

3D LIDAR–camera intrinsic and extrinsic calibration: Identifiability and analytical least-squares-based initialization

The International Journal of Robotics Research 31(4) 452–467 © The Author(s) 2012 Reprints and permission: sagepub.co.uk/journalsPermissions.nav DOI: 10.1177/0278364911435689 ijr.sagepub.com



Faraz M Mirzaei, Dimitrios G Kottas and Stergios I Roumeliotis

Abstract

In this paper we address the problem of estimating the intrinsic parameters of a 3D LIDAR while at the same time computing its extrinsic calibration with respect to a rigidly connected camera. Existing approaches to solve this nonlinear estimation problem are based on iterative minimization of nonlinear cost functions. In such cases, the accuracy of the resulting solution hinges on the availability of a precise initial estimate, which is often not available. In order to address this issue, we divide the problem into two least-squares sub-problems, and analytically solve each one to determine a precise initial estimate for the unknown parameters. We further increase the accuracy of these initial estimates by iteratively minimizing a batch nonlinear least-squares cost function. In addition, we provide the minimal identifiability conditions, under which it is possible to accurately estimate the unknown parameters. Experimental results consisting of photorealistic 3D reconstruction of indoor and outdoor scenes, as well as standard metrics of the calibration errors, are used to assess the validity of our approach.

Keywords

Sensing and perception, computer vision, range sensing, calibration and identification

1. Introduction

As demonstrated in the Defense Advanced Research Projects Agency (DARPA) Urban Challenge, commercially available high-speed 3D LIDARs, such as the Velodyne, have made autonomous navigation and mapping within dynamic environments possible. In most applications, however, another sensor is employed in conjunction with the 3D LIDAR to assist in localization and place recognition. In particular, spherical cameras are often used to provide visual cues and to construct photorealistic maps of the environment. In these scenarios, accurate extrinsic calibration of the six degrees of freedom (d.o.f.) transformation between the two sensors is a prerequisite for optimally combining their measurements.

Several methods exist for calibrating a 2D laser scanner with respect to a camera. The work of Zhang and Pless (2004) relies on the observation of a planar checkerboard by both sensors. In particular, corners are detected in the images and planar surfaces are extracted from the laser measurements. The detected corners are used to determine the normal vector and distance of the planes where the laser-scan endpoints lie. Using this geometric constraint, the estimation of the transformation between the two sensors is formulated as a nonlinear least-squares problem and solved iteratively. A simplified linear least-squares solution is also provided to initialize the iterative nonlinear algorithm. More recently, Naroditsky et al. (2011) have presented a minimal approach for calibrating a 2D laser scanner with respect to a camera, using only six measurements of a planar calibration board. The computed transformation is then used in conjunction with RAndom Sample Consensus (RANSAC) (Fischler and Bolles 1981) to initialize an iterative least-squares refinement.

The existing 2D laser scanner–camera calibration methods have been extended to 3D LIDARs by Unnikrishnan and Hebert (2005) and Pandey et al. (2010). In both works, a geometric constraint similar to that presented by Zhang and Pless (2004) is employed to form a nonlinear least-squares cost function which is iteratively minimized to estimate the LIDAR–camera transformation. In addition, Unnikrishnan and Hebert (2005) have presented an initialization method for the iterative minimization based on a simplified

Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN, USA

Corresponding author:

Faraz M Mirzaei, Department of Computer Science and Engineering, University of Minnesota, 4-192 Keller Hall, 200 Union Street SE, Minneapolis, MN 55455, USA. Email: faraz@umn.edu linear least-squares formulation. Specifically, the estimation of relative rotation and translation are decoupled, and then each of them is computed from a geometric constraint between the planar segments detected in the measurements of both the 3D LIDAR and the camera. An alternative 3D LIDAR-camera calibration approach is described by Scaramuzza et al. (2007), where several point correspondences are manually selected in images and their associated LIDAR scans. Then, the Perspective n-point Pose (PnP) estimation algorithm of Quan and Lan (1999) is employed to find the transformation between the camera and the 3D LIDAR based on these point correspondences. In a different approach, presented by Stamos et al. (2008), the structural edges extracted from 3D LIDAR scans are matched with the vanishing points of the corresponding 2D images to compute a coarse 3D LIDAR-camera transformation, followed by an iterative least-squares refinement.

The main limitation of the above methods is that they assume the 3D LIDAR to be intrinsically calibrated. If the LIDAR's intrinsic calibration is not available or sufficiently accurate, then the calibration accuracy as well as the performance of subsequent LIDAR–camera data fusion significantly degrades. Pandey et al. (2010) have partially addressed this issue for the Velodyne 3D LIDAR by first calibrating only some of its intrinsic parameters. However, the suggested intrinsic calibration procedure is also iterative, and no method is provided for initializing it. While several of the intrinsic parameters of a LIDAR may be initialized using the technical drawings of the device (if available), other parameters, such as the offset in the range measurements induced by the delay in the electronic circuits, cannot be determined in this way.

To address these limitations, in this work we propose a novel algorithm for jointly estimating the intrinsic parameters of a revolving-head 3D LIDAR as well as the LIDAR– camera transformation. Specifically, we use measurements of a calibration plane at various configurations to establish geometric constraints between the LIDAR's intrinsic parameters and the LIDAR–camera 6 d.o.f. relative transformation. We process these measurement constraints to estimate the calibration parameters as follows: First, we analytically compute an initial estimate for the intrinsic and extrinsic calibration parameters in two steps. Subsequently, we employ a batch iterative (nonlinear) leastsquares method to refine the accuracy of the estimated parameters.

In particular, to analytically compute an initial estimate, we relax the estimation problem by seeking to determine the transformation between the camera and each one of the conic laser scanners within the LIDAR, along with its intrinsic parameters. As a first step, we formulate a nonlinear least-squares problem to estimate the 3 d.o.f. rotation between each conic laser scanner and the camera, as well as a subset of the laser scanner's intrinsic parameters. The optimality conditions of this nonlinear least-squares problem form a system of polynomial equations, which we solve analytically using an algebraic-geometry approach to find all of its critical points. Amongst these, the one that minimizes the least-squares cost function corresponds to the global minimum and provides us with the initial estimates for the relative rotation and the first set of intrinsic LIDAR parameters. In the next step, we use a linear least-squares algorithm to compute the initial estimate for the relative

and the remaining intrinsic parameters. Once all initial estimates are available, we finally perform a batch iterative joint optimization of the LIDAR–camera transformation and the LIDAR's intrinsic parameters. As part of our contributions, we also study the identifiability properties of the problem and present the minimal necessary conditions for concurrently estimating the LIDAR's intrinsic parameters and the LIDAR–camera transformation. Our experimental results demonstrate that our proposed method significantly improves the accuracy of the intrinsic calibration parameters of the LIDAR, as well as the LIDAR–camera transformation.

translation between the camera and the conic laser scanners,

The remainder of this paper is structured as follows. The calibration problem is formulated in Section 2, and the proposed solution is presented in Section 3. In Section 4 the identifiability of the problem is investigated and in Section 5, an experimental comparison of our method with the approach of Pandey et al. (2010) is provided, and photorealistic 3D reconstruction of indoor and outdoor scenes using the estimated calibration parameters are presented. Finally, in Section 6, the conclusions of this work are drawn, and suggestions for future research directions are provided.

2. Problem formulation

A revolving-head 3D LIDAR consists of K conic laser scanners mounted on a rotating head so that they span a 360° panoramic (azimuth) view (see Figure 1). Each laser scanner has a horizontal offset from the axis of rotation, and a vertical offset from adjacent laser scanners. In addition, each laser scanner points to a different elevation angle, such that, collectively, all of the laser scanners cover a portion of the vertical field of view. Therefore, once the LIDAR's head completes a full rotation, each laser scanner has swept a cone in space specified by its elevation angle. Let $\{L\}$ be the LIDAR's fixed frame of reference whose z-axis is the axis of rotation of the sensor's head (see Figure 1). Also, let $\{L_i\}, i = 1, \ldots, K$, be the coordinate frame corresponding to the *i*th laser scanner, such that its origin is at the center of the associated cone on the z-axis of $\{L\}$ with vertical offset h_i from the origin of $\{L\}$, its z-axis aligned with that of $\{L\}$, and its x-axis defining an angle θ_{oi} with the x-axis of $\{L\}$.

We determine the direction of the *k*th shot of the *i*th laser beam from its corresponding elevation angle, ϕ_i , and azimuth measurement, θ_{ik} , and denote it with

$${}^{L_i}\bar{\mathbf{p}}_k \triangleq \begin{bmatrix} \cos\phi_i \cos\phi_{ik} \\ \cos\phi_i \sin\phi_{ik} \\ \sin\phi_i \end{bmatrix}.$$
(1)



Fig. 1. A revolving-head 3D LIDAR consists of *K* laser scanners, pointing to different elevation angles, and rotating around a common axis. The intrinsic parameters of the LIDAR describe the measurements of each laser scanner in its coordinate frame, $\{L_i\}$, and the transformation between the LIDAR's fixed coordinate frame, $\{L\}$, and $\{L_i\}$. Note that besides the physical offset of the laser scanners from the axis of rotation, the value of ρ_{oi} may depend on the delay in the electronic circuits of the LIDAR.

The distance measured by the *k*th shot of the *i*th laser scanner is represented by ρ_{ik} . The real distance to the object that reflects the *k*th shot of the *i*th laser beam is $\alpha_i(\rho_{ik} + \rho_{oi})$, where α_i is the scale factor, and ρ_{oi} is the range offset due to the delay in the electronic circuits of the LIDAR and the offset of each laser scanner from its cone's center. In this way, the position of the *k*th point measured by the *i*th laser scanner is described by

$${}^{L_i}\mathbf{p}_{ik} = \alpha_i(\rho_{ik} + \rho_{oi}){}^{L_i}\bar{\mathbf{p}}_k.$$
 (2)

The transformation between $\{L_i\}$ and $\{L\}$ (i.e. h_i and θ_{oi}), the scale α_i , offset ρ_{oi} , and elevation angle ϕ_i , for i = 1, ..., K, comprise the intrinsic parameters of the LIDAR that must be precisely known for any application, including photorealistic reconstruction of the surroundings. Since the intrinsic parameters supplied by the manufacturer are typically not accurate (except for the elevation angle ϕ_i), in this work we estimate them along with the transformation with respect to a camera.¹

We assume that an *intrinsically calibrated* camera is rigidly connected to the LIDAR, and our objective is to determine the 6 d.o.f. relative transformation between the two, as well as the intrinsic parameters of the LIDAR. For this purpose, we employ a planar calibration board with fiducial markers² at M different configurations to establish geometric constraints between the measurements of the LIDAR and the camera, their relative transformation, and the LIDAR's intrinsic parameters.

