatrix} \quad (\text{to be solved w.r.t. } x^*) \quad (24)$$

Proof: From Lemma 1, we know that when the duality gap is zero, it holds $M(\lambda^*)[x^* \ 1] = 0$. Then, recalling the structure of $M(\lambda^*)$ from (18), it's easy to see that (24) only rewrites the condition $M(\lambda^*)[x^* \ 1] = 0$, moving the constant terms to the right-hand side. ■

Fig. 2. Statistics regarding the verification technique V1 (first row), the computation of the primal solution via the dual (second row), and the verification technique V2 (third row). Results are shown for different levels of translation noise σ_T (first column), rotation noise σ_R (second column), and size of the problem (third column). The rightmost column shows the CPU time required by V1 and V2 and the datasets used for the Monte Carlo simulations.

Note that Lemma 1 ensures that the linear system (24) admits a solution. Proposition 5 allows finding an optimal solution to the primal problem (12). If this satisfies the determinant constraints in (6), it follows that this solution is also optimal for the original PGO problem (6). While currently we cannot prove that the determinant constraints are always satisfied, this was always the case in our experiments.

V. EXPERIMENTS

In this section we show that the first verification technique in Proposition 4 (referred to as V1) enables accurate quantification of the sub-optimality gap of a candidate SLAM solution, but that current SDP solvers do not scale well with increasing problem size. The second verification technique in Proposition 4 (referred to as V2) provides a convenient alternative for large-scale problems: it is reliable and less computationally demanding. Finally, we provide empirical evidence that when the duality gap is zero we can compute an optimal solution from the solution of the dual problem, as suggested by Proposition 5. The SDP (22) is solved using SDPA [22].

Effectiveness of V1. Proposition 4(V1) ensures that for any candidate solution $\hat{x}$, $f_{ML}(\hat{x}) - d^*$ represents an upper-bound on the sub-optimality gap $f_{ML}(\hat{x}) - f_{ML}^*$ of $\hat{x}$. However, it is possible that this bound is very loose, in which case it would be of little practical utility. In this section, we show that d^* is close to f_{ML}^* in practice, hence $f_{ML}(\hat{x}) - d^*$ is a very good measure of the sub-optimality gap of the candidate $\hat{x}$.

In our experiments we compute the “optimal” solution f_{ML}^* of (6) by refining the *chordal* initialization of [23], [24] with 10

Gauss-Newton (GN) iterations. While one cannot guarantee *a priori* that this approach always produces the optimal estimate, using the results of this paper we will be able to check optimality a posteriori.

We evaluated how close d^* is to f_{ML}^* on the cube dataset of Fig. 2(b4). In this dataset, the odometric trajectory is simulated as the robot travels on a 3D grid world, and random loop closures are added between nearby nodes, with probability 0.3. Relative pose measurements are obtained by contaminating the true relative poses with zero-mean Gaussian noise, with standard deviation σ_T and σ_R for the translational and rotational noise, respectively. Statistics are computed over 10 runs: for each run we create a cube with random connectivity and random measurement noise. We consider an example with $n = 5^3$ poses and varying noise levels σ_T and σ_R .

Fig. 2(a1) shows d^* and f_{ML}^* for different translational noise levels, fixing $\sigma_R = 0.05\text{rad}$. The figure shows that $d^* = f_{ML}^*$ (zero duality gap) independently on the translational noise level, hence d^* is a very good proxy of f_{ML}^* .

Fig. 2(a2) shows d^* and f_{ML}^* for different rotational noise, fixing $\sigma_T = 0.1\text{m}$. In this case the duality gap $f_{ML}^* - d^*$ is more sensitive to the noise level, and for large rotational noise d^* becomes smaller than f_{ML}^* . However, the gap $f_{ML}^* - d^*$ remains small, and is within 20% of f_{ML}^* in all cases.

Fig. 2(a3) shows d^* and f_{ML}^* for different sizes of the cube dataset, fixing $\sigma_R = 0.05\text{rad}$ and $\sigma_T = 0.1\text{m}$. Again in this case $d^* = f_{ML}^*$ (zero duality gap) independently of the size of the dataset. However, we observe that the SDP (22) becomes intractable for larger problem sizes, as shown in Fig. 2(a4).

