

A Linear Approximation for Graph-based Simultaneous Localization and Mapping

Original

A Linear Approximation for Graph-based Simultaneous Localization and Mapping / Carlone, L., Aragues, R., Castellanos, J.A., Bona, B.. - (2011). (Robotics: Science and Systems Los Angeles June 27 - June 30, 2011) [10.15607/RSS.2011.VII.006].

Availability:

This version is available at: 11583/2454375 since:

Publisher:

MIT Press

Published

DOI:10.15607/RSS.2011.VII.006

Terms of use:

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

Publisher copyright

(Article begins on next page)

A Linear Approximation for Graph-based Simultaneous Localization and Mapping

Luca Carlone^{*}, Rosario Aragues[†], José A. Castellanos[†] and Basilio Bona[‡]

^{*}CSPP, Laboratorio di Meccatronica, Politecnico di Torino, Torino, Italy - Email: luca.carlone@polito.it

[†]Departamento de Informática e Ingeniería de Sistemas, Instituto de Investigación en Ingeniería de Aragón, Universidad de Zaragoza, Zaragoza, Spain - Email: {raragues, jacaste}@unizar.es

[‡]Dipartimento di Automatica e Informatica, Politecnico di Torino, Torino, Italy - Email: basilio.bona@polito.it

Abstract—This article investigates the problem of Simultaneous Localization and Mapping (SLAM) from the perspective of linear estimation theory. The problem is first formulated in terms of graph embedding: a graph describing robot poses at subsequent instants of time needs to be embedded in a three-dimensional space, assuring that the estimated configuration maximizes measurement likelihood. Combining tools belonging to linear estimation and graph theory, a closed-form approximation to the full SLAM problem is proposed, under the assumption that the relative position and the relative orientation measurements are independent. The approach needs no initial guess for optimization and is formally proven to admit solution under the SLAM setup. The resulting estimate can be used as an approximation of the actual nonlinear solution or can be further refined by using it as an initial guess for nonlinear optimization techniques. Finally, the experimental analysis demonstrates that such refinement is often unnecessary, since the linear estimate is already accurate.

I. INTRODUCTION

In the landscape of Simultaneous Localization and Mapping problems it is possible to distinguish between *online* and *full* SLAM approaches [29]. The former methods estimate the SLAM posterior by recursively including the most recent measurements in the posterior density. This is typically the case of a robot that acquires proprioceptive and exteroceptive measurements from the environment and, at each time step, incrementally includes such information in the posterior describing robot pose and a representation of the surrounding environment. Full SLAM approaches, instead, tackle the problem of retrieving the whole posterior, taking into account at the same time all the available measurements; this batch solution may occur, for instance, when the data, acquired during robot operation, have to be processed off-line to produce a meaningful representation of the scenario.

From the algorithmic point of view, literature on SLAM can be roughly divided in three mainstream approaches: Gaussian filter-based, Particle filters-based and graphical approaches. The first category includes Extended Kalman Filter SLAM (EKF-SLAM) [26], Sparse Extended Information Filter [28], delayed-state EKF-SLAM [9] and related approaches. EKF-like techniques require the linearization of both process and measurement models; the first-order approximation, however, often comes at a price of inconsistency and divergence [3]. Moreover the complexity of the filter is quadratic in the number of landmarks, hence a naive implementation prevents large scale mapping [29]. Finally, we remark that the effects of

partial observability in standard EKF-SLAM implementation are not fully understood, this being witnessed by recent results [15]. The second class of approaches, namely Rao-Blackwellized particle filters [8], is based on a factorization of the posterior and was successfully employed for building both landmark-based and occupancy grid map representations of the environment [29]. Also in this case, fundamental limitations are well known to the robotic community [21]: a finite number of particles may not be able to correctly represent the trajectory posterior, and this issue becomes critical when the uncertainty in the posterior increases. The third category encompasses the techniques in which measurements acquired during robot motion are modeled as constraints in a graphical model. The poses assumed by the robot during its motion correspond to vertices of a directed graph, whereas the constraints are modeled as edges in such graph. Hence SLAM can be stated in terms of maximum likelihood and solved through nonlinear optimization techniques [20]. Although graph-based approaches are recognized to outperform other SLAM techniques in terms of accuracy, they present the classical drawbacks of nonlinear optimization techniques: convergence to a global minimum of the cost function cannot be guaranteed in general and, if the initial guess for optimization is outside the basin of attraction of the global optimum, the iterative process is likely to be stuck in a local minimum.