2.1. Noise-free geometric constraints

At the *j*th configuration of the calibration board, j = 1, ..., M, (see Figure 2), the fiducial markers whose positions are known with respect to the calibration board's frame of reference $\{B_j\}$, are first detected in the camera's image. The 6 d.o.f. transformation between $\{C\}$ and $\{B_j\}$ is then computed using a PnP algorithm (Quan and Lan 1999; Ansar and Daniilidis 2003; Hesch and Roumeliotis 2011),



Fig. 2. Geometric constraint between the *j*th plane, the camera $\{C\}$, and the *i*th laser scanner, $\{L_i\}$. Each laser beam is described by a vector ${}^{L_i}\mathbf{p}_{ijk}$. The plane is described by its normal vector ${}^{C}\mathbf{n}_j$ and its distance d_i both expressed with respect to the camera.

from which the normal vector and the distance of the target plane in the camera's frame are extracted as

$${}^{C}\bar{\mathbf{n}}_{j} \triangleq {}^{C}_{B_{j}} \mathbf{C} \begin{bmatrix} 0 & 0 & -1 \end{bmatrix}^{\mathrm{T}}$$
(3)

$$l_j \triangleq {}^C \bar{\mathbf{n}}_j^{\mathrm{T}C} \mathbf{t}_{B_j} \tag{4}$$

where ${}_{B_j}^C \mathbf{C}$ and ${}^{C} \mathbf{t}_{B_j}$ represent the relative rotation and translation between the camera and the calibration board at the *j*th configuration. Consequently, in the absence of noise, any point ${}^{C}\mathbf{p}$ that lies on the *j*th plane satisfies

$${}^{C}\bar{\mathbf{n}}_{j}^{\mathrm{T}\,C}\mathbf{p}-d_{j}=0.$$
(5)

We now turn our attention to the LIDAR point measurements reflected from the *j*th calibration plane and identified based on the depth discontinuity. Let us denote such points as ${}^{L_i}\mathbf{p}_{ijk}$, $k = 1..., N_{ij}$, measured by the LIDAR's *i*th laser scanner (see (2)). Transforming these points to the camera's frame, and substituting them in (5) yields:

$${}^{C}\bar{\mathbf{n}}_{j}^{\mathrm{T}}\left({}^{C}_{L_{i}}\mathbf{C}^{L_{i}}\mathbf{p}_{ijk}+{}^{C}\mathbf{t}_{L_{i}}\right)-d_{j}=0 \xrightarrow{(2)}$$
(6)

$$\alpha_i(\rho_{ijk}+\rho_{oi})^C \bar{\mathbf{n}}_{j\ L_i}^{\mathrm{T}\ C} \mathbf{C}^{L_i} \bar{\mathbf{p}}_{ijk} + {}^C \bar{\mathbf{n}}_{j\ C}^{\mathrm{T}\ C} \mathbf{t}_{L_i} - d_j = 0 \quad (7)$$

where ${}_{L_i}^C \mathbf{C}$ and ${}^C \mathbf{t}_{L_i}$ are the relative rotation and translation between the camera and the *i*th laser scanner.

2.2. Geometric constraints in the presence of noise

In the presence of noise, the geometric constraint in (7) is not exactly satisfied. Instead, we will have

$$\boldsymbol{\chi}_{i}(\rho_{ijk}+\rho_{oi})^{C}\bar{\mathbf{n}}_{j\ L_{i}}^{TC}\mathbf{C}^{L_{i}}\bar{\mathbf{p}}_{ijk}+{}^{C}\bar{\mathbf{n}}_{j}^{TC}\mathbf{t}_{L_{i}}-d_{j}=\epsilon_{ijk} \qquad (8)$$

where ϵ_{ijk} is the residual due to the noise in the image and the LIDAR measurements. The covariance of this residual is

$$\sigma_{\epsilon_{ijk}}^{2} = \left(\alpha_{i}^{C} \bar{\mathbf{n}}_{j}^{T} {}_{Li}^{C} \mathbf{C}^{L_{i}} \bar{\mathbf{p}}_{ijk}\right)^{2} \sigma_{\rho}^{2} + \mathbf{h}_{p_{ijk}} \mathbf{R}_{p} \mathbf{h}_{p_{ijk}}^{T} + \mathbf{h}_{n} \mathbf{R}_{n_{j}} \mathbf{h}_{n}^{T} + \sigma_{d_{j}}^{2}$$

$$(9)$$

where σ_{ρ} is the standard deviation of noise in ρ_{ijk} , and σ_{dj} is the standard deviation of the uncertainty in d_j . The covariance of the uncertainty in the laser beam directions, ${}^{L_i}\bar{\mathbf{p}}_{ijk}$, and the plane normal vector, ${}^{C}\bar{\mathbf{n}}_{j}$, expressed in their local tangent planes (see Hartley and Zisserman 2004, Appendix 6.9.3), are represented by \mathbf{R}_p and \mathbf{R}_{nj} , respectively. The corresponding Jacobians are

$$\mathbf{h}_{p_{ijk}} = \alpha_i (\rho_{ijk} + \rho_{oi})^C \bar{\mathbf{n}}_{j\ L_i}^T \mathbf{C} \mathbf{C} \mathbf{H}_{\bar{\mathbf{p}}_{ijk}} \begin{bmatrix} \mathbf{I}_2 & \mathbf{0}_{2 \times 1} \end{bmatrix}^T$$
(10)

$$\mathbf{h}_{n_j} = \begin{pmatrix} \alpha_i (\rho_{ijk} + \rho_{oi})^{L_i} \bar{\mathbf{p}}_{ijk\ L_i}^{\mathrm{T}\ C} \mathbf{C}^{\mathrm{T}} + {}^C \mathbf{t}_{L_i}^{\mathrm{T}} \end{pmatrix} \mathbf{H}_{\bar{\mathbf{n}}_j} \begin{bmatrix} \mathbf{I}_2 & \mathbf{0}_{2 \times 1} \end{bmatrix}^{\mathrm{T}}$$
(11)

where $\mathbf{H}_{\bar{\mathbf{u}}}$ is the 3 × 3 Householder matrix associated with the unit vector $\bar{\mathbf{u}}$ (Golub and van Loan 1996), defined as

$$\mathbf{H}_{\tilde{\mathbf{u}}} \triangleq \mathbf{I}_3 - 2\frac{\mathbf{v}\mathbf{v}^{\mathrm{T}}}{\mathbf{v}^{\mathrm{T}}\mathbf{v}}, \quad \mathbf{v} \triangleq \tilde{\mathbf{u}} + \operatorname{sign}(\mathbf{e}_3^{\mathrm{T}}\tilde{\mathbf{u}})\mathbf{e}_3, \quad \mathbf{e}_3 \triangleq [0\ 0\ 1]^{\mathrm{T}}.$$
(12)

Note that the characteristics of ϵ_{ijk} depend not only on the uncertainty of the measurements, but also on the unknown calibration parameters.

2.3. Structural constraints

In addition to the camera and laser scanner measurements, the following constraints can also be used to increase the accuracy of the calibration process. Specifically, since the *z*-axis of $\{L_i\}$ is aligned with the *z*-axis of $\{L\}$, while their *x*-axes form an angle θ_{oi} , the following constraint holds for all $_{Li}^{C}$ C:

$${}^{C}_{Li}\mathbf{C} = {}^{C}_{L}\mathbf{C}\mathbf{C}_{z}(\theta_{oi}) \tag{13}$$

where $C_z(\theta_{oi})$ represents a rotation around the *z*-axis by an angle θ_{oi} . In addition, the origin of each laser-scanner frame lies on the *z*-axis of {*L*} with vertical offset of h_i from the origin of {*L*}, resulting in the following constraint:

$${}_{L}^{C}\mathbf{C}^{\mathrm{T}}\left({}^{C}\mathbf{t}_{L_{i}}-{}^{C}\mathbf{t}_{L}\right)=\left[0\ 0\ h_{i}\right]^{\mathrm{T}}.$$
(14)

3. Algorithm description

In order to estimate the unknown calibration parameters, we form a constrained nonlinear least-squares cost function from the residuals of the geometric constraint over all point and plane observations (see (8)). To minimize this least-squares cost, one has to employ iterative minimizers such as the Levenberg–Marquardt (Press et al. 1992), that Algorithm 1 Estimate intrinsic LIDAR and extrinsic LIDAR–camera calibration parameters.

- 1: for *j*th configuration of the calibration plane do
- 2: Record an image and a LIDAR snapshot.
- 3: Detect the known fiducial markers on the image.
- 4: Compute ${}^{C}\bar{\mathbf{n}}_{i}$ and d_{i} using a PnP algorithm.
- 5: **for** *i*th laser scanner **do**
- 6: Identify the laser points hitting the calibration plane using depth discontinuity.
- 7: Compute the contributions of *j*th plane's observation to the rotation-offset optimality equations (see (23)).
- 8: end for

9: end for

- 10: **for** *i*th laser scanner **do**
- 11: Solve the optimality equations in (23) to compute critical points of (21).
- 12: Estimate $\hat{\mathbf{s}}_i$ and $\hat{\rho}_{oi}$ as the critical point that minimizes (22).
- 13: Solve the linear least-squares problem in (31) to estimate ${}^{C}\mathbf{t}_{L_{i}}$ and α_{i} .
- 14: **end for**
- 15: Refine the estimates for all of the unknowns and enforce (13) and (14) by iteratively minimizing (32).

require a precise initial estimate to ensure convergence. To provide accurate initialization, in this section we present a novel analytical method to estimate the LIDAR–camera transformation and all intrinsic parameters of the LIDAR (except the elevation angle ϕ_i which is precisely known from the manufacturer). In order to reduce the complexity of the initialization process, we temporarily drop the constraints in (13) and (14) and seek to determine the transformation between the camera and each of the laser scanners (along with each scanner's intrinsic parameters) independently (see Sections 3.1–3.3). Once an accurate initial estimate is computed, we lastly perform an iterative nonlinear least-squares refinement that explicitly considers (13) and (14), and increases the calibration accuracy (see Section 3.4).

