

On the Inclusion of Determinant Constraints in Lagrangian Duality for 3D SLAM

Roberto Tron*, David M. Rosen†, Luca Carlone‡

*University of Pennsylvania, Philadelphia, PA, USA. Email: tron@seas.upenn.edu

†Massachusetts Institute of Technology, Cambridge, MA, USA. Email: dmrosen@mit.edu

‡Georgia Institute of Technology, Atlanta, GA, USA. Email: luca.carlone@gatech.edu

Abstract—A recent paper in 3D Pose Graph Optimization (PGO) shows how a dual Lagrangian formulation of the problem can be used to verify (and possibly certify) the quality of a given solution [1]. A limitation of this approach is that, in the derivation, the authors relax the positive determinant constraint for the rotations. As a consequence, when the approach fails to certify an optimal solution (i.e., when the duality gap is non-zero), one cannot determine if this is due to the relaxation or if it is an intrinsic feature of the problem at hand.

In this paper we extend the results of [1] by including the determinant constraints in the derivation of the dual, thus showing that their relaxation is unnecessary. We show experimentally that this complete formulation does not lead to tangible differences with respect to the original, relaxed version. This indicates that the reasons for failures in providing a certificate of optimality are intrinsic to the problem, and that the determinant constraints are somehow redundant in common PGO instances.

I. INTRODUCTION

Pose Graph Optimization (PGO) is one of the most popular formulations of the SLAM problem. The goal of PGO is to estimate a set of poses (sampled along the trajectory of a mobile robot) from relative measurements. This is done by solving a non-convex optimization problem, where the non-convexity is mainly due to the presence of unknown rotations.

Despite the empirical success of state-of-the-art techniques for PGO (see [1] for a literature review) current approaches are not able to guarantee the computation of a global optimal solution in general problem instances. This is due to the non-convexity of PGO, which implies that iterative optimization techniques can be trapped in local minima, resulting in wrong pose estimates. We argue that the transition of SLAM from research topic to industrial technology requires techniques with *guaranteed performance*. For instance, in autonomous vehicles, 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, preventing the reliable operation of mobile robots.

Driven by this motivation, previous works [1–3] use Lagrangian duality as a tool to obtain clear performance guarantees. The dual problem can be used for “verification” purposes [1, 2], i.e., given a PGO estimate (e.g., returned by an iterative solver), we can use the (convex) dual problem to evaluate the quality of this estimate and possibly certify its optimality. Moreover, the dual problem enables the computation of a guaranteed optimal solution in particular cases (when the duality gap is zero) [1, 2]. This second use

currently seems less appealing, since it requires solving a large semidefinite program (SDP); despite the fact that SDPs are convex problems, current solvers do not scale well and prevent real time operation. On the other hand, the verification can be done without solving the SDP [1], and this makes duality appealing from a practical standpoint.

The present paper provides an extension of the formulation in [1], which derives a dual formulation for 3D PGO by relaxing the constraints that rotations have positive unit determinant. Effectively, this changes the domain of the optimization problem from rotation matrices (belonging to the group $SO(3)$) to orthogonal matrices (belonging to the group $O(3)$). In this paper we demonstrate that this relaxation is unnecessary, and we show how to include the determinant constraints in the derivation of the dual (Sections III-IV). Moreover, we present a short summary of the results of [1] (Section V), omitting the proofs. We conclude the paper with an experimental comparison between the dual problem of [1] and the extended version proposed in this addendum (Section VI). Our Monte Carlo analysis shows that including the determinant constraints does not significantly impact the results, suggesting that these constraints are redundant in common PGO instances. We provide an intuitive explanation for this empirical observation in Section VII, and we draw conclusions in Section VIII.

II. 3D POSE GRAPH OPTIMIZATION

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 $\bar{\mathbf{x}}_{ij}$.