In this article we address the full SLAM problem in planar setups. Combining tools of linear estimation and graph theory, a closed-form approximation to SLAM is retrieved, under mild assumptions on the structure of the involved covariance matrices. The linear approximation requires no initial guess; moreover, experimental evidence demonstrates that it is accurate in practice, hence it can be used in place of the exact solution or for bootstrapping nonlinear techniques.

II. RELATED WORK

The seminal paper [20] paved the way for several results towards the objective of a sustainable solution to full SLAM. The setup considered by Lu and Milios was based on scan matching for retrieving relative constraints between robot poses, but the use of a Gauss-Newton scheme for optimization prevented its applicability in large scale scenarios. This framework was extended to the general case of a graph containing both robot poses and landmark positions in [27]. Thrun and

Montemerlo showed that it is possible to marginalize variables corresponding to landmarks, hence reducing the problem to the pose estimation setup discussed in [20]. Moreover they enabled the estimation of large maps (10^8 landmarks), applying a conjugate gradient-based optimization. Konolige investigated a reduction scheme for the purpose of improving the computational effort of nonlinear optimization [17]: the optimization process was restricted to poses involved in at least one loop constraint, providing a remarkable advantage for graphs with low connectivity. Frese *et al.* proposed a multilevel relaxation approach for full SLAM [10], allowing to considerably reduce the computational time of optimization by applying a multi-grid algorithm. A further breakthrough in the literature of graph-based approaches came with the use of *incremental pose parametrization*, proposed in [23]. Olson *et al.* showed how a clever selection of the optimization variables can greatly simplify the problem structure, enabling fast computation. Moreover they proposed not to optimize all the constraints at the same time, but to sequentially select each constraint and refine nodes' configuration so to reduce the residual error for such constraint. Grisetti *et al.* extended such framework, taking advantage of the use of stochastic gradient descent in planar and three-dimensional scenarios [11]. Moreover they proposed a parametrization based on a spanning tree of the graph, for speeding up the correction of the residual errors over the network. This approach was shown to be quite resilient to the *orientation wraparound problem* (which causes the configuration estimate to be stuck in a local minimum [11]) and currently constitutes the state-of-the-art.

All the aforementioned techniques are iterative, in the sense that, at each iteration, they solve a local convex approximation of the original problem, and use such local solution to update the configuration. This process is then repeated until the optimization variable converges to a local minimum of the cost function. In particular, when a linear approximation of the residual errors in the cost function is considered, the problem becomes an unconstrained quadratic problem and the local correction can be obtained as solution of a system of linear equations, known as *normal equation* [7]. In such a case the source of complexity stems from the need of repeatedly solving large scale linear systems. A discussion on the performances of direct and iterative methods for solving large scale equations, with application to SLAM, can be found in [7]. Some original attempts to exploit the mathematical structure of the optimization problem and the properties of the underlying graph have been recently proposed [5, 7, 11].

Although the state-of-the-art approaches have been demonstrated to produce impressive results on real world problems, their iterative nature requires the availability of an initial guess for nonlinear optimization, which needs to be sufficiently accurate for the technique to converge to a global solution of the problem. In this work we provide a linear approximation of the full SLAM solution that can be used as initial guess for an iterative nonlinear approach. We remark that the idea of a linear initialization is not new in the literature. For instance, in computer vision, linear methods based on algebraic errors

(e.g., [13, 24]) are often employed for bootstrapping nonlinear techniques (i.e., bundle adjustment [13]). Similar intuitions can be also found within the SLAM community, see [6].

Notation and preliminaries. I_n denotes the $n \times n$ identity matrix, $\mathbf{0}_n$ denotes a (column) vector of all zeros of dimension n . $M_{n \times m}$ denotes a matrix with n rows and m columns and \otimes denotes the Kronecker product. The cardinality of a generic set S is written as $|S|$. A *directed graph* \mathcal{G} is a pair $(\mathcal{V}, \mathcal{E})$, where \mathcal{V} is a finite set of elements, called *vertices* or *nodes*, and \mathcal{E} is a set containing ordered pairs of nodes. A generic element $e \in \mathcal{E}$, referred to as *edge*, is in the form $e = (i, j)$, meaning that edge e , incident on nodes i and j , leaves i (*tail*) and is directed towards node j (*head*) [16]. The number of nodes and edges are denoted with $n + 1$ (the reason for this choice will be clear later) and m , respectively, i.e., $|\mathcal{V}| = n + 1$ and $|\mathcal{E}| = m$. The *incidence matrix* \mathcal{A} of a directed graph is a matrix in $\mathbb{R}^{(n+1) \times m}$ in which each column contains the information of an edge in \mathcal{E} ; in particular the column corresponding to the edge $e = (i, j)$, has the i -th element equal to -1 , the j -th element equal to $+1$ and all the others equal to zero. A directed graph is *connected* (or *weakly connected*) if there exists an undirected path (regardless edges' orientation) connecting any pair of nodes. A *spanning tree* of \mathcal{G} is a subgraph with n edges that contains all the nodes in \mathcal{G} . The edges of \mathcal{G} that do not belong to a given spanning tree of the graph are usually referred to as *chords* [4].