3.1. Analytical estimation of offset and relative rotations

Note that the term ${}^{C}\bar{\mathbf{n}}_{j}^{T C}\mathbf{t}_{L_{i}} - d_{j}$ in (7) is constant for all points *k* of the *i*th laser scanner that hit the calibration plane at its *j*th configuration. Therefore, subtracting two noise-free constraints of the form (7) for the points ${}^{L_{i}}\mathbf{p}_{ijk}$ and ${}^{L_{i}}\mathbf{p}_{ijl}$, and dividing the result by the non-zero scale, α_{i} , yields

$${}^{C}\bar{\mathbf{n}}_{j\ L_{i}}^{\mathrm{T}\ C}\mathbf{C}(\mathbf{u}_{jkl}^{i}+\rho_{oi}\mathbf{v}_{jkl}^{i})=0$$
(15)

where $\mathbf{u}_{jkl}^{i} \triangleq \rho_{ijk}^{L_{i}} \bar{\mathbf{p}}_{ijk} - \rho_{ijl}^{L_{i}} \bar{\mathbf{p}}_{ijl}$ and $\mathbf{v}_{jkl}^{i} \triangleq {}^{L_{i}} \bar{\mathbf{p}}_{ijk} - {}^{L_{i}} \bar{\mathbf{p}}_{ijl}$. Note that this constraint involves as unknowns only the relative rotation of the *i*th laser scanner with respect to the camera, ${}^{L_{i}}_{L_{i}}\mathbf{C}$, and its offset, ρ_{oi} . Let us express the former, $_{L_i}^{C}$ **C**, using the Cayley–Gibbs–Rodriguez (CGR) parameterization (Shuster 1993), i.e.

$$_{L_i}^{C} \mathbf{C}(\mathbf{s}) = \frac{\mathbf{C}(\mathbf{s}_i)}{1 + \mathbf{s}_i^{\mathrm{T}} \mathbf{s}_i}$$
(16)

$$\bar{\mathbf{C}}(\mathbf{s}_i) \triangleq \left((1 - \mathbf{s}_i^{\mathrm{T}} \mathbf{s}_i) \mathbf{I}_3 + 2\lfloor \mathbf{s}_i \times \rfloor + 2\mathbf{s}_i \mathbf{s}_i^{\mathrm{T}} \right)$$
(17)

where $\mathbf{s}_i^{\mathrm{T}} = [s_{i1} \ s_{i2} \ s_{i3}]$ is the vector of CGR parameters that represent the relative orientation of the *i*th laser scanner with respect to the camera, and

$$\lfloor \mathbf{s} \times \rfloor \triangleq \begin{bmatrix} 0 & -s_3 & s_2 \\ s_3 & 0 & -s_1 \\ -s_2 & s_1 & 0 \end{bmatrix}$$
(18)

is the corresponding skew-symmetric matrix. Substituting (16) in (15), and multiplying both sides with the non-zero term $1 + \mathbf{s}_i^{\mathrm{T}} \mathbf{s}_i$ yields

$${}^{C}\bar{\mathbf{n}}_{j}^{\mathrm{T}}\bar{\mathbf{C}}(\mathbf{s}_{i})(\mathbf{u}_{jkl}^{i}+\rho_{oi}\mathbf{v}_{jkl}^{i})=0.$$
(19)

This algebraic constraint holds exactly in the absence of noise. In that case, the method presented in Section 4 can be employed to recover the unknowns given the minimum number of measurements. In the presence of noise, however, (19) becomes

$${}^{C}\bar{\mathbf{n}}_{j}^{\mathrm{T}}\bar{\mathbf{C}}(\mathbf{s}_{i})\left(\mathbf{u}_{jkl}^{i}+\rho_{oi}\mathbf{v}_{jkl}^{i}\right)=\eta_{jkl}^{i}$$
(20)

where η_{jkl}^i is a non-zero residual. In this case, we estimate \mathbf{s}_i and ρ_{oi} by solving the following nonlinear least-squares problem:

$$\hat{\mathbf{s}}_{i}, \hat{\rho}_{oi} = \min_{\substack{\mathbf{s}_{i}, \rho_{oi}}} \mathcal{C}_{i} \tag{21}$$

$$\mathcal{C}_{i} \triangleq \frac{1}{2} \sum_{j=1}^{M} \sum_{k=1}^{\frac{N_{ij}}{2}} \sum_{l=\frac{N_{ij}}{2}+1}^{N_{ij}} \left({}^{C} \bar{\mathbf{n}}_{j}^{\mathrm{T}} \bar{\mathbf{C}}(\mathbf{s}_{i}) \left(\mathbf{u}_{jkl}^{i} + \rho_{oi} \mathbf{v}_{jkl}^{i} \right) \right)^{2}$$
(22)

where, without loss of generality, we have assumed that N_{ij} is even. Note that the N_{ij} points from the *i*th laser scanner, and the *j*th configuration of the calibration plane are divided into two mutually exclusive groups so as to ensure that each point appears in the least-squares cost only once and hence avoid noise correlations.

When a sufficient number of plane configurations are observed, we employ a recently proposed algebraic method to directly solve this nonlinear least-squares problem without requiring initialization (Trawny and Roumeliotis 2010). Specifically, we first form the following polynomial system describing the optimality conditions of (21):

$$f_{\ell\ell} = \frac{\partial C_i}{\partial s_{\ell\ell}} = 0, \quad \ell = 0, \dots, 3, \quad \text{and} \quad s_{\ell} \triangleq \rho_{oi}.$$
 (23)

Note that the cost function in (22) is a polynomial of degree six in the elements of \mathbf{s}_i and ρ_{oi} . Therefore, (23) consists of four polynomials of degree five in four variables. The

detailed derivation of these polynomials is provided in Appendix A. This polynomial system has 243 solutions that comprise the critical points of the least-squares cost function C_i , and can be computed using the eigenvalue decomposition of the so-called multiplication matrix (see Section 3.2). The globally optimal solution of the least-squares problem is the critical point that minimizes (22), and it is selected through direct evaluation of the cost function C_i . We point out that the computational complexity of solving (23) and finding the global minimum does not increase with the addition of measurements, since the degree and number of polynomials expressing the optimality conditions are fixed regardless of the number of calibration-plane configurations and laser-scanner points reflected from them. Moreover, computing the contribution of all points to the coefficients of the polynomials $f_{i\ell}$, $\ell = 0, ..., 3$, increases only linearly with the number of measurements.

3.2. Polynomial system solver

Several methods exist for solving the polynomials describing the optimality conditions of (23). Amongst them, numerical methods, such as Newton-Raphson, need proper initialization and may not find all of the solutions. Symbolic reduction methods based on the computation of the system's Gröbner basis are capable of finding all roots without any initialization (Cox et al. 2004). However, they can only be used for integer or rational coefficients since their application to floating-point numbers suffers from quick accumulation of round-off errors, which, in turn, results in incorrect solutions (Cox et al. 2004). Instead, we employ a method developed by Auzinger and Stetter (1988) that computes a generalization of the companion matrix to systems of multivariate polynomial equations, namely the multiplication matrix, whose eigenvalues are the roots of the associated polynomial system. In the following, we briefly describe an efficient method for computing the multiplication matrix. In Appendix C, a simple example is provided to demonstrate how this method is applied in practice.

Let us denote a monomial in $\mathbf{x} = [x_1 \cdots x_n]^T$ by $x^{\mathbf{y}} \triangleq x_1^{\gamma_1} x_2^{\gamma_2} \cdots x_n^{\gamma_n}, \ \gamma_i \in \mathbb{Z}_{\geq 0}$, with degree $\sum_{i=1}^n \gamma_i$. A polynomial of degree d in \mathbf{x} is denoted by $f = \mathbf{c}^T \mathbf{x}_d$ where \mathbf{x}_d is the $\binom{n+d}{n}$ -dimensional vector of monomials of degree up to and including d, and \mathbf{c} is the vector of coefficients of equal size. We assume that the given system of equations has n polynomials, denoted by $f_i = \mathbf{c}_i^T \mathbf{x}_{d_i} = 0, \ i = 1, \dots, n$, each of them with degree d_i . The total degree of the polynomial system is $d \triangleq \max_i d_i$. By padding the coefficient vectors of f_i with zeros, and stacking them together in \mathbf{C} , we can present the polynomial system in the matrix form of $\mathbf{C}\mathbf{x}_d = \mathbf{0}$.

A system of polynomial equations defines an *ideal I* as the set of all of the polynomials that can be generated as $\sum_i f_i h_i$ where h_i is any polynomial in **x**. Clearly the elements of the ideal become zero at the solutions of the original (generator) polynomial system. The Gröbner basis $G \triangleq \langle g_1, \ldots g_l \rangle$ of an ideal is a finite subset of the ideal such that (i) the remainder of the division of any polynomial to it is unique, (ii) any polynomial whose division by the Gröbner basis results in zero remainder, is a member of the associated ideal. The first property can be expressed as

$$\varphi(\mathbf{x}) = r(\mathbf{x}) + \sum_{i=1}^{t} g_i(\mathbf{x}) h_i(\mathbf{x})$$
(24)

where φ is any polynomial in **x**, h_i are the quotient polynomials, and r is the unique remainder. We hereafter use the name 'remainder' as the remainder of the division of a polynomial by the Gröbner basis. The Gröbner basis for an ideal generated from polynomials with integer or rational numbers can be computed using implementations of the so-called Buchberger's algorithm (Cox et al. 2004) on symbolic software packages such as Macaulay2 or Maple. Computation of the Gröbner basis for polynomials with floating-point coefficients is much more difficult due to quick accumulation of round-off errors in the Buchberger's algorithm.