Primal optimal solution via the dual. Here we demonstrate experimentally that one can recover a primal optimal solution $\mathbf{x}^*$ for the SLAM problem from the dual optimal solution λ^* of (22) whenever strong duality holds.

For each of the experimental trials described previously, we also computed an estimate $\hat{\mathbf{x}}$ for the SLAM solution directly from the optimal solution λ^* for the dual program (22) using equation (24). Figs. 2(b1)-(b3) compare the value $f(\hat{\mathbf{x}})$ of this estimate against the value f_{ML}^* of the solution obtained by solving the SLAM problem directly using the chordal initialization, and against the cost of the odometric guess.

We can see that for those experimental conditions in which strong duality holds (those experiments in Figs. 2(a1)-(a3) for which the green and red bars are the same height), the estimate $\hat{\mathbf{x}}$ in fact achieves the (certified) globally optimal cost f_{ML}^* and is therefore a globally optimal solution for the SLAM problem, as guaranteed by Proposition 5. More interestingly, we also find that even in those cases when strong duality does *not* hold (in which case Proposition 5 no longer guarantees that $\hat{\mathbf{x}}$ is a primal optimal solution), the quality of the candidate $\hat{\mathbf{x}}$ degrades gracefully (i.e. the gap $f(\hat{\mathbf{x}}) - f_{\text{ML}}^*$ increases gradually) with increasing noise levels. In particular, we find that $\hat{\mathbf{x}}$ outperforms the odometric guess in all tested cases.

These experiments confirm that we can extract a primal optimal solution $\mathbf{x}^*$ for the SLAM problem *directly* from the solution λ^* of the convex dual problem whenever strong duality holds. However, this approach is unfortunately not currently practical as a general-purpose SLAM technique, due to the high computational cost of solving large scale SDPs.

Effectiveness of V2. In this section we show that V2 is a computationally tractable verification approach, and it preserves the desirable properties of V1, i.e., it is able to discern optimal estimates from suboptimal ones. In contrast to V1, V2 does not quantify the sub-optimality gap, but can only give a binary answer: either it certifies that the estimate $\hat{\mathbf{x}}$ is optimal (by producing a dual certificate $\hat{\lambda}$), or it is inconclusive. Before moving to the real tests, we consider the cube scenario discussed in the previous section.

We perform the following test: for each realization of the cube scenario, we compute the optimal solution $\mathbf{x}^*$ (attaining f_{ML}^*) as in the previous tests, and a sub-optimal solution $\mathbf{x}^\dagger$ (attaining $f^\dagger > f^*$) by bootstrapping the Gauss-Newton method with a random initial guess. Then, we apply the verification technique to both $\mathbf{x}^*$ and $\mathbf{x}^\dagger$ and see if they pass the optimality test. Using the results we can compute the *precision* and *recall* of our classification:

$$\text{precision} = \frac{|\checkmark \mathbf{x}^*|}{|\checkmark \mathbf{x}^*| + |\checkmark \mathbf{x}^\dagger|} \quad \text{recall} = \frac{|\checkmark \mathbf{x}^*|}{|\checkmark \mathbf{x}^*| + |\text{X} \mathbf{x}^*|} \quad (25)$$

where $|\checkmark \mathbf{x}^*|$ denotes the number of tests in which an optimal solution $\mathbf{x}^*$ was accepted as optimal by V2, $|\checkmark \mathbf{x}^\dagger|$ is the number of tests in which a suboptimal solution was accepted as optimal, and $|\text{X} \mathbf{x}^*|$ is the number of tests in which V2 was not able to certify the optimality of an optimal solution.