We consider the following PGO formulation:

$$f^* = \min_{\substack{\{\mathbf{t}_i \in \mathbb{R}^3\} \\ \{\mathbf{R}_i \in SO(3)\}}} \sum_{(i,j) \in \mathcal{E}} \omega_i^2 \|\mathbf{t}_j - \mathbf{t}_i - \mathbf{R}_i \bar{\mathbf{t}}_{ij}\|^2 \quad (\text{PGO}) \\ + \frac{\omega_R^2}{2} \|\mathbf{R}_j - \mathbf{R}_i \bar{\mathbf{R}}_{ij}\|_F^2 \quad (1)$$

in which we compute the translations $\{\mathbf{t}_i \in \mathbb{R}^3\}$ and the rotations $\{\mathbf{R}_i \in SO(3)\}$ by minimizing the residual errors

with respect to the given measurements $(\bar{\mathbf{R}}_{ij}, \bar{\mathbf{t}}_{ij}), \forall (i, j) \in \mathcal{E}$. In (1), $\|\cdot\|_F^2$ is the (squared) Frobenius matrix norm (sum of the squares of the entries), and ω_t^2, ω_R^2 are the inverse of the translation and rotation measurement covariances, which are assumed isotropic (the derivation of the maximum likelihood estimator is given in [1]). The norm $\|\mathbf{R}_a - \mathbf{R}_b\|_F^2$ is usually referred to as the *chordal distance* between two rotations \mathbf{R}_a and \mathbf{R}_b [4]. The main difference between (1) and formulations in related work is the use of the chordal distance in place of the commonly used *geodesic distance*. The chordal distance has been already proposed in a SLAM context in [5, 6]

The advantage of the formulation (1) is that it has a quadratic objective function. This facilitates the derivation of the Lagrangian dual problem (Section IV). Before presenting the dual, we reformulate (1) in a more convenient form. This is done in the following section.

III. 3D SLAM AS A QUADRATIC PROBLEM WITH QUADRATIC EQUALITY CONSTRAINTS

In this section, we rewrite (1) in order to (i) have vector variables (the rotations \mathbf{R}_i are matrices), (ii) express the constraints $\mathbf{R}_i \in \text{SO}(3)$ as quadratic equality constraints, (iii) anchor one of the poses to the origin of the reference frame (this is standard in PGO solvers). This reformulation makes the derivation of the dual problem straightforward.

A. Vectorization

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 (1) becomes (using $\|\mathbf{R}\|_F = \|\mathbf{R}^\top\|_F$ in the first expression):

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

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.

B. Constraints in PGO

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:

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

$$\det(\mathbf{R}_i) = 1 \Leftrightarrow \mathbf{R}_i^{(1)} \times \mathbf{R}_i^{(2)} = \mathbf{R}_i^{(3)} \quad (4)$$

where \times denotes the cross product. In other words, the rows of a rotation matrix have to be orthonormal (3), and

have to satisfy the right-hand rule (4). The orthogonality constraints (3) can be written in terms of the vector \mathbf{r}_i :

$$(\mathbf{R}_i^{(u)})^\top \mathbf{R}_i^{(v)} = 0 \Leftrightarrow \mathbf{r}_i^\top \mathbf{E}_{uv} \mathbf{r}_i = 0 \quad u \neq v, \quad (5)$$

$$(\mathbf{R}_i^{(u)})^\top \mathbf{R}_i^{(u)} = 1 \Leftrightarrow \mathbf{r}_i^\top \mathbf{E}_{uu} \mathbf{r}_i = 1 \quad (6)$$

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. To facilitate the following derivation, we include a slack variable y and a constraint:

$$y^2 = 1 \quad (7)$$

which allows writing (6) equivalently as:

$$\mathbf{r}_i^\top \mathbf{E}_{uu} \mathbf{r}_i = 1 \Leftrightarrow \mathbf{r}_i^\top \mathbf{E}_{uu} \mathbf{r}_i - y^2 = 0 \quad (8)$$

(the derivation is easier if most constraints are homogeneous).

While in the previous work [1] the authors relaxed the determinant constraint (4), in this paper we show that it is actually possible to include it in the derivation of the dual problem. For this purpose, we define $\mathbf{e}_u \in \mathbb{R}^3$ as the vector that is zero everywhere except the u -th entry which is 1, and we use the following equalities (note that $\mathbf{R}_i^{(k)}$ are treated as column vectors):