The remainders of the polynomials that are not in an ideal are instrumental in finding the solutions (i.e. variety) of that ideal. It can be shown that all such remainders can be expressed as a linear combination of a specific (unique) group of monomials that comprise the so-called *normal set* (Cox et al. 2004). The normal set can be easily obtained from the Gröbner basis of an ideal, and under mild conditions,³ its cardinality equals the number of solutions (real and complex) of the ideal, and it will contain the monomial 1 (Cox et al. 2004, p.43). The important point here is that the normal set is generically fixed across different instantiations of the polynomials. Therefore, we can compute the normal set of an instance of the problem (e.g. integer or rational coefficients) and use it when the coefficients are floating point.

Let us assume that the cardinality of the normal set is *s*, and represent its monomials in a vector form \mathbf{x}_B . Then multiplication of \mathbf{x}_B with a generic polynomial $\varphi(\mathbf{x})$ yields (see (24)):

$$\varphi(\mathbf{x}) \cdot \mathbf{x}_{B} = \mathbf{M}_{\varphi} \mathbf{x}_{B} + \begin{bmatrix} h_{11} & \cdots & h_{1t} \\ \vdots & & \vdots \\ h_{s1} & \cdots & h_{st} \end{bmatrix} \begin{bmatrix} g_{1} \\ \vdots \\ g_{t} \end{bmatrix}$$
(25)

where h_{ij} are polynomials in **x** and g_i are the elements of the Gröbner basis. In this expression, \mathbf{M}_{φ} is called the *multiplication matrix* associated with φ . This relationship holds since the remainder of any polynomial (including $x^{\gamma}\varphi(\mathbf{x})$, $x^{\gamma} \in \mathbf{x}_{B}$) can be written as a linear combination of \mathbf{x}_{B} . Now, if we evaluate (25) at $\mathbf{x} = \mathbf{p}$, a solution of the ideal, all g_i become zero, and we obtain

$$\varphi(\mathbf{p}) \cdot \mathbf{p}_B = \mathbf{M}_{\varphi} \mathbf{p}_B \tag{26}$$

where \mathbf{p}_B is \mathbf{x}_B evaluated at \mathbf{p} . Clearly, \mathbf{p}_B is an eigenvector of \mathbf{M}_{φ} , and $\varphi(\mathbf{p})$ is the associated eigenvalue. Therefore, if we select $\varphi(\mathbf{x})$ equal to one of the variables (e.g. x_i), we

can read off the x_i -coordinate of the solutions as the eigenvalues of \mathbf{M}_{φ} . Furthermore, depending on the ordering of the monomials when computing the Gröbner basis, \mathbf{x}_B may include all first-order monomials x_1, \ldots, x_n . In that case, one can simultaneously read off all of the coordinates of the solutions, for an arbitrary choice of φ , as long as it is non-zero and distinct at each solution of the ideal.

When the Gröbner basis is available (such as in polynomial systems with integer coefficients), one can use it directly to compute remainders of $\varphi(\mathbf{x}) \cdot \mathbf{x}_B$, and construct \mathbf{M}_{φ} . This is not possible, however, when working with polynomials with floating-point coefficients. Therefore we employ the method proposed by Byröd et al. (2008) to compute \mathbf{M}_{φ} . We first note that some of the monomials of $\varphi(\mathbf{x}) \cdot \mathbf{x}_B$ remain in \mathbf{x}_B , while some others do not. We form the vector \mathbf{x}_R from the latter monomials, and write

$$\varphi(\mathbf{x}) \cdot \mathbf{x}_{B} = \mathbf{M}_{\varphi}' \begin{bmatrix} \mathbf{x}_{R} \\ \mathbf{x}_{B} \end{bmatrix}$$
(27)

where \mathbf{M}'_{φ} is called the unreduced multiplication matrix. Our objective is to express the remainders of \mathbf{x}_R as a linear combination of \mathbf{x}_B without using the Gröbner basis. For this purpose, we expand each original polynomial f_i by multiplying it with all of the monomials up to degree $\ell - d_i$ (ℓ to be determined later). Clearly all of these new expanded polynomials belong to the ideal generated by the original polynomials, and they have monomials up to degree ℓ . Thus, we can write them collectively in matrix form as $\mathbf{C}_e \mathbf{x}_{\ell} = 0$. We reorder \mathbf{x}_{ℓ} and \mathbf{C}_e as

$$\mathbf{C}_{e}\mathbf{x}_{\ell} = \begin{bmatrix} \mathbf{C}_{E} & \mathbf{C}_{R} & \mathbf{C}_{B} \end{bmatrix} \begin{bmatrix} \mathbf{x}_{E} \\ \mathbf{x}_{R} \\ \mathbf{x}_{B} \end{bmatrix} = 0$$
(28)

where \mathbf{x}_{E} are the monomials that belong neither to \mathbf{x}_{R} nor to \mathbf{x}_{B} . Multiplying (28) with \mathbf{N}^{T} , the left null space of \mathbf{C}_{E} , and decomposing $\mathbf{N}^{T}\mathbf{C}_{R} = \mathbf{Q}\mathbf{R} = [\mathbf{Q}_{1} \ \mathbf{Q}_{2}]\begin{bmatrix} \mathbf{R}_{1} \\ \mathbf{0} \end{bmatrix}$ using QR factorization, yields:

$$\begin{bmatrix} \mathbf{N}^{\mathrm{T}}\mathbf{C}_{R} & \mathbf{N}^{\mathrm{T}}\mathbf{C}_{B} \end{bmatrix} \begin{bmatrix} \mathbf{x}_{R} \\ \mathbf{x}_{B} \end{bmatrix} = \mathbf{Q} \begin{bmatrix} \mathbf{R}_{1} & \mathbf{Q}_{1}^{\mathrm{T}}\mathbf{N}^{\mathrm{T}}\mathbf{C}_{B} \\ \mathbf{0} & \mathbf{Q}_{2}^{\mathrm{T}}\mathbf{N}^{\mathrm{T}}\mathbf{C}_{B} \end{bmatrix} \begin{bmatrix} \mathbf{x}_{R} \\ \mathbf{x}_{B} \end{bmatrix} = \mathbf{0}.$$
(29)

If ℓ is selected sufficiently large, \mathbf{R}_1 will be full rank (Reid and Zhi 2009), which allows us to solve (29) and find \mathbf{x}_R as a function of \mathbf{x}_B , i.e. $\mathbf{x}_R = -\mathbf{R}_1^{-1}\mathbf{Q}_1^{\mathrm{T}}\mathbf{N}^{\mathrm{T}}\mathbf{C}_B\mathbf{x}_B$. Substituting this relationship in (27) yields the multiplication matrix:

$$\mathbf{M}_{\varphi} = \mathbf{M}_{\varphi}' \begin{bmatrix} \mathbf{I}_{s} \\ -\mathbf{R}_{1}^{-1} \mathbf{Q}_{1}^{\mathrm{T}} \mathbf{N}^{\mathrm{T}} \mathbf{C}_{B} \end{bmatrix}.$$
 (30)

For solving equations (23), we had to expand the polynomials up to degree $\ell = 15$ and arrived at a multiplication matrix \mathbf{M}_{φ} of dimensions 243 × 243. Finally, we mention that it is possible to compute the multiplication matrix without explicit computation of the normal set. Further details on this subject and also on possible numerical instabilities and their remedies are given by Byröd et al. (2008); Reid and Zhi (2009); Trawny et al. (2009).

3.3. Analytical estimation of scale and relative translation

Once the relative rotation, $_{L_i}^{C}\mathbf{C}$, and offset, ρ_{oi} , of each laser scanner, $i = 1, \ldots, K$, is computed, we use linear least-squares to determine the relative translation and scale from (7). Specifically, we stack together all of the measurement constraints on the *i*th laser scanner's scale and relative translation (from different points and calibration-plane configurations), and write them in a matrix form as

$$\begin{bmatrix} {}^{C}\bar{\mathbf{n}}_{1}^{\mathrm{T}} & (\rho_{i11} + \rho_{oi}) {}^{C}\bar{\mathbf{n}}_{1}^{\mathrm{T}} {}^{C}_{L_{i}} \mathbf{C}^{L_{i}} \bar{\mathbf{p}}_{i11} \\ {}^{C}\bar{\mathbf{n}}_{1}^{\mathrm{T}} & (\rho_{i12} + \rho_{oi}) {}^{C}\bar{\mathbf{n}}_{1}^{\mathrm{T}} {}^{C}_{L_{i}} \mathbf{C}^{L_{i}} \bar{\mathbf{p}}_{i12} \\ \vdots & \vdots \\ {}^{C}\bar{\mathbf{n}}_{M}^{\mathrm{T}} & (\rho_{iMN_{iM}} + \rho_{oi}) {}^{C}\bar{\mathbf{n}}_{M}^{\mathrm{T}} {}^{C}_{L_{i}} \mathbf{C}^{L_{i}} \bar{\mathbf{p}}_{iMN_{iM}} \end{bmatrix} \begin{bmatrix} {}^{C}\mathbf{t}_{L_{i}} \\ \alpha_{i} \end{bmatrix} = \begin{bmatrix} d_{1} \\ d_{1} \\ \vdots \\ d_{M} \end{bmatrix}.$$
(31)

Under the condition that the coefficient matrix on the lefthand side of this equality is full rank (see Section 5), we can easily obtain the *i*th laser scanner's scale factor, α_i , and relative translation, ${}^{C}\mathbf{t}_{L_i}$, by solving (31).

3.4. Iterative refinement

Once the initial estimates for the transformation between the camera and the laser scanners, and the intrinsic parameters of the LIDAR are known (Sections 3.1–3.3), we employ an iterative refinement method to enforce the constraints in (13) and (14). Specifically, we choose the coordinate frame of one of the laser scanners (e.g. the first laser scanner) as the LIDAR's fixed coordinate frame, i.e. $\{L\} = \{L_1\}$. Then for $\{L_i\}$, i = 2, ..., K, we employ the estimated relative transformation with respect to the camera (i.e. ${}_{L_i}^C$ C and ${}^C \mathbf{t}_{L_i}$) to obtain the relative transformations between $\{L_i\}$ and $\{L\}$. From these relative transformations, we only use the *z* component of the translation to initialize each laser scanner's vertical offset, h_i (see (14)), and the *yaw* component of the rotation to initialize each laser scanner's θ_{oi} (see (13)).