We point out that Proposition 4(V2) guarantees that our certification approach always has precision equal to 1; however, recall may be less than 1, and will be for cases in which strong duality does not hold. We plot precision/recall for different

		$f_{\text{ML}}(\hat{\mathbf{x}})$	$d(\hat{\lambda})$	μ	Time	V2
sphere $n=2500$ $m=4949$	Init.	$5.7595 \cdot 10^2$	$5.75 \cdot 10^2$	$-1.80 \cdot 10^{-4}$	609	✓
	Odom.	$5.8019 \cdot 10^2$	$4.38 \cdot 10^2$	$-1.45 \cdot 10^{-4}$	597	X
sphere-a $n=2200$ $m=8647$	Init.	$1.2485 \cdot 10^6$	$1.25 \cdot 10^6$	$-9.64 \cdot 10^{-3}$	332	✓
	Odom.	$3.0413 \cdot 10^6$	$3.04 \cdot 10^6$	$-1.01 \cdot 10^2$	2.5	X
torus $n=5000$ $m=9048$	Init.	$1.2114 \cdot 10^4$	$1.21 \cdot 10^4$	$-7.85 \cdot 10^{-2}$	39	✓
	Odom.	$2.7666 \cdot 10^4$	$2.76 \cdot 10^4$	$-1.04 \cdot 10^2$	3.7	X
cube $n=8000$ $m=22236$	Init.	$4.216 \cdot 10^4$	$4.22 \cdot 10^4$	$-1.24 \cdot 10^{-2}$	1646	✓
	Odom.	$2.7465 \cdot 10^5$	$2.74 \cdot 10^5$	$-9.98 \cdot 10^1$	68.2	X
garage $n=1661$ $m=6275$	Init.	$6.2994 \cdot 10^{-1}$	$6.11 \cdot 10^{-1}$	$-6.78 \cdot 10^{-2}$	12.3	✓
	Odom.	$6.2997 \cdot 10^{-1}$	$3.53 \cdot 10^{-1}$	$-6.75 \cdot 10^{-2}$	11.2	X
cubicle $n=5750$ $m=16869$	Init.	$6.2481 \cdot 10^2$	$6.25 \cdot 10^2$	$-4.76 \cdot 10^{-1}$	16.8	✓
	Odom.	$6.2484 \cdot 10^2$	$6.16 \cdot 10^2$	$-4.75 \cdot 10^{-1}$	15.5	✓
rim $n=10195$ $m=29743$	Init.	$1.235 \cdot 10^4$	$1.23 \cdot 10^4$	$-9.77 \cdot 10^1$	7.7	X
	Odom.	$1.6985 \cdot 10^4$	$-2.8 \cdot 10^4$	$-9.38 \cdot 10^1$	7.1	X

TABLE I
VERIFICATION TECHNIQUE V2 ON LARGE-SCALE SLAM DATASETS.

levels of translational and rotational noise, and for different sizes of the dataset, as shown in Figs. 2(c1)-(c2)-(c3). We observe that only for larger rotational noise Fig. 2(c2) the recall decreases; these are exactly the cases in which the duality gap is nonzero. (For $\sigma_R = 0.2\text{rad}$, V2 was not able to certify optimality in any case, which means that the precision becomes undefined ($\frac{0}{0}$) and for this reason we do not show the corresponding data point in Fig. 2(c2).) Finally, Fig. 2(d2) shows the CPU time required by V2. This is the time required to solve the linear system (23), and check if $M(\hat{\lambda}) \succeq 0$. V2 is computationally cheap, as it does not require solving the SDP.

We conclude the experimental part of this paper by testing the performance of the verification technique V2 on large-scale SLAM datasets. We consider the same datasets as [23]: the sphere, sphere-a, torus and cube are simulated datasets, while the garage, cubicle and rim are real datasets. For the results in this paper we substituted the covariances in the datasets of the scenarios sphere, sphere-a, garage, cubicle and rim with isotropic ones, as required by the PGO formulation (6).