$$\begin{aligned} \mathbf{R}_i^{(1)} \times \mathbf{R}_i^{(2)} = \mathbf{R}_i^{(3)} & \Leftrightarrow \mathbf{e}_u^\top \mathbf{R}_i^{(1)} \times \mathbf{R}_i^{(2)} = \mathbf{e}_u^\top \mathbf{R}_i^{(3)}, \quad u=1,2,3 \\ (\text{using } a \times b = [a]_\times b) & \Leftrightarrow \mathbf{e}_u^\top [\mathbf{R}_i^{(1)}]_\times \mathbf{R}_i^{(2)} = \mathbf{e}_u^\top \mathbf{R}_i^{(3)}, \\ (\text{using } [a]_\times b = -[b]_\times a) & \Leftrightarrow -(\mathbf{R}_i^{(1)})^\top [\mathbf{e}_u]_\times \mathbf{R}_i^{(2)} = \mathbf{e}_u^\top \mathbf{R}_i^{(3)}, \\ (\text{including } y) & \Leftrightarrow -(\mathbf{R}_i^{(1)})^\top [\mathbf{e}_u]_\times \mathbf{R}_i^{(2)} = y \mathbf{e}_u^\top \mathbf{R}_i^{(3)}, \\ (\text{using } \mathbf{r}_i) & \Leftrightarrow \mathbf{r}_i^\top \mathbf{S}_u \mathbf{r}_i + y \mathbf{e}_{6+u}^\top \mathbf{r}_i = 0, \end{aligned} \quad (9)$$

where we assumed that $y = 1$ (we will comment more on this in Section III-D), $[a]_\times \in \mathbb{R}^{3 \times 3}$ is a skew-symmetric matrix built from the entries of a vector $\mathbf{a} \in \mathbb{R}^3$, and where \mathbf{S}_u is a 9×9 selection matrix composed of 3×3 blocks that are zero everywhere except the 3×3 block in position $(1, 2)$, which is $[\mathbf{e}_u]_\times$; moreover, since the last three entries of \mathbf{r}_i coincide with $\mathbf{R}_i^{(3)}$, in (9) we also used $\mathbf{e}_u^\top \mathbf{R}_i^{(3)} = \mathbf{e}_{6+u}^\top \mathbf{r}_i$.

Combining the objective function (2) and the constraints (5), (7), (8), (9), we rewrite the PGO problem (1) as:

$$\begin{aligned} f^* = \min_{\{\mathbf{r}_i, \mathbf{t}_i\}} & \sum_{(i,j) \in \mathcal{E}} \omega_t^2 \|\mathbf{t}_j - \mathbf{t}_i - \mathbf{T}_{ij} \mathbf{r}_i\|^2 + \frac{\omega_R^2}{2} \|\mathbf{r}_j - \mathbf{Q}_{ij} \mathbf{r}_i\|^2 \\ \text{subject to} & \left. \begin{aligned} \mathbf{r}_i^\top \mathbf{E}_{uv} \mathbf{r}_i - y^2 &= 0, & u = v \\ \mathbf{r}_i^\top \mathbf{E}_{uv} \mathbf{r}_i &= 0, & u \neq v \\ \mathbf{r}_i^\top \mathbf{S}_u \mathbf{r}_i + y \mathbf{e}_{6+u}^\top \mathbf{r}_i &= 0 \\ y^2 &= 1 \end{aligned} \right\} \begin{matrix} u, v=1,2,3 \\ i=1, \dots, n \end{matrix} \end{aligned} \quad (10)$$

In order to write (10) in a more compact matrix notation, we define the vector $\check{\mathbf{x}} = [\mathbf{t}_1^\top, \dots, \mathbf{t}_n^\top, \mathbf{r}_1^\top, \dots, \mathbf{r}_n^\top, y]^\top \in \mathbb{R}^{12n+1}$. Using this notation, (10) becomes:

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

where the matrices $\check{\mathbf{A}}$ and $\check{\mathbf{E}}_{iuv}$, and $\check{\mathbf{S}}_{iu}$ are obtained by stacking the coefficient matrices in (10), with suitable zero blocks for padding, while $\check{\mathbf{U}}$ is zero everywhere, except a single entry in the bottom-right corner, which is equal to 1.

C. Anchoring

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*.