We then formulate the following constrained minimization problem to enforce (13) and (14):

$$\min \sum_{i,j,k} \frac{\left[\alpha_{i}(\rho_{ijk} + \rho_{oi})^{C} \bar{\mathbf{n}}_{j\ L_{i}}^{T\ C} \mathbf{C}^{L_{i}} \bar{\mathbf{p}}_{ijk} + {}^{C} \bar{\mathbf{n}}_{j}^{T\ C} \mathbf{t}_{L_{i}} - d_{j}\right]^{2}}{\sigma_{\epsilon_{ijk}}^{2}}$$

subject to: ${}_{L_{i}}^{C} \mathbf{C} = {}_{L}^{C} \mathbf{C} \mathbf{C}_{z}(\theta_{oi})$
 ${}_{L}^{C} \mathbf{C}^{T} ({}^{C} \mathbf{t}_{L_{i}} - {}^{C} \mathbf{t}_{L}) = [0\ 0\ h_{i}]^{T}$ (32)

where the optimization variables are α_i , ρ_{oi} , θ_{oi} , h_i , i = 2, ..., K, α_1 , ρ_{o1} , ${}^c\mathbf{t}_L$, ${}^c_L\mathbf{C}$, and $\sigma^2_{\epsilon_{ijk}}$ is defined in (9).⁴ Note that the constraints in (32) should be taken into account using the method of Lagrange multipliers. Alternatively, we minimized a reformulation of (32) that uses a minimal parameterization of the unknowns to avoid the constraints (and, hence, the Lagrange multipliers). The details of this alternative cost function are provided in Appendix D.

4. Identifiability conditions

In this section, we examine the conditions under which the unknown LIDAR–camera transformation and the intrinsic parameters of the LIDAR are identifiable, and thus can be estimated using the algorithms in Sections 3.1–3.4.

4.1. Observation of one plane

Suppose that we are provided with LIDAR measurements that lie only on one plane whose normal vector is denoted as ${}^{C}\bar{\mathbf{n}}_{1}$. In this case, it is easy to show that the measurement constraint in (6) does not change if ${}^{C}_{L_{i}}\mathbf{C}$ is perturbed by a rotation around ${}^{C}\bar{\mathbf{n}}_{1}$, represented by the rotation matrix \mathbf{C}' :

$${}^{C}\bar{\mathbf{n}}_{1}^{\mathrm{T}}\mathbf{C}'{}_{L_{i}}^{C}\mathbf{C}{}^{L_{i}}\mathbf{p}_{i1k} + {}^{C}\bar{\mathbf{n}}_{1}^{\mathrm{T}C}\mathbf{t}_{L_{i}} - d_{1} = 0 \qquad (33)$$

$$\implies {}^{C}\bar{\mathbf{n}}_{1\ L_{i}}^{\mathrm{T}\ C}\mathbf{C}^{L_{i}}\mathbf{p}_{i1k} + {}^{C}\bar{\mathbf{n}}_{1}^{\mathrm{T}\ C}\mathbf{t}_{L_{i}} - d_{1} = 0.$$
(34)

The second equation is obtained from the first, since ${}^{C}\bar{\mathbf{n}}_{1}$ is an eigenvector of \mathbf{C}' , thus ${}^{C}\bar{\mathbf{n}}_{1}^{T}\mathbf{C}' = {}^{C}\bar{\mathbf{n}}_{1}^{T}$. Therefore, when observing only one plane, any rotation around the plane's normal vector is unidentifiable. Similarly, if we perturb ${}^{C}\mathbf{t}_{L_{i}}$ by a translation parallel to the plane, represented by \mathbf{t}' , the measurement constraint does not change:

$${}^{C}\bar{\mathbf{n}}_{1\ L_{i}}^{\mathrm{T}\ C}\mathbf{C}^{L_{i}}\mathbf{p}_{i1k} + {}^{C}\bar{\mathbf{n}}_{1}^{\mathrm{T}\ (C}\mathbf{t}_{L_{i}} + \mathbf{t}') - d_{1} = 0 \qquad (35)$$

$$\implies {}^{C}\bar{\mathbf{n}}_{1\ L_{i}}^{T\ C}\mathbf{C}^{L_{i}}\mathbf{p}_{i1k} + {}^{C}\bar{\mathbf{n}}_{1}^{T\ C}\mathbf{t}_{L_{i}} - d_{1} = 0.$$
(36)

This relationship holds since ${}^{c}\bar{\mathbf{n}}_{1}^{T}\mathbf{t}' = 0$. Therefore, when observing only one plane, any translation parallel to the plane's normal is unidentifiable.

4.2. Observation of two planes

Consider now that we are provided with measurements from two planes, described by ${}^{C}\bar{\mathbf{n}}_{1}$, d_{1} , ${}^{C}\bar{\mathbf{n}}_{2}$, d_{2} . If we perturb the laser scanner's relative translation with $\mathbf{t}'' \propto {}^{C}\bar{\mathbf{n}}_{1} \times {}^{C}\bar{\mathbf{n}}_{2}$ (see (35)), none of the measurement constraints will change, since ${}^{C}\bar{\mathbf{n}}_{1}^{T}\mathbf{t}'' = {}^{C}\bar{\mathbf{n}}_{2}^{T}\mathbf{t}'' = 0$. Therefore, we conclude that the relative translation cannot be determined if only two planes are observed.

4.3. Observation of three planes

In this section, we prove that when three planes with linearly independent normal vectors are observed, we can determine all of the unknowns. For this purpose, we first determine the relative orientation ${}_{L_i}^C \mathbf{C}$ and the offset ρ_{oi} and then find the scale α_i and relative translation ${}^C \mathbf{t}_{L_i}$. Let us assume that the *i*th laser scanner has measured four points on each plane, denoted as $(\rho_{ijk}, {}^{L_i} \mathbf{\bar{p}}_{ijk}), j = 1, 2, 3, k = 1, ..., 4$. Each of these points provides one constraint of the form (7). We first eliminate the unknown relative translation and scale, by subtracting the constraints for point k = 1 from k = 2, point k = 2 from k = 3, and point k = 3 from k = 4, and obtain

$${}^{C}\bar{\mathbf{n}}_{j\ L_{i}}^{\mathrm{T}\ C}\mathbf{C}\left(\mathbf{u}_{j12}^{i}+\rho_{oi}\,\mathbf{v}_{j12}^{i}\right)=0\tag{37}$$

$${}^{C}\bar{\mathbf{n}}_{j\ L_{i}}^{\mathrm{T}\ C}\mathbf{C}\left(\mathbf{u}_{j23}^{i}+\rho_{oi}\,\mathbf{v}_{j23}^{i}\right)=0\tag{38}$$

$${}^{C}\bar{\mathbf{n}}_{j\ L_{i}}^{\mathrm{T}\ C}\mathbf{C}\left(\mathbf{u}_{j34}^{i}+\rho_{oi}\,\mathbf{v}_{j34}^{i}\right)=0\tag{39}$$

where $\mathbf{u}_{jkl}^i \triangleq \rho_{ijk}^{L_i} \bar{\mathbf{p}}_{ijk} - \rho_{ijl}^{L_i} \bar{\mathbf{p}}_{ijl}$, $\mathbf{v}_{jkl}^i \triangleq {}^{L_i} \bar{\mathbf{p}}_{ijk} - {}^{L_i} \bar{\mathbf{p}}_{ijl}$, and j = 1, 2, 3. Note that ${}^{L_i} \bar{\mathbf{p}}_{ijk}$ and ${}^{L_i} \bar{\mathbf{p}}_{ijl}$ lie on the intersection of the unit sphere and the cone specified by the beams of the *i*th laser scanner. Since the intersection of a co-centric unit sphere and a cone is always a circle, we conclude that all \mathbf{v}_{jkl}^i for a given *i* belong to a plane and have only 2 d.o.f. Thus, we can write \mathbf{v}_{j34}^i as a linear combination of \mathbf{v}_{j12}^i and \mathbf{v}_{j23}^i , i.e.

$$\mathbf{v}_{j34}^{i} = a \, \mathbf{v}_{j12}^{i} + b \, \mathbf{v}_{j23}^{i} \tag{40}$$

for some known scalars *a* and *b*. Substituting this relationship in (39), and using (37)-(38) to eliminate the terms containing ρ_{oi} yields

$${}^{C}\bar{\mathbf{n}}_{j\ L_{i}}^{\mathrm{T}\ C}\mathbf{C}\left(\mathbf{u}_{j34}^{i}-a\,\mathbf{u}_{j12}^{i}-b\,\mathbf{u}_{j23}^{i}\right)=0$$
(41)

for j = 1, 2, 3. The only unknown in this equation is the relative orientation ${}_{L_i}^{C}\mathbf{C}$ of the *i*th laser scanner. These equations are identical to those for orientation estimation using line-to-plane correspondences, which is known to have at most eight solutions that can be analytically computed when ${}^{C}\bar{\mathbf{n}}_{j}, j = 1, 2, 3$, are linearly independent (Chen 1991). Once ${}_{L_i}^{C}\mathbf{C}$ is known, we can use any of (37)–(39) to compute the offset ρ_{oi} . Finally, the scale and the relative translation can be obtained from (31).

5. Experiments

5.1. Setup

In order to validate the proposed calibration method, we conducted a series of experiments using a Velodyne revolving-head 3D LIDAR and a Ladybug2 spherical vision system. The Velodyne consists of 64 laser scanners that collectively span 27° of the vertical field of view. The Ladybug consists of six rigidly connected and intrinsically calibrated cameras equipped with wide-angle lenses (see Figure 2). The extrinsic calibration between the different cameras is also provided by the manufacturer with high accuracy. Therefore, the measurements from any of the cameras can be easily transformed to the Ladybug's fixed frame of reference. We rigidly connected the Velodyne and the Ladybug, and recorded measurements of a $36" \times 40"$ calibration plane at 18 different configurations (see Figure 3). By processing the Ladybug's images using the PnP algorithm of Hesch and Roumeliotis (2011), we





Fig. 3. (a) A view of the calibration environment. Note the Velodyne–Ladybug pair at the bottom-right of the picture. The configuration (i.e. position and orientation) of the calibration board (center-right of the picture) changed for each data capture. (b) A typical LIDAR snapshot, constructed using the LIDAR's intrinsic parameters provided by the manufacturer. The extracted calibration plane is shown in green. The red dots specify the extracted points corresponding to the laser scanner 20.

computed the normal vector and the distance of the calibration plane at each configuration. We then identified the approximate location of the calibration plane in the LIDAR scans based on a coarse prior estimate for the relative rotation of the Velodyne and the Ladybug. Within these approximate locations, we detected the LIDAR data points reflected from the calibration plane, based on their depth discontinuity.