Table I shows the results of the application of V2 to the SLAM datasets. The first column shows the cost obtained by applying a GN method starting from the initialization [23] (rows “Init.” in the table) and from the odometric guess (rows “Odom.”). These are the candidate solutions that we want to check, using V2. The columns “ $d(\hat{\lambda})$ ” and “ μ ” show intermediate results of V2. In particular, $d(\hat{\lambda})$ is the same one described in Proposition 4(V2), while μ is the smallest eigenvalue of $M(\hat{\lambda})$. Recall that V2 certifies optimality when $f_{\text{ML}}(\hat{\mathbf{x}}) = d(\hat{\lambda})$ and $\mu \geq 0$ (if the smallest eigenvalue is non-negative, then $M(\hat{\lambda}) \succeq 0$). We specify a tolerance in these tests, since the GN estimate $\hat{\mathbf{x}}$ will not attain the optimal solution $\mathbf{x}^*$ exactly. Consequently, the solution $\hat{\lambda}$ of (23) is not exact, so we consider that $f_{\text{ML}}(\hat{\mathbf{x}}) = d(\hat{\lambda})$ if $|f_{\text{ML}}(\hat{\mathbf{x}}) - d(\hat{\lambda})|/f_{\text{ML}}(\hat{\mathbf{x}}) < 20\%$. Similarly, we add some tolerance to the condition $\mu \geq 0$, and accept $\mu \geq -1$. Recall that μ is expected to be slightly negative, as the objective in (22) tends to push $M(\lambda)$ towards the boundary of the positive definite cone, and the SDP is solved numerically.

Let us start our analysis from the sphere dataset. When using the initialization (“Init.” row), the GN method attains an objective $5.7595 \cdot 10^2$. In this case, the two conditions for V2 are satisfied and we can certify the optimality of the resulting estimate (green check-mark in the rightmost column).

When using the odometric initialization the resulting cost is larger (red entry in the column $f_{\text{ML}}(\hat{x})$), hence the estimate is suboptimal. V2 is able to identify the suboptimality, since $d(\hat{\lambda})$ becomes much smaller than $f_{\text{ML}}(\hat{x})$. Hence V2 correctly decides not to certify the optimality of the odometric estimate (red “X” in the last column). Similar considerations hold for the second scenario, the challenging sphere-a: in this case the initialized estimate (shown in Fig. 1, top left) is accepted as optimal. The odometric estimate, instead, is trapped in a local minimum (Fig. 1, bottom left), and the optimality test V2 correctly rejects the estimate since it leads to a very negative μ . Similar considerations hold for the torus, cube, and the cubicle datasets: our technique is able to discern optimal solutions from suboptimal estimates in all cases.

For the garage dataset, the initialized estimate is classified as optimal, while the odometric estimate, which has a similar cost ($6.2997 \cdot 10^{-1}$ vs $6.2994 \cdot 10^{-1}$), is rejected as suboptimal. We observed the corresponding trajectory estimates (see [17]) and, while they are both visually correct and have very similar costs, they do not overlap. This may indicate the presence of regions in the cost functions that are nearly flat, i.e., for which different estimates can have similar cost. Allowing for extra GN iterations, the odometric estimate converges to the same cost as the initialized one, and our technique is able to certify its optimality. Empirically, we observed that estimates that are suboptimal because they need extra iterations to converge tend to fail the check on $d(\hat{\lambda})$ (compare with sphere and garage), while estimates that converge to a wrong minimum tend to fail the check on μ (compare with sphere-a, torus, and cube).

For the rim dataset, both estimates (“init” and “odom”) are rejected by V2. While the initialized estimate is visually correct, currently, we cannot conclude anything as it might be that there exists a better estimate (i.e., GN failed), or our technique failed because the duality gap was nonzero.

Finally, in the column “Time”, we report the CPU time (in seconds) required to perform the second verification technique. Most of the time here is spent computing the smallest eigenvalue μ of $M(\hat{\lambda})$, which was obtained using Matlab’s `eigs`, specifying -100 as guess for the eigenvalue. The CPU time depends on the size of the problem, but also depends on the distance between our guess (-100) and the closest eigenvalue. We leave a more thorough investigation of these computational aspects for future work.

VI. CONCLUSION

We show that Lagrangian duality is an effective tool to assess the quality of a given SLAM estimate. We propose two techniques to judge if an estimate is globally optimal (i.e., it is the ML estimate), and we show that when the duality gap is zero one can compute an optimal SLAM solution from the dual problem. The performance of our verification techniques is extensively tested in real and simulated datasets, including large-scale benchmarks. Many theoretical questions remain open, e.g., under which conditions the duality gap is zero? is it possible to derive bounds on the duality gap? We also leave as future work more practical problems, such as the one of exploring more efficient solvers for large SDPs, possibly exploiting problem structure.