Preserving an objective function that is homogeneous in the variables is convenient for the derivation of the dual problem; hence, rather than setting $\mathbf{r}_1 = \text{rows}(\mathbf{I}_3)$ (which would lead to a constant terms within the square in the objective), we set $\mathbf{r}_1 = y \text{rows}(\mathbf{I}_3)$ (more discussion on this in Section III-D). Anchoring the first pose modifies (11) as follows:

$$\begin{aligned} f^* = \min_{\mathbf{x}} \quad & \|\mathbf{Ax}\|^2 && \text{(Primal problem)} \\ \text{subject to} \quad & \begin{cases} \mathbf{x}^\top \mathbf{E}_{iuv} \mathbf{x} = 0, & u = v \\ \mathbf{x}^\top \mathbf{E}_{iuv} \mathbf{x} = 0, & u \neq v \\ \mathbf{x}^\top \mathbf{S}_{iu} \mathbf{x} = 0 \\ \mathbf{x}^\top \mathbf{U} \mathbf{x} = 1 \end{cases} && \left. \begin{array}{l} u, v = 1, 2, 3 \\ i = 1, \dots, n-1 \end{array} \right\} \end{aligned}$$

where $\mathbf{x} \in \mathbb{R}^{12(n-1)+1}$ 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 adjusting the last column to accommodate the terms multiplying y (essentially, $\text{rows}(\mathbf{I}_3)$); \mathbf{E}_{iuv} (resp. \mathbf{S}_{iu} , \mathbf{U}) are the same as $\check{\mathbf{E}}_{iuv}$ (resp. $\check{\mathbf{S}}_{iu}$, $\check{\mathbf{U}}$) but without the rows and columns corresponding to the first pose.

D. Effect of the slack variable y

Observe that the constraint $y^2 = 1$ implies that $y = \pm 1$. Hence, when $y = -1$, it seems that we get “wrong” constraints:

- 1) The constraint $\mathbf{r}_1 = y \text{rows}(\mathbf{I}_3)$ becomes $\mathbf{R}_1 = -\mathbf{I}_3$.
- 2) The constraint (10) implies $\det(\mathbf{R}_i) = -1$ for all $i = 1, \dots, n$.

However, noting that the objective in (12) is homogeneous, we can easily see that $\|\mathbf{Ax}\| = \|-\mathbf{Ax}\|$. Hence, we can change the sign of y and of \mathbf{R}_i , $i = 1, \dots, n$ to satisfy the right constraints without altering the cost. This means that, without loss of generality, we can assume $y = 1$. This makes the constraints in our homogeneous formulation equivalent to the original ones.

IV. THE DUAL PROBLEM

Standard optimization theory tells us that to every constrained optimization problem (called the *primal* problem), we can associate a *dual* problem, and it provides useful results to relate the two problems [7, 8]. In this section we derive the Lagrangian dual problem of the primal problem (12).

The first step to derive the dual is to build the *Lagrangian*:

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) = \quad & \|\mathbf{Ax}\|^2 - \left(\sum_{i=1}^{n-1} \sum_{u,v=1,2,3} \lambda_{iuv}^\perp \mathbf{x}^\top \mathbf{E}_{iuv} \mathbf{x} \right) && (13) \\ & - \left(\sum_{i=1}^{n-1} \sum_{u=1,2,3} \lambda_{iu}^{\det} \mathbf{x}^\top \mathbf{S}_{iu} \mathbf{x} \right) - \lambda^y \left(\mathbf{x}^\top \mathbf{U} \mathbf{x} - 1 \right) \end{aligned}$$

where λ_{iuv}^\perp , λ_{iu}^{\det} , and λ^y are the *Lagrange multipliers* or *dual variables* associated to the orthogonality, determinant, and $y^2 = 1$ constraints in (12), respectively; we stack all the dual variables in a vector $\boldsymbol{\lambda}$. The Lagrangian can be understood as a function that includes the objective function in (12) and penalty terms corresponding to each constraint in (12). We note that the only difference (besides minor notation changes) with respect to [1] is the presence of the terms “ $\lambda_{iu}^{\det} \mathbf{x}^\top \mathbf{S}_{iu} \mathbf{x}$ ”, which are due to the determinant constraints.