Once the Velodyne's measurements for each configuration of the calibration plane were available, we used the method described in Section 2 to accurately estimate the LIDAR's intrinsic parameters and the LIDAR–camera transformation. Note, however, that in order to increase the robustness of our algorithm to outliers, we did not directly use the raw laser points measured by the LIDAR. Instead, for each laser scanner, we fit small line segments to the intersection of the laser scanner's beam and the calibration plane, and used the endpoints of these line segments as the LIDAR's measurements.⁵

5.2. Implemented methods

We compared the accuracy and consistency of the calibration parameters estimated by our proposed algorithm (denoted as **AlgBLS**), with the results of the approach presented by Pandey et al. (2010) (denoted as **PMSE**), and with those when using the calibration parameters provided by the manufacturer (intrinsic parameters only, denoted as **Factory**). Note that the **PMSE** only calibrates the offset in the range measurements of each laser scanner (i.e. ρ_{oi} ; see (2)), while for the rest of the parameters it uses the **Factory** values.

We implemented the PMSE as follows: for each calibration-plane configuration, we transformed the laser points reflected from the plane surface to the LIDAR's Euclidean frame (see (1), (2), (13), and (14)) based on the Factory parameters, and fitted a plane to them using RANSAC (Hartley and Zisserman 2004). In the next step, we employed least-squares to minimize the distance of the laser points from the fitted planes by optimizing over the range offsets, ρ_{oi} . The Euclidean coordinates of the laser points are then adjusted accordingly, and processed to fit new planes using RANSAC; these re-fitted planes are used, in turn, to re-estimate the range offsets. This process is continued until convergence, or until a maximum number of iterations is reached. Once the range offsets were calibrated, we minimized a least-square cost function similar to (32), but only over the extrinsic calibration parameters (i.e. ${}^{C}_{L}\mathbf{C}$ and ${}^{C}\mathbf{t}_{L}$).

5.3. Consistency of intrinsic parameters

In order to evaluate the consistency of the estimated intrinsic parameters, we collected a new test dataset comprising the image observations and LIDAR snapshots of the calibration plane at 17 different configurations. We then extracted the LIDAR points belonging to the calibration plane following the procedure described in Section 6.1, and computed their Euclidean coordinates using the intrinsic parameters of the Factory, the PMSE, and the AlgBLS. A sample reconstruction from each of these methods is shown in Figure 4, where the inaccuracy in some of the intrinsic LIDAR parameters of the Factory is clearly visible. To quantitatively compare the estimated intrinsic parameters, we fitted planes to the LIDAR points and computed the (signed) residual distance⁶ of each point to the corresponding plane, from which the histograms in Figure 5 are obtained. The detailed statistics of the signed distances for all three methods are presented in Table 1. It is evident from these statistics that the **AlgBLS** results in a significantly smaller median value, indicating lower skewness and bias in the errors of the **AlgBLS** than those of the **Factory** and **PMSE**. This, in effect, shows that the **AlgBLS** leads to more consistent calibration across different laser scanners of the LIDAR.

5.4. Comparison of intrinsic and extrinsic parameters

The evaluations in Section 5.3 are incomplete since they do not reflect the accuracy of the extrinsic camera-LIDAR parameters. To complement the above results, we transformed the laser points in the test datasets to the Ladybug's frame of reference using the extrinsic calibration parameters estimated by the AlgBLS and the PMSE. Then, we computed the signed residual distance of the points reflected by the calibration planes, from the planes as detected by the Ladybug by evaluating (5) for each point. The histograms of these errors (i.e. signed distances) are shown in Figure 6 and their statistics are provided in Table 2. As evident from these results, the AlgBLS yields superior accuracy compared with the PMSE. In particular, higher median and mean error indicates a bias in the estimates of the PMSE. This may be explained by the fact that the PMSE does not calibrate the scale of the LIDAR measurements.

In order to ensure the repeatability of the proposed calibration procedure, we collected three independent datasets, and employed the **AlgBLS** to determine the calibration parameters from each one. A sample of these estimates is provided in Table 3. The parameters' estimates are relatively stable, except for the *z* component of ${}^{C}\mathbf{t}_{L}$ which shows a few centimeters of variation. This can be explained by the fact that due to the limited and mostly horizontal field of view of the LIDAR, the geometric constraints in (7) provide less information along the vertical axis of the LIDAR.

Finally, Table 3 provides a selection of the calibration parameters obtained by (i) iteratively minimizing the least-squares cost function of (32) with an inaccurate initialization (labeled as Bad Init.), (ii) the PMSE algorithm, and (iii) the Factory intrinsic parameters. Note that with inaccurate initialization, the estimate for the z component of ${}^{C}\mathbf{t}_{L}$ has converged to a local minimum, far from the actual value of the translation between the Velodyne and the Ladybug. The PMSE method performs optimization only over the range offset, and uses the **Factory** for the rest of the parameters (labeled with N.A. in Table 3). In this table, note the inaccurate range offset estimate of the laser scanner 20 as provided by the Factory. The point cloud corresponding to a plane obtained using these estimates is shown in Figures 4(a) and 4(b). In particular, the laser points from scanner 20 (marked in red) are easily distinguishable as they suffer from \sim 30 cm bias. These errors become smaller when using the calibration parameters of PMSE method (see Figures 4(c) and 4(d)), but they are still noticeable. In



Fig. 4. Consistency of the intrinsic parameters. (a), (b) LIDAR points reflected from the calibration plane in a test dataset viewed from the front and side, respectively. The points' Euclidean coordinates were computed using the **Factory** parameters. Note the considerable bias of the points from the laser scanner 20 shown in red (grid size 10 cm). (c), (d) The same LIDAR points, front and side view, respectively, when their Euclidean coordinates are computed using the **PMSE** intrinsic parameters. (e), (f) The same LIDAR points, front and side view, respectively, when their Euclidean coordinates are computed using the **AlgBLS** intrinsic parameters. Note that the points from the laser scanner 20 (shown in red) no longer exhibit a significant bias (grid size 10 cm).

Table 1. Statistics of the signed distance error of the laser points from the fitted planes.

	Mean (mm)	Median (mm)	Standard deviation (mm)			
Factory	-0.85	0.11	65			
PMSE	2.56	-0.61	152			
AlgBLS	-0.67	-0.05	55			



Fig. 5. Histograms of the signed distance between laser points reflected from the calibration target and the corresponding fitted plane. The laser points' Euclidean coordinates in each of the above plots are computed using the intrinsic LIDAR parameters determined by three different methods.

contrast, the reconstruction obtained using the calibration parameters estimated by the **AlgBLS** (Figures 4(e) and 4(f)) exhibit considerably lower systematic error.

5.5. Photorealistic reconstructions

To further demonstrate the accuracy of the estimated calibration parameters by the **AlgBLS**, we employed the



Fig. 6. Histograms of the signed distance errors of the laser points from the calibration planes detected by the Ladybug.

following procedure to create photorealistic reconstructions of several indoor and outdoor scenes from the University of Minnesota campus (see Figures 7–8):

- Transform the LIDAR points to the Ladybug's frame of reference using the **AlgBLS** intrinsic and extrinsic calibration parameters.
- Overlay the spherical image of the Ladybug on the LIDAR points, and assign each pixel a depth according to the corresponding LIDAR point. If a pixel has no

	Mean (mm)	Median (mm)	Standard deviation (mm)			
PMSE	40	55	151			
AlgBLS	-4.3	-1.4	28			

Table 2. Statistics of the signed distance error of the laser points from the calibration planes detected by the Ladybug.

Table 3. A sample of the estimated calibration parameters. The first three rows show a selection of calibration parameters as estimated by **AlgBLS** for three different datasets. The column labeled as *rpy* represents roll, pitch, and yaw of $_{L}^{C}$ C. Rows 4–6 show a selection of calibration parameters that are obtained by other methods. Note that the scale parameters α_i are unitless, and N.A. indicates non-applicable fields.

	Number of planes	rpy (deg)	$^{C}\mathbf{t}_{L}$ (cm)	α1	$\rho_{o1} (\mathrm{cm})$	α_{20}	$\rho_{o20} \ ({\rm cm})$	θ_{o20} (deg)	<i>h</i> ₂₀ (cm)
AlgBLS 1	18	[1.31 1.27 -88.45]	[-2.5 0.04 -20.10]	1.02	99.49	1.01	100.88	36.91	-3.71
AlgBLS 2	17	[1.62 0.79 -88.40]	[-1.64 0.19 -12.92]	1.01	101.83	1.00	101.82	36.03	-4.98
AlgBLS 3	18	[1.26 0.84 -88.43]	$[-2.07 \ -0.12 \ -17.95]$	1.01	101.05	0.98	105.94	36.05	-0.85
Bad Init.	18	[0.34 - 0.63 - 87.93]	[-1.64 - 0.50 - 117.08]	1.01	98.99	0.99	104.49	34.95	98.31
PMSE	18	[3.19 -0.97 -88.32]	[1.40 -2.12 -53.64]	N.A.	91.49	N.A.	90.85	N.A.	N.A.
Factory	18	N.A.	N.A.	1	95.50	1	66.53	34.95	-0.24

corresponding LIDAR point, compute an approximate depth by linearly interpolating the nearest laser points.

• Render 3D surfaces from the 3D pixels using Delaunay triangulation (de Berg et al. 2008).

In Figures 7(b)–7(c) and 8(b)–8(c), a selection of the rendered surfaces in Matlab using the estimated calibration parameters by **AlgBLS** are shown for indoor and outdoor scenes. Note that white gaps in the reconstructed surfaces result from relatively large patches of missing LIDAR measurements due to occlusion or specular reflection of the laser beams from glass and shiny surfaces.