III. PROBLEM FORMULATION

Let $\mathcal{V} = \{v_0, \dots, v_n\}$ be a set of $n + 1$ nodes (representing subsequent poses assumed by a mobile robot) and let $\mathcal{P} = \{p_0, \dots, p_n\}$ denote a corresponding set of absolute poses in a planar setup, i.e., $p_i = [\rho_i^\top \theta_i]^\top \in \mathbb{R}^3$, where $\rho_i = [x_i \ y_i]^\top \in \mathbb{R}^2$ is the Cartesian position of the i -th node, and θ_i is its orientation. We shall call \mathcal{P} a *configuration* of nodes. Suppose that it is possible to measure the relative pose between some nodes, say nodes (i, j) ; in particular node i can measure the pose of j in its local reference frame \mathcal{R}_i , i.e., $\tilde{\xi}_{ij} = p_j \ominus p_i$, where \ominus is a standard pose compounding operator that can be rewritten in explicit form as:

$$p_j \ominus p_i \doteq \begin{bmatrix} R_i^\top (\rho_j - \rho_i) \\ \theta_j - \theta_i \end{bmatrix}, \quad (1)$$

being $R_i \in \mathbb{R}^{2 \times 2}$ a planar rotation matrix of an angle θ_i . Since relative pose measurements are affected by noise, the measured quantities are in the form $\xi_{ij} = \tilde{\xi}_{ij} + \epsilon_{ij}$, where $\epsilon_{ij} \in \mathbb{R}^3$ is a zero mean Gaussian noise, i.e., $\epsilon_{ij} \sim \mathcal{N}(\mathbf{0}, P_{ij})$, being P_{ij} a 3 by 3 covariance matrix. In practice, we can distinguish two kinds of relative pose measurements:

- *odometric constraints*: relative measurements between poses assumed by the robot at subsequent instants of time. These constraints are connected to measurements of the ego-motion (odometry) of the robot and are provided by proprioceptive sensors (wheel odometry, IMU, etc.) or by exteroceptive sensors-based techniques (scan matching, visual features registration, etc.) [29];
- *loop closing constraints*: are connected to place revisiting episodes. This phase clearly requires the use of exteroceptive sensors, among which it is worth mentioning vision sensors and laser range finders.

In this context, we assume the constraints to be given, since the reliable determination of both odometric constraints and

loop closing constraints is still an active research topic [25], whose implications are out of the scope of the present article.

The problem can be naturally modeled using graph formalism: each node in the set \mathcal{V} corresponds to a vertex of a directed graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ (often referred to as *pose graph*), where \mathcal{E} (graph edges) is the set containing the unordered node pairs (i, j) such that a relative pose measurement exists between i and j . By convention, if an edge is directed from node i to node j , the corresponding relative measurement is expressed in the reference frame of node i . We denote with Ξ the set of all available measurements, i.e., $\Xi = \{\xi_{ij}, (i, j) \in \mathcal{E}\}$. The cardinality of the set Ξ is $|\Xi| = |\mathcal{E}| = m$. The relationship between odometric (and loop closing) constraints and graph edges is discussed in detail in Section IV-E.

From the knowledge of relative pose measurements and the corresponding uncertainty, the robot is required to estimate the configuration \mathcal{P} in a given reference frame \mathcal{R}_0 . By convention, we set the initial pose of the robot to be the origin of such reference frame, i.e., $p_0 = [0 \ 0 \ 0]^\top$. In topological graph theory the problem is also referred to as *graph embedding*, *realization* or *drawing*, depending on the context and on problem constraints [12]. Therefore our objective is to determine an estimated configuration $\mathcal{P}^* = \{p_1^*, \dots, p_n^*\}$, that maximizes the likelihood of the observations. It is common in literature to assume independence among observations, hence we can write the log-likelihood function of a configuration as:

$$\ln \mathcal{L}(\mathcal{P}|\Xi) = \sum_{(i,j) \in \mathcal{E}} \ln \pi(\xi_{ij}|\mathcal{P}), \quad (2)$$

where $\pi(\xi_{ij}|\mathcal{P})$ is the conditional probability density of the measurement ξ_{ij} given the node configuration. Since we assumed that the involved densities are Gaussians, it is easy to demonstrate that maximizing the likelihood function (2) is equivalent to minimize the sum of the weighted residual errors:

$$f(\mathcal{P}) = \sum_{(i,j) \in \mathcal{E}} (p_j \ominus p_i - \xi_{ij})^\top P_{ij}^{-1} (p_j \ominus p_i - \xi_{ij}). \quad (3)$$

The full SLAM problem is hence formulated as a minimization of the nonlinear non-convex cost function (3), i.e., the optimal configuration is $\mathcal{P}^* = \arg \min f$. The nonlinearity of the function $f(\mathcal{P})$ is due to the structure of the pose compounding operator (1) and, in particular, to the nonlinear terms in the orientations of the robot.

IV. A LINEAR ESTIMATION FRAMEWORK FOR SLAM

We first observe that each relative pose measurement is composed by a two-valued vector (*relative position*) corresponding to the first two components of ξ_{ij} , see (1), and a scalar (*relative orientation*). Therefore we can rewrite each measurements as $\xi_{ij} = [\Delta_{ij}^l \ \delta_{ij}]^\top$, where the superscript l denotes that the relative position vector is expressed in a local frame. According to the notation introduced so far the cost function (3) can be rewritten as:

$$\sum_{(i,j) \in \mathcal{E}} \begin{bmatrix} R_i^\top (\rho_j - \rho_i) - \Delta_{ij}^l \\ (\theta_j - \theta_i) - \delta_{ij} \end{bmatrix}^\top P_{ij}^{-1} \begin{bmatrix} R_i^\top (\rho_j - \rho_i) - \Delta_{ij}^l \\ (\theta_j - \theta_i) - \delta_{ij} \end{bmatrix}.$$

In this context we assume that the relative position information and the relative orientation measurements are independent, i.e. $P_{ij} = \text{diag}(P_{\Delta_{ij}^l}, P_{\delta_{ij}})$. This hypothesis is a technical requirement for attacking the problem, see also Remark 2 in Section V. Under this hypothesis the cost function $f(\mathcal{P})$ becomes:

$$\sum_{(i,j) \in \mathcal{E}} \begin{bmatrix} R_i^\top (\rho_j - \rho_i) - \Delta_{ij}^l \\ (\theta_j - \theta_i) - \delta_{ij} \end{bmatrix}^\top P_{\Delta_{ij}^l}^{-1} \begin{bmatrix} R_i^\top (\rho_j - \rho_i) - \Delta_{ij}^l \\ (\theta_j - \theta_i) - \delta_{ij} \end{bmatrix} + \sum_{(i,j) \in \mathcal{E}} [(\theta_j - \theta_i) - \delta_{ij}]^\top P_{\delta_{ij}}^{-1} [(\theta_j - \theta_i) - \delta_{ij}]. \quad (4)$$

In order to put the previous formulation in a more compact form, let us number the nodes pairs, for which a relative pose measurement is available, from 1 to m ; let us stack all the relative position measurements in the vector $\Delta^l = [(\Delta_1^l)^\top, (\Delta_2^l)^\top, \dots, (\Delta_m^l)^\top]^\top$, and all the relative orientation measurements in the vector $\delta = [\delta_1, \delta_2, \dots, \delta_m]^\top$. Repeating the same procedure for all the positions and orientations assumed by the robot we get the *nodes' position* $\rho^+ = [\rho_0^\top, \rho_1^\top, \dots, \rho_n^\top]^\top$ and the *nodes' orientation* $\theta^+ = [\theta_0, \theta_1, \dots, \theta_n]^\top$. Therefore it can be easily verified that equation (4) can be rewritten as:

$$f(\mathcal{P}) = (\mathcal{A}_2^\top \rho^+ - R \Delta^l)^\top (R P_{\Delta^l} R^\top)^{-1} (\mathcal{A}_2^\top \rho^+ - R \Delta^l) + (\mathcal{A}^\top \theta^+ - \delta)^\top P_\delta^{-1} (\mathcal{A}^\top \theta^+ - \delta), \quad (5)$$