To make the notation more compact, we define the matrix:

$$\mathbf{W}(\boldsymbol{\lambda}) \doteq \mathbf{A}^\top \mathbf{A} - \sum_{i=1}^{n-1} \left(\sum_{\substack{u,v= \\ 1,2,3}} \lambda_{iuv}^\perp \mathbf{E}_{iuv} + \sum_{\substack{u= \\ 1,2,3}} \lambda_{iu}^{\det} \mathbf{S}_{iu} \right) - \lambda^y \mathbf{U} \quad (14)$$

Using $\mathbf{W}(\boldsymbol{\lambda})$ we write the Lagrangian succinctly as:

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) = \mathbf{x}^\top \mathbf{W}(\boldsymbol{\lambda}) \mathbf{x} + \lambda^y \quad (15)$$

(12) The *dual function* $d(\boldsymbol{\lambda})$ is the infimum of the Lagrangian with respect to \mathbf{x} :

$$d(\boldsymbol{\lambda}) = \inf_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) = \inf_{\mathbf{x}} \mathbf{x}^\top \mathbf{W}(\boldsymbol{\lambda}) \mathbf{x} + \lambda^y \quad (16)$$

For any choice of $\boldsymbol{\lambda}$ the dual function provides a lower bound for the optimal value of the primal problem [7, Section 5.1.3]. Therefore, the *Lagrangian dual problem* looks for a *maximum* of the dual function over $\boldsymbol{\lambda}$:

$$d^* \doteq \max_{\boldsymbol{\lambda}} d(\boldsymbol{\lambda}) = \max_{\boldsymbol{\lambda}} \inf_{\mathbf{x}} \mathbf{x}^\top \mathbf{W}(\boldsymbol{\lambda}) \mathbf{x} + \lambda^y \quad (17)$$

The infimum over \mathbf{x} of $\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda})$ drifts to $-\infty$ unless $\mathbf{W}(\boldsymbol{\lambda}) \succeq 0$. Therefore we can safely restrict the maximization to vectors $\boldsymbol{\lambda}$ that are such that $\mathbf{W}(\boldsymbol{\lambda}) \succeq 0$; these are called *dual-feasible*. Moreover, at any dual-feasible $\boldsymbol{\lambda}$, the \mathbf{x} minimizing the Lagrangian are those that make $\mathbf{x}^\top \mathbf{W}(\boldsymbol{\lambda}) \mathbf{x} = 0$ (recall that $\mathbf{x}^\top \mathbf{W}(\boldsymbol{\lambda}) \mathbf{x}$ cannot be negative when $\mathbf{W}(\boldsymbol{\lambda}) \succeq 0$). Therefore, (17) reduces to the following *dual problem*

$$\begin{aligned} d^* = \max_{\boldsymbol{\lambda}} \quad & \lambda^y, && \text{(Dual problem)} \\ \text{s.t.:} \quad & \mathbf{W}(\boldsymbol{\lambda}) \succeq 0. && (18) \end{aligned}$$

The importance of the dual problem is twofold. First, it holds

$$d^* \leq f^* \quad (19)$$

This property is called *weak duality*, see, e.g., [7, Section 5.2.2]. For particular problems the inequality (19) becomes an equality, and in such cases we say that *strong duality* holds. The difference between f^* and d^* is called *duality gap* and, by (19), the gap is always positive.

The second important property of the Lagrangian dual problem (18) is that, since $d(\boldsymbol{\lambda})$ is concave (point-wise minimum

of affine functions), the dual (18) is always convex in λ , regardless the convexity properties of the primal problem. The dual PGO problem (18) is a semidefinite program (SDP) and we can compute the optimal solution λ^* of (18) using off-the-shelf solvers, e.g., [9, 10]. In the following section we discuss the relations between the primal and the dual problem, and elucidate on the practical use of the dual.

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

In this section we provide a more practical view of the results of [1], applying them to the extended formulation proposed in this addendum. We do not give proofs, since the inclusion of the determinant constraints only added extra terms in the definition of the matrix $\mathbf{W}(\lambda)$, hence the proof of the results stated here remains identical to the ones given in [1].

As stated in the introduction, our objective is two-fold. We want to design verification techniques and possibly find globally optimal solutions for PGO. Let us start with verification.

Verification. Assume that we computed an estimate for the poses in the pose graph, namely $(\hat{\mathbf{R}}_i, \hat{\mathbf{t}}_i)$, $i = 1, \dots, n$, using a state-of-the-art iterative solver, e.g., iSAM2 [11] or g2o [12]. Then we want to ask a basic question: *can we quantify how far is this estimate from the global minimum of the cost function, and possibly certify its optimality?*

We call this candidate solution $\hat{\mathbf{x}}$, assuming that the poses $(\hat{\mathbf{R}}_i, \hat{\mathbf{t}}_i)$, $i = 1, \dots, n$, have been “vectorized” as in Section III. The following result provides a first tool for verification.