The reconstructions using the calibration parameters of **PMSE** and **Factory** look quite similar to Figures 7(b)–7(c) and 8(b)-8(c) when zoomed out, and hence we omit the corresponding figures. However, Figures 7(d)-7(g) and 8(d)-8(g) show close-up views of the same surfaces with renderings obtained by using AlgBLS and alternative methods. In particular, in Figures 7(e) and 8(e), the calibration parameters estimated by minimizing (32) given an imprecise initialization are employed to render the image. In creating Figures 7(f) and 8(f), we have only estimated the extrinsic calibration parameters, and used the Factory intrinsic parameters. Finally Figures 7(g) and 8(g) are rendered using the estimated calibration parameters by PMSE. The superior quality of the renderings obtained using the calibration parameters of the AlgBLS is clear in these images. In particular, in Figures 8(d)-8(g), the lines corresponding to the corner of the wall are estimated based on the LIDAR measurements, and depicted as green dashed line. Given a perfect calibration, the shadow edge of the overlaid image should match the extracted corner. Clearly, while the rendering by AlgBLS provides an almost perfect match between the two, the other rendered images

suffer from mismatches caused by inaccurate estimates of the calibration parameters. Furthermore, the intrinsic calibration parameters provided by the manufacturer often lead to spikes in the reconstructed surface, as shown by the green arrows in 7(f) and 8(f).

6. Conclusions and future work

In this paper, we presented a novel method for intrinsic calibration of a revolving-head 3D LIDAR and extrinsic calibration with respect to a camera. Specifically, we developed an analytical method for computing a precise initial estimate for both the LIDAR's intrinsic parameters and the LIDAR-camera transformation. Subsequently, we used these estimates to initialize an iterative nonlinear leastsquares refinement of all of the calibration parameters. In addition, we presented an identifiability analysis to determine the minimal conditions under which it is possible to estimate the calibration parameters. Experimental results from both indoor and outdoor scenes are used to demonstrate the achieved accuracy of the calibration process by photorealistic reconstruction of the observed areas. Optimally combining multiple images and LIDAR scans over consecutive time steps for mapping large areas while at the same time increasing the 3D points' resolution and revealing occluded areas, is part of our ongoing research work.

Notes

 When the technical drawings of the LIDAR are available, a coarse initial estimate for h_i, θ_{oi}, and φ_i can be readily obtained. Computing an initial estimate for ρ_{oi} and α_i,







Fig. 7. Photorealistic reconstruction of an indoor scene (best viewed in color). (a) Panoramic image of an indoor scene with several corridors provided by the Ladybug. (b), (c) The corresponding photorealistic reconstruction using the calibration parameters obtained from the **AlgBLS** algorithm, viewed from two different directions. The green rectangle in (b) marks the close-up area shown in the last row. (d) The close-up view rendered using the parameters estimated by the **AlgBLS** algorithm. (e) The close-up view rendered using the parameters estimated by the **AlgBLS** algorithm. (e) The close-up view rendered using the parameters estimated by the **AlgBLS** algorithm. (e) The close-up view rendered using the intrinsic parameters provided by the manufacturer. (g) The close-up view rendered using the algorithm proposed by Pandey et al. (2010). The white gaps are the regions where at least one of the sensors did not return meaningful measurements (e.g. due to occlusion, specular reflections, or limited resolution and field of view). Note that the depth of the scene can be inferred from the dotted grids.

however, is significantly more challenging even for the manufacturer, since their values do not solely depend on the physical dimensions of the device.

- 2. For example, see Alvar fiducial markers at http://www. vtt.fi/multimedia/alvar.html
- 3. These conditions are: (i) the ideal must be radical; (ii) its variety must be non-empty and zero dimensional (Cox et al. 2004). These conditions are generally satisfied for the current problem.
- 4. In general, the optimization should be performed over ϕ_i as well. However, in our experiments, we observed that the provided value of ϕ_i by the manufacturer is sufficiently accurate, and thus excluded it from the calibration.
- 5. In general, the intersection of the cone induced by the laser scanner's beam with a plane results in a conic section, and not a straight line. However, in practical situations this conic section can be well approximated with consecutive straight line segments.

- 6. We assigned a positive (negative) sign to the distance of the points in front (behind) the plane.
- 7. In other words, except for singular choices of the coefficients (e.g. zero coefficients), the normal set remains the same. For a more precise definition of *genericity*, refer to Cox et al. (2004).

Funding

This work was supported by the University of Minnesota through the Digital Technology Center (DTC), the National Science Foundation (grant number IIS-0811946), and Air Force Office of Scientific Research (grant number FA9550-10-1-0567).

Acknowledgement

A short version of this work was presented at the 15th International Symposium on Robotics Research (ISRR), Flagstaff,



Fig. 8. Photorealistic reconstruction of an outdoor scene (best viewed in color). (a) Panoramic view of an outdoor scene. (b), (c) Photorealistic reconstruction of the scene viewed from two different directions. The green rectangle in (b) marks the close-up area shown in the last row. (d) The close-up view obtained using the parameters estimated by the **AlgBLS** algorithm. (e) The close-up view obtained using the parameters estimated by an iterative least-squares refinement with inaccurate initialization. (f) The close-up view obtained using the intrinsic parameters provided by the manufacturer. The green arrow points to the spikes created due to inaccurate range offset parameters. (g) The close-up view obtained using the algorithm proposed by Pandey et al. (2010). The green dashed lines mark the corner of the wall detected from the LIDAR points. The white gaps are the regions where at least one of the sensors did not return meaningful measurements (e.g. due to occlusion, specular reflections, or limited resolution and field of view). Note that some of the occlusions are due to the trees, the lamp post, and different elevations of the grassy area.

AZ, 28 August–1 September 2011. The authors would like to thank Sean Bowman for writing the C++ interface for simultaneous collection of Velodyne and Ladybug measurements, Joel Hesch for providing the code for constructing and painting Delaunay triangles, and Igor Melnyk for helping to conduct the experiments.

References

- Ansar A and Daniilidis K (2003) Linear pose estimation from points or lines. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25: 578–589.
- Auzinger W and Stetter HJ (1988) An elimination algorithm for the computation of all zeros of a system of multivariate

polynomial equations. In *Proceedings of the International Conference on Numerical Mathematics*, Singapore, pp. 11–30.

- Byröd M, Josephson K and Åström K (2008) A column-pivoting based strategy for monomial ordering in numerical Gröbner basis calculations. In *Proceedings of the European Conference* on Computer Vision, Marseille, France, pp. 130–143.
- Chen HH (1991) Pose determination from line-to-plane correspondences: existence condition and closed-form solutions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13: 530–541.
- Cox D, Little J and O'Shea D (2004) Using Algebraic Geometry. Berlin: Springer.
- de Berg M, Cheong O, van Kreveld M and Overmars M (2008) Computational Geometry: Algorithms and Applications. Berlin: Springer-Verlag.
- Fischler MA and Bolles RC (1981) Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* 24: 381–395.
- Golub GH and van Loan CF (1996) *Matrix Computations*. Baltimore, MD: Johns Hopkins University Press.
- Hartley RI and Zisserman A (2004) Multiple View Geometry in Computer Vision, 2nd edn. Cambridge: Cambridge University Press.
- Hesch JA and Roumeliotis SI (2011) A direct least-squares (DLS) solution for PnP. In *Proceedings of the International Conference on Computer Vision*, Barcelona, Spain.
- Naroditsky O, Patterson A, IV, and Daniilidis K (2011) Automatic alignment of a camera with a line scan lidar system. In *Proceedings IEEE International Conference on Robotics and Automation*, Shanghai, China.
- Pandey G, McBride J, Savarese S and Eustice R (2010) Extrinsic calibration of a 3D laser scanner and an omnidirectional camera. In *Proceedings IFAC Symposium on Intelligent Autonomous Vehicles*, Lecce, Italy.
- Press WH, Teukolsky SA, Vetterling WT and Flannery BP (1992) *Numerical Recipes in C.* Cambridge: Cambridge University Press.
- Quan L and Lan Z-D (1999) Linear *n*-point camera pose determination. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21: 774–780.
- Reid G and Zhi L (2009) Solving polynomial systems via symbolic–numeric reduction to geometric involutive form. *Journal of Symbolic Computation* 44: 280–291.
- Scaramuzza D, Harati A and Siegwart R (2007) Extrinsic self calibration of a camera and a 3D laser range finder from natural scenes. In *Proceedings IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Diego, CA, pp. 4164–4169.
- Shuster MD (1993) A survey of attitude representations. *Journal* of Astronautical Science 41: 439–517.
- Stamos I, Liu L, Chen C, Wolberg G, Yu G and Zokai S (2008) Integrating automated range registration with multiview geometry for the photorealistic modeling of large-scale scenes. *International Journal of Computer Vision* 78: 237–260.
- Trawny N and Roumeliotis SI (2010) On the global optimum of planar, range-based robot-to-robot relative pose estimation. In *Proceedings IEEE International Conference on Robotics and Automation*, Anchorage, AK, pp. 3200–3206.
- Trawny N, Zhou XS and Roumeliotis SI (2009) 3D relative pose estimation from six distances. In *Proceedings of Robotics Science and Systems*, Seattle, WA.

- Unnikrishnan R and Hebert M (2005) *Fast Extrinsic Calibration* of a Laser Rangefinder to a Camera. Technical report, Carnegie Mellon University, Robotics Institute.
- Zhang Q and Pless R (2004) Extrinsic calibration of a camera and laser range finder (improves camera calibration). In Proceedings IEEE/RSJ International Conference Intelligent Robots and Systems, Sendai, Japan, pp. 2301–2306.