where \mathcal{A} is the incidence matrix of the graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, $\mathcal{A}_2 = \mathcal{A} \otimes I_2$ is an expanded form of the incidence matrix, see [1], $P_{\Delta^l} = \text{diag}(P_{\Delta_1^l}, P_{\Delta_2^l}, \dots, P_{\Delta_m^l})$, $P_\delta = \text{diag}(P_{\delta_1}, P_{\delta_2}, \dots, P_{\delta_m})$, and $R = R(\theta)$ is a block diagonal matrix containing the rotation matrices that transform the corresponding local measurements in the global frame, i.e., the non-zero entries of R are in positions $(2k-1, 2k-1)$, $(2k-1, 2k)$, $(2k, 2k-1)$, $(2k, 2k)$, $k = 1, \dots, m$, and the k -th diagonal block contains the rotation matrix converting the k -th relative position measurement in the global frame.

Equation (5) can be further simplified by recalling the choice of the reference frame $\rho_0 = [0 \ 0]^\top$ and $\theta_0 = 0$, leading to the following relation:

$$f(\mathcal{P}) = (\mathcal{A}_2^\top \rho - R \Delta^l)^\top (R P_{\Delta^l} R^\top)^{-1} (\mathcal{A}_2^\top \rho - R \Delta^l) + (\mathcal{A}^\top \theta - \delta)^\top P_\delta^{-1} (\mathcal{A}^\top \theta - \delta), \quad (6)$$

where ρ and θ are the vectors obtained from ρ^+ and θ^+ by deleting the elements ρ_0 and θ_0 , respectively, whereas \mathcal{A} and \mathcal{A}_2 are *reduced* incidence matrices, obtained by deleting the rows corresponding to θ_0 and ρ_0 from \mathcal{A} and \mathcal{A}_2 , respectively.

We observe that the minimization of the cost function (6) is equivalent to find the solution (in the least-squares sense) to the following system of equations:

$$\begin{cases} \mathcal{A}_2^\top \rho = R(\theta) \Delta^l \\ \mathcal{A}^\top \theta = \delta \end{cases}. \quad (7)$$

The system (7) includes m vector-valued relations (or constraints) on relative positions and m scalar constraints. When a solution exactly satisfies all the constraints in (7), the corresponding minimum of the cost function (6) is zero, otherwise, the searched solution has to minimize the weighted residual errors on the constraints. The nonlinear nature of the problem at hand is connected with the matrix $R(\theta)$, which contains trigonometrical functions of nodes' orientation.

A. Solving SLAM with Linear Estimation Theory

We now provide a procedure that allows to retrieve an approximation of the actual nonlinear solution of (6). The procedure can be summarized in three phases:

- 1) Solve the following linear estimation problem:

$$\mathcal{A}^\top \theta = \delta, \quad (8)$$

from which the suboptimal orientation estimate $\hat{\theta}$ and its covariance matrix $P_{\hat{\theta}}$ can be obtained.

- 2) Estimate the relative position measurements in the global reference frame:

$$z = \begin{bmatrix} \hat{R} & \mathbf{0}_{2m \times n} \\ \mathbf{0}_{2m \times n}^\top & I_n \end{bmatrix} \begin{bmatrix} \Delta^l \\ \hat{\theta} \end{bmatrix} = \begin{bmatrix} g_1(\Delta^l, \theta) \\ g_2(\theta) \end{bmatrix}_{\theta=\hat{\theta}}, \quad (9)$$

with $\hat{R} = R(\hat{\theta})$; compute the corresponding uncertainty (preserving the correlation with the orientation estimate):

$$P_z = H \begin{bmatrix} P_{\Delta^l} & \mathbf{0}_{2m \times n} \\ \mathbf{0}_{2m \times n}^\top & P_{\hat{\theta}} \end{bmatrix} H^\top, \quad (10)$$

where H is the Jacobian of the transformation in (9):

$$H \doteq \begin{bmatrix} \frac{\partial g_1}{\partial \Delta^l} & \frac{\partial g_1}{\partial \theta} \\ \frac{\partial g_2}{\partial \Delta^l} & \frac{\partial g_2}{\partial \theta} \end{bmatrix} = \begin{bmatrix} \hat{R} & J \\ \mathbf{0}_{2m \times n}^\top & I_n \end{bmatrix}. \quad (11)$$

- 3) Solve the linear estimation problem in the unknown $p = [\rho^\top \theta^\top]^\top$, given z (9) and P_z (10):

$$z = \begin{bmatrix} A_2^\top & \mathbf{0}_{2m \times n} \\ \mathbf{0}_{2m \times n}^\top & I_n \end{bmatrix} \begin{bmatrix} \rho \\ \theta \end{bmatrix} = B^\top p, \quad (12)$$

from which the solution $p^* = [(\rho^*)^\top (\theta^*)^\top]^\top$ and the corresponding uncertainty can be retrieved.