Proposition 1 (Verification of Primal Objective): Given a candidate solution $\hat{\mathbf{x}}$ for the primal problem (12), and calling d^* the optimal objective of the dual problem (18), if $f(\hat{\mathbf{x}}) = d^*$, then the duality gap is zero and $\hat{\mathbf{x}}$ is an optimal solution of (12). Moreover, even if the duality gap is nonzero, $f(\hat{\mathbf{x}}) - d^* \geq f(\hat{\mathbf{x}}) - f^*$, meaning that $f(\hat{\mathbf{x}}) - d^*$ is an upper-bound for the sub-optimality gap of $\hat{\mathbf{x}}$.

Proposition 1 (cf. [1, Proposition 2]) ensures that the candidate $\hat{\mathbf{x}}$ is optimal when $f(\hat{\mathbf{x}}) = d^*$. Moreover, even in the case in which we get $f(\hat{\mathbf{x}}) > d^*$, the quantity $f(\hat{\mathbf{x}}) - d^*$ can be used as an indicator of how far $\hat{\mathbf{x}}$ is from the global optimum.

Our derivation also enables a more sophisticated verification technique. This technique, given in Proposition 3, is based on the following result (cf. [1, Lemma 1 and Proposition 3]).

Lemma 2 (Primal optimal solution and zero duality gap): If the duality gap is zero ($d^* = f^*$), then any primal optimal solution \mathbf{x}^* of (12) is in the null space of the matrix $\mathbf{W}(\lambda^*)$, where λ^* is the solution of the dual (18), i.e., $\mathbf{W}(\lambda^*)\mathbf{x}^* = \mathbf{0}$.

Proposition 3 (Verification of Primal Optimal Solution): Given a candidate solution $\hat{\mathbf{x}}$ for the primal problem (12), if the solution $\hat{\lambda}$ of the linear system

$$\mathbf{W}(\hat{\lambda})\hat{\mathbf{x}} = \mathbf{0} \quad (\text{to be solved w.r.t. } \hat{\lambda}) \quad (20)$$

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

This second verification technique is more convenient in practice, since it does not require solving the SDP (18), but

only requires solving a sparse linear system and then verifying that the sparse matrix $\mathbf{W}(\hat{\lambda})$ is positive definite (this can be checked by computing the smallest eigenvalue of $\mathbf{W}(\hat{\lambda})$).

Optimal solution. In [1], it has been shown that, when the duality gap is zero, we can compute a guaranteed, globally optimal solution to PGO. The result stems directly from Lemma 2 and can be stated as follows (cf. [1, Proposition 5]).

Proposition 4: If the duality gap is zero and λ^* is an optimal solution of (18), then an optimal solution \mathbf{x}^* of (12) can be computed by solving the following linear system:

$$\mathbf{W}(\lambda^*)\mathbf{x}^* = \mathbf{0} \quad (\text{to be solved w.r.t. } \mathbf{x}^*) \quad (21)$$

where the last entry in \mathbf{x}^* (corresponding to y) is fixed to 1.

Despite the appeal of having a (non-iterative) technique that computes a globally optimal solution for SLAM, the current use of (21) is pretty limited, as the computation of \mathbf{x}^* requires to compute λ^* by solving the SDP (18). While SDPs are convex problems (hence can be solved in polynomial time), current SDP solvers are fairly slow and do not scale to large problems, as the ones usually encountered in SLAM.

VI. EXPERIMENTS

Section V shows that when the duality gap is zero we are able to provide strong results on 3D PGO: we are able to design fast verification techniques that do not require solving the SDP underlying the dual problem (Proposition 3), and we can compute provably optimal solutions (Proposition 4).

In [1] we showed that the duality gap is zero in many large-scale real SLAM problems. Moreover, we provided a Monte Carlo analysis that highlights the influence of different parameters (noise level, number of poses) on the duality gap.