Appendix A

Notation

- $\{B_j\}$ coordinate frame of reference corresponding to the calibration board at the *j*th configuration
- $\{C\}$ camera's coordinate frame of reference
- ${}_{Y}^{X}\mathbf{C}$ rotation matrix representing the relative orientation of $\{Y\}$ w.r.t. $\{X\}$
- d_j distance of the calibration plane, at its *j*th configuration from the origin of $\{C\}$
- h_i vertical offset of the *i*th laser scanner w.r.t. $\{L\}$
- \mathbf{I}_n $n \times n$ identity matrix
- $\{L\}$ LIDAR's coordinate frame of reference
- $\{L_i\}$ coordinate frame of reference corresponding to the *i*th laser scanner, i = 1, ..., K
- ${}^{C}\bar{\mathbf{n}}_{j}$ normal vector of the calibration plane, at its *j*th configuration, w.r.t. {*C*}
- $L_i \mathbf{p}_{ijk}$ kth intrinsically corrected point, belonging to the calibration board at its *j*th configuration, measured by the *i*th laser scanner, w.r.t. { L_i }
- ${}^{X}\mathbf{t}_{Y}$ relative position of $\{Y\}$ w.r.t. $\{X\}$
- w.r.t. with respect to
- $\mathbf{0}_{m \times n}$ $m \times n$ matrix of zeros
- α_i scale factor of the *i*th laser scanner
- ϕ_i elevation angle of the *i*th laser scanner
- θ_{oi} azimuth angle between coordinate frames $\{L\}$ and $\{L_i\}$
- θ_{ik} azimuth angle of the *k*th shot of the *i*th laser scanner w.r.t. $\{L_i\}$
- ρ_{ik} range measurement of the *k*th shot of the *i*th laser scanner w.r.t. $\{L_i\}$
- ρ_{oi} range offset of the *i*th laser scanner

Appendix B

The optimality conditions of (22) are as follows:

$$f_{i\ell} = \frac{\partial \mathcal{C}_i}{\partial s_{i\ell}} = \sum_{j=1}^{M} \sum_{k=1}^{N_{ij}} \sum_{l=\frac{N_{ij}}{2}+1}^{N_{ij}} \left({}^{C} \bar{\mathbf{n}}_j^{\mathrm{T}} \bar{\mathbf{C}}(\mathbf{s}_i) \left(\mathbf{u}_{jkl}^i + \rho_{oi} \mathbf{v}_{jkl}^i \right) \right) \\ \cdot \underbrace{\frac{\partial}{\partial s_{i\ell}} \left({}^{C} \bar{\mathbf{n}}_j^{\mathrm{T}} \bar{\mathbf{C}}(\mathbf{s}_i) \left(\mathbf{u}_{jkl}^i + \rho_{oi} \mathbf{v}_{jkl}^i \right) \right)}_{J_{i\ell}} = 0.$$
(42)

For $\ell = 1, 2, 3, J_{i\ell}$ is

$$J_{i\ell} = {}^{\scriptscriptstyle C} \bar{\mathbf{n}}_{j}^{\rm T} \, \mathbf{D}_{\ell}(\mathbf{s}_{i}) \left(\mathbf{u}_{jkl}^{i} + \rho_{oi} \mathbf{v}_{jkl}^{i} \right) \tag{43}$$

where

$$\mathbf{D}_{\ell}(\mathbf{s}_{i}) = -2s_{i\ell}\mathbf{I}_{3} + 2\lfloor\mathbf{e}_{\ell}\times\rfloor + 2\mathbf{e}_{\ell}\mathbf{s}_{i}^{\mathrm{T}} + 2\mathbf{s}_{i}\mathbf{e}_{\ell}^{\mathrm{T}} \qquad (44)$$

$$\mathbf{e}_1 \triangleq \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^{\mathrm{T}}, \ \mathbf{e}_2 \triangleq \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}^{\mathrm{T}}, \ \mathbf{e}_3 \triangleq \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^{\mathrm{T}}$$
 (45)

and for $\ell = 0$, i.e. $s_{i0} \triangleq \rho_{oi}, J_{i0} = {}^{C} \bar{\mathbf{n}}_{j}^{T} \bar{\mathbf{C}}(\mathbf{s}_{i}) \mathbf{v}_{jkl}^{i}$.

Appendix C

Consider the following simple example polynomials in $\mathbf{x} = [x_1 \ x_2]^{\mathrm{T}}$:

$$f_1 = x_1 + x_1 x_2 + 5 \tag{46}$$

$$f_2 = x_1^2 + x_2^2 - 10. (47)$$

These equations are of degree $d_1 = d_2 = 2$. The Gröbner basis of this polynomial system (using graded reverse lex ordering (Cox et al. 2004)) is

$$g_1 = x_1 x_2 + x_1 + 5 \tag{48}$$

$$g_2 = x_1^2 + x_2^2 - 10 \tag{49}$$

$$g_3 = x_2^3 + x_2^2 - 5x_1 - 10x_2 - 10$$
 (50)

and, consequently, its normal set is

$$\{1, x_2, x_1, x_2^2\}.$$
 (51)

Note that this normal set is *generically*⁷ the same for different coefficients of the system in (46)–(47). For example the following system yields the same normal set:

$$f_1' = 1.5x_1 + e^{-1}x_2x_1 + 0.5 \tag{52}$$

$$f_2' = 2.3x_1^2 + \frac{4}{3}x_2^2 - \pi.$$
 (53)

Let us arrange the normal set in the vector form $\mathbf{x}_B = [1, x_2, x_1, x_2^2]^T$ and choose $\varphi(\mathbf{x}) = x_2$. Then multiplying

 $\varphi(\mathbf{x})$ with \mathbf{x}_B and expressing the result in terms of \mathbf{x}_B and \mathbf{x}_R (see (27)) yields:

$$\varphi(\mathbf{x})\mathbf{x}_{B} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \underbrace{\begin{bmatrix} 1 \\ x_{2} \\ x_{1} \\ x_{2}^{2} \end{bmatrix}}_{\mathbf{x}_{B}} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \underbrace{\begin{bmatrix} x_{1}x_{2} \\ x_{2}^{3} \\ \mathbf{x}_{R}}_{\mathbf{x}_{R}}.$$
(54)

In order to express \mathbf{x}_R in terms of \mathbf{x}_B , we expand the polynomials f_1 and f_2 up to degree $\ell = 3$ by multiplying each of them with $\{1, x_2, x_1\}$. As a result, we obtain $\mathbf{C}_e = [\mathbf{C}_E \ \mathbf{C}_R \ \mathbf{C}_B]$ where

$$\mathbf{C}_{E} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{C}_{R} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}$$
$$\mathbf{C}_{B} = \begin{bmatrix} 5 & 0 & 1 & 0 \\ 0 & 5 & 0 & 0 \\ 0 & 0 & 5 & 0 \\ -10 & 0 & 0 & 1 \\ 0 & -10 & 0 & 0 \\ 0 & 0 & -10 & 0 \end{bmatrix}.$$

Note that C_E corresponds to $\mathbf{x}_E = [x_1x_2^2 \ x_1^2x_2 \ x_1^2 \ x_1^3]^T$, i.e. the monomials that appear neither in \mathbf{x}_B nor in \mathbf{x}_R . Following the algebraic manipulations of (27)–(29), we obtain the following multiplication matrix:

$$\mathbf{M}_{x_2} = \begin{bmatrix} 0 & 0 & -5 & 10 \\ 1 & 0 & 0 & 10 \\ 0 & 0 & -1 & 5 \\ 0 & 1 & 0 & -1 \end{bmatrix}.$$
 (55)

In the next step, we compute the left eigenvectors of \mathbf{M}_{x_2} , and then scale them such that their first elements become 1 (corresponding to the first element in \mathbf{x}_B). Consequently, the solutions of the polynomial system are the elements of the eigenvectors that correspond to x_1 and x_2 in \mathbf{x}_B . Specifically, the set of solutions (i.e. variety) of (46)–(47) is

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \in \left\{ \begin{bmatrix} -1.2856 \\ 2.8891 \end{bmatrix}, \begin{bmatrix} -3.1026 \\ 0.6116 \end{bmatrix}, \\ \begin{bmatrix} 2.1941 + 1.2056i \\ -2.7504 + 0.9618i \end{bmatrix}, \begin{bmatrix} 2.1941 - 1.2056i \\ -2.7504 - 0.9618i \end{bmatrix} \right\}.$$
(56)

Appendix D

One approach for enforcing the constraints in (32) is to use the method of Lagrange multipliers. It is possible, however, to re-parameterize the cost function and minimally express it over the optimization variables. In this way, the constraints in (32) will be automatically satisfied. Specifically, we consider the *k*th *intrinsically corrected* point measured by the *i*th laser scanner from the *j*th configuration of the calibration plane as:

$${}^{L}\mathbf{p}_{ijk} = \begin{bmatrix} \alpha_i(\rho_{ijk} + \rho_{oi})\cos\phi_i\cos(\theta_{ijk} + \theta_{oi})\\ \alpha_i(\rho_{ijk} + \rho_{oi})\cos\phi_i\sin(\theta_{jik} + \theta_{oi})\\ \alpha_i(\rho_{ijk} + \rho_{oi})\sin\phi_i + h_i \end{bmatrix}.$$
 (57)

This relationship is obtained by substituting (1) in (2), and then transforming the result to the LIDAR's frame of reference {L}. Note that the intrinsic LIDAR parameters are already expressed in their minimal form; thus the constraints in (32) are redundant and can be removed. In particular, (14) is satisfied since h_i is added to the z component of the point's measurement and (13) is satisfied since θ_{oi} is added to the azimuth of the point's measurement. Also, we set $h_1 = \theta_{o1} = 0$, since we have assumed $\{L_1\} \equiv \{L\}$. Based on (57), we define the following *unconstrained* minimization problem:

min
$$\sum_{i,j,k} \frac{\left({}^{C}\bar{\mathbf{n}}_{j}^{\mathrm{T}}{}_{L}^{C}\mathbf{C}^{L}\mathbf{p}_{ijk} + {}^{C}\bar{\mathbf{n}}_{j}^{\mathrm{T}}{}^{C}\mathbf{t}_{L} - d_{j}\right)^{2}}{\sigma_{\epsilon_{ijk}}^{2}} \qquad (58)$$

over ${}^{C}\mathbf{t}_{L}$, ${}^{C}_{L}\mathbf{C}$, α_{i} , ρ_{oi} , θ_{oi} , h_{i} , i = 2, ..., K, α_{1} , and ρ_{o1} . Finally, we minimize this cost function using the Levenberg–Marquardt algorithm (Press et al. 1992).