REFERENCES

- [1] F. Lu and E. Milios, “Globally consistent range scan alignment for environment mapping,” *Autonomous Robots*, pp. 333–349, Apr 1997.
- [2] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, “g2o: A general framework for graph optimization,” in *Proc. of the IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, May 2011.
- [3] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, “iSAM2: Incremental smoothing and mapping using the Bayes tree,” *Intl. J. of Robotics Research*, vol. 31, pp. 217–236, Feb 2012.
- [4] E. Olson, J. Leonard, and S. Teller, “Fast iterative alignment of pose graphs with poor initial estimates,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, May 2006, pp. 2262–2269.
- [5] G. Grisetti, C. Stachniss, and W. Burgard, “Non-linear constraint network optimization for efficient map learning,” *Trans. on Intelligent Transportation systems*, vol. 10, no. 3, pp. 428–439, 2009.
- [6] D. Rosen, M. Kaess, and J. Leonard, “RISE: An incremental trust-region method for robust online sparse least-squares estimation,” *IEEE Trans. Robotics*, 2014.
- [7] F. Dellaert and A. Stroupe, “Linear 2D localization and mapping for single and multiple robots,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, May 2002.
- [8] L. Carlone, R. Aragues, J. Castellanos, and B. Bona, “A fast and accurate approximation for planar pose graph optimization,” *Intl. J. of Robotics Research*, 2014.
- [9] L. Carlone and F. Dellaert, “Duality-based verification techniques for 2D SLAM,” in *Intl. Conf. on Robotics and Automation (ICRA)*, accepted, 2015.
- [10] D. Rosen, C. DuHadway, and J. Leonard, “A convex relaxation for approximate global optimization in simultaneous localization and mapping,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2015.
- [11] R. Tron, B. Afsari, and R. Vidal, “Intrinsic consensus on $SO(3)$ with almost global convergence,” in *IEEE Conf. on Decision and Control*, 2012.
- [12] L. Carlone and A. Censi, “From angular manifolds to the integer lattice: Guaranteed orientation estimation with application to pose graph optimization,” *IEEE Trans. Robotics*, 2014.
- [13] S. Huang, Y. Lai, U. Frese, and G. Dissanayake, “How far is SLAM from a linear least squares problem?” in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2010, pp. 3011–3016.
- [14] H. Wang, G. Hu, S. Huang, and G. Dissanayake, “On the structure of nonlinearities in pose graph SLAM,” in *Robotics: Science and Systems (RSS)*, 2012.
- [15] L. Carlone, “Convergence analysis of pose graph optimization via Gauss-Newton methods,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2013, pp. 965–972.
- [16] K. Khosoussi, S. Huang, and G. Dissanayake, “Novel insights into the impact of graph structure on SLAM,” in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2014.
- [17] L. Carlone, D. Rosen, G. Calafiore, J. Leonard, and F. Dellaert, “Lagrangian duality in 3D SLAM: Verification techniques and optimal solutions. supplemental material.” [Online]. Available: www.lucacarlone.com/index.php/resources/research/duality3DPGO
- [18] N. Boumal, A. Singer, and P. Absil, “Cramer-rao bounds for synchronization of rotations,” *Information and Inference*, vol. 3, no. 1, pp. 1–39, 2014.
- [19] R. Hartley, J. Trumpf, Y. Dai, and H. Li, “Rotation averaging,” *IJCV*, vol. 103, no. 3, pp. 267–305, 2013.
- [20] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
- [21] G. Calafiore and L. E. Ghaoui, *Optimization Models*. Cambridge University Press, 2014.
- [22] M. Yamashita, K. Fujisawa, and M. Kojima, “Implementation and evaluation of SDPA 6.0 (semidefinite programming algorithm 6.0),” *Optimization Methods and Software*, vol. 18, no. 4, pp. 491–505, Aug. 2003.
- [23] 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,” in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2015.
- [24] D. Martinec and T. Pajdla, “Robust rotation and translation estimation in multiview reconstruction,” in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2007, pp. 1–8.