The goal of this experimental section is to compare the results of [1] and the extended formulation proposed in this addendum, and check if the addition of the determinant constraints “enlarged” the domain in which the duality gap is zero. We do this comparison on the same simulation setup of [1]. We use the cube dataset of Fig. 1(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.

In our experiments we compute the “optimal” solution f^* of (1) by refining the *chordal* initialization of [5, 13] 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*. For the solution of the dual problem (18), we use SDPA [10]. The solver returns the dual optimal objective d^* and the dual optimal variables λ^* .

Duality gap. Here we compare the value of the primal optimal objective f^* , against the value of the dual optimal




Fig. 1: (a1)-(a3): comparison between f^* , d^* (eq. (18)), and $d_{[1]}^*$ [1, eq. (22)], for different levels of translation noise σ_T (first column), rotation noise σ_R (second column), and size of the problem (third column). (a4): CPU time required to solve the SDP in (18) (dashed blue line) and in [1, eq. (22)] (solid red line). (b1)-(b3): comparison between f^* , the primal optimal solution computed according to Proposition 4 (“dual”), and the primal solution computed as in [1, Proposition 5] (“dual [1]”), (b4): cube dataset used for the Monte Carlo simulations.

objective d^* in (18), and the dual objective of the relaxed formulation in [1, eq. (22)], which we denote as $d_{[1]}^*$. Fig. 1(a1) shows $d_{[1]}^*$, d^* , and f^* for different translational noise levels, fixing $\sigma_R = 0.05\text{rad}$. The figure shows that $d_{[1]}^* = d^* = f^*$ (zero duality gap) independently on the translational noise. The duality gap was already observed to be zero using the formulation [1], hence the inclusion of the determinant constraints does not imply particular benefits. The same conclusion can be drawn from Fig. 1(a3), which compares $d_{[1]}^*$, d^* and f^* for different sizes of the cube dataset, fixing $\sigma_R = 0.05\text{rad}$ and $\sigma_T = 0.1\text{m}$. Again, $d_{[1]}^* = d^* = f^*$ (zero duality gap) independently of the size of the dataset, hence the relaxed formulation [1] already performed well.

Fig. 1(a2) shows $d_{[1]}^*$, d^* , and f^* for different rotational noise, fixing $\sigma_T = 0.1\text{m}$. In this case the duality gap $f^* - d^*$ is more sensitive to the noise level, and for large rotational noise both $d_{[1]}^*$ and d^* become smaller than f^* . While d^* is slightly better (i.e., larger) than $d_{[1]}^*$, the difference is small, and the inclusion of the determinant constraints does not significantly enlarge the domain in which the duality gap is zero.

Fig. 1(a4) reports the average CPU time required to solve the SDP in eq. (18) (dashed blue line) and the SDP in [1, eq. (22)] (solid red line). Interestingly, the computation cost of solving the complete dual problem (including the determinant constraints) is slightly smaller than the cost of the relaxed formulation [1]. Therefore, the inclusion of the determinant constraints does not imply an extra computational burden.

Primal optimal solution via the dual. In this paragraph we compare the performance of the primal optimal solution computed according to Proposition 4 against the primal solution computed according to [1, Proposition 5]. Recall that the

two formulations only differ by the extra terms we included in the matrix $W(\lambda^*)$ and which correspond to the determinant constraints. The ideal outcome of our analysis is that the extended formulation performs better, since it can find more instances in which the duality gap is zero and Proposition 4 is guaranteed to produce an optimal solution. Unfortunately, the following experiments show that this is not the case.

Figs. 1(b1)-(b3) compare the cost of the solution obtained from (i) Proposition 4 (label: “dual”), (ii) Proposition 5 in [1] (label: “dual [1]”), and the primal solution f^* . Fig. 1(b1) and Fig. 1(b3) compare the costs for increasing levels of translational noise and dataset size, respectively. We can see that the bars corresponding to the two dual problems lead to very similar costs, hence there is no clear advantage in including the determinant constraints in the formulation. We attribute the small mismatch between the heights of the bars on the left of Fig. 1(b1) to numerical errors, as inaccuracies in the solution of the SDP (18) (which may depend on the stopping conditions used in SDPA [10]) propagate to the primal solution computed according to Proposition 4.

Fig. 1(b2) compares the costs for increasing levels of rotation noise. In this case we see that the bars corresponding to the two dual problems become larger than f^* (i.e., they produce suboptimal solutions) for rotation noise larger than 0.1rad; this corresponds to the noise levels for which the duality gap becomes non-zero (cf. Fig. 1(a2)); indeed, when the duality gap is non-zero, the results in Proposition 4 and [1, Proposition 5] do not apply, hence we may get suboptimal solutions. Interestingly, the suboptimal solutions of Proposition 4 (which includes the determinant constraints) have larger cost with respect to the relaxed formulation [1, Proposition 5].

We conclude this section by remarking that, using Proposition 3, we can classify a candidate solution (e.g., distinguish optimal versus suboptimal estimates) whenever the duality gap is zero. To evaluate the efficacy of this classification, in [1], we provided an empirical analysis of the relaxed formulation by reporting precision-recall curves. We repeated the same tests for the extended formulation of this paper and we obtained exactly the same precision-recall curves, which essentially means that whenever the complete formulation proposed in this paper succeeded to certify optimality, also the relaxed formulation [1] was able to correctly classify the candidate. This again confirms that the introduction of the determinant constraints does not have significant impact in practice.

VII. DISCUSSION

Our results seem to suggest that the inclusion of the determinant constraints, in practical instances of PGO, do not provide any significant advantage from the point of view of the duality gap. This, in turn, implies that the determinant constraints are not active at the optimal (primal) solution, and that the orthogonality constraints are sufficient to define the optimal solution. We argue that this is due to the fact that, in practical instances of the problem (in which rotation measurement noise is reasonably small) one cannot obtain better solutions by flipping the sign of some rotations and turning them into reflections. Intuitively, this would happen only with highly self-inconsistent sets of measurements which contain a large number of outliers (for which the least-squares formulation in (1) might be not suitable). Nonetheless, our analysis shows that the relaxation of the determinant constraints, in practice, is not the determining factor to cause a non-zero duality gap.

VIII. CONCLUSION

This technical addendum shows how to include determinant constraints in the derivation of the dual SLAM problem, originally proposed in [1]. Determinant constraints can be formulated as quadratic equality constraints and do not fundamentally alter the structure of the primal and dual problem. As such, the results given in [1] apply to the extended formulation proposed in this paper without any significant modification.

While this extension completes the derivation given in [1], current experiments do not highlight any tangible advantage from adding the determinant constraints: the range of operation in which the duality gap is empirically zero remains practically the same for both formulations. We conclude that the relaxation of the determinant constraints is not a cause of failures in the certification of optimality.

Future work includes performing further numerical evaluations (e.g., using datasets corrupted by outliers), and investigating alternative methods to reduce/quantify the duality gap.

REFERENCES

[1] L. Carlone, D. Rosen, G. Calafiore, J.J. Leonard, and F. Dellaert. Lagrangian duality in 3D SLAM: Verification techniques and optimal solutions. In *the Intl. Conf. on*

Intelligent Robots and Systems, submitted, *ArXiv preprint available at: <http://arxiv.org/abs/1506.00746>*, 2015.

[2] L. Carlone and F. Dellaert. Duality-based verification techniques for 2D SLAM. In *Intl. Conf. on Robotics and Automation (ICRA)*, accepted, 2015.

[3] G. Calafiore, L. Carlone, and F. Dellaert. Pose graph optimization in the complex domain: Lagrangian duality, conditions for zero duality gap, and optimal solutions. In *arXiv preprint: 1505.03437*, 2015.

[4] R. Hartley, J. Trumpf, Y. Dai, and H. Li. Rotation averaging. *IJCV*, 103(3):267–305, 2013.

[5] 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.

[6] D.M. Rosen 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)*, 2015.

[7] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[8] G.C. Calafiore and L. El Ghaoui. *Optimization Models*. Cambridge University Press, 2014.

[9] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming. URL <http://cvxr.com/cvx>.

[10] M. Yamashita, K. Fujisawa, and M. Kojima. Implementation and evaluation of SDPA 6.0 (semidefinite programming algorithm 6.0). *Optimization Methods and Software*, 18(4):491–505, August 2003.

[11] 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*, 31:217–236, Feb 2012.

[12] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g2o: A general framework for graph optimization. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, May 2011.

[13] D. Martinec and T. Pajdla. Robust rotation and translation estimation in multiview reconstruction. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2007.