In the rest of this section we will show how the proposed procedure provides an approximate solution to SLAM with graphical models. For sake of readability each phase of the procedure will be discussed in detail, hence the proof of the result will be clear at the end of Section IV-D.

B. Phase 1: Linear Orientation Estimation

Observing the structure of problem (7) it can be seen that the second part of the system, including the relative orientation measurements, is linear in the unknown variable θ . The first phase of the procedure requires the solution of such subset of constraints; this is a standard linear estimation problem: given the matrix A , the measurements δ and the corresponding covariance matrix P_δ , the objective is to provide an estimate $\hat{\theta}$ of the unknown θ . According to well known results in linear estimation theory [22], the Best Linear Unbiased Estimator (BLUE) for θ in (8) and the corresponding covariance are:

$$\hat{\theta} = (AP_\delta^{-1}A^\top)^{-1}AP_\delta^{-1}\delta, \quad P_{\hat{\theta}} = (AP_\delta^{-1}A^\top)^{-1}. \quad (13)$$

The following results establish existence and uniqueness of the solution of problem (8). We recall that similar results can be found in literature, see [1] and the references therein.

Lemma 1 (Connectivity of pose graph): The pose graph \mathcal{G} modeling robot poses, including odometric and loop closing constraints is connected.

Proof. The proof is trivial since the path connecting all the nodes is actually the trajectory traveled by the robot; in this sense the edges corresponding to the odometric constraints assure graph connectivity. \square

Proposition 1 (Existence of the solution for phase 1): The solution of the first step in the procedure IV-A is unique if \mathcal{G} is connected and a node orientation is supposed to be known.

Proof. The incidence matrix \mathcal{A} of a connected graph \mathcal{G} has rank n (with $|\mathcal{V}| = n + 1$) and becomes full rank as soon as one row is deleted, see [4]. As we have just shown in (6), assuming a node to have known orientation (e.g., the first node is assumed to be the absolute reference frame) allows to reduce the incidence matrix by deleting

the corresponding row. Therefore, A is full rank and P_δ^{-1} is positive definite by definition, then $AP_\delta^{-1}A^\top$ is positive definite, hence invertible, see [14]. If the matrix is invertible, the solution of the linear problem can be uniquely determined as $\hat{\theta} = (AP_\delta^{-1}A^\top)^{-1}AP_\delta^{-1}\delta$. \square

It is worth observing that once the absolute orientation of the robot is known, also the first equation in (7) becomes linear in the unknown ρ . Therefore, using $\hat{\theta}$ as the actual nodes' orientation, we can also compute an estimate for nodes' position $\hat{\rho}$, using linear estimation framework:

$$\hat{\rho} = \left[A_2 \left(\hat{R}P_{\Delta^l} \hat{R}^\top \right)^{-1} A_2^\top \right]^{-1} A_2 \left(\hat{R}P_{\Delta^l} \hat{R}^\top \right)^{-1} \hat{R} \Delta^l, \quad (14)$$

where $\hat{R} = R(\hat{\theta})$. Also in this case, the proof of existence and uniqueness of the solution directly stems from Proposition 1. However, the first equation in (7) also constrains the orientations of the robot, thus the estimate $\hat{p} = [\hat{\rho}^\top \hat{\theta}^\top]^\top$ constitutes a suboptimal solution, in which the influence of the first equation on the estimated orientations is neglected.

C. Phase 2: First-order Error Propagation

The second phase simply provides an estimate of the relative position measurements in the global reference frame: the vector z will be in the form $z = [(\Delta^g)^\top \hat{\theta}^\top]^\top$ where $\Delta^g = \hat{R} \Delta^l$ is the vector containing the relative node position expressed in the absolute reference frame \mathcal{R}_0 . The corresponding covariance matrix can be obtained by a first-order propagation of the uncertainty, hence we can rewrite (10) in explicit form as:

$$P_z = \begin{bmatrix} P_{\Delta^g} + JP_{\hat{\theta}}J^\top & JP_{\hat{\theta}} \\ P_{\hat{\theta}}J^\top & P_{\hat{\theta}} \end{bmatrix},$$

where $P_{\Delta^g} \doteq \hat{R}P_{\Delta^l}\hat{R}^\top$, and J and $P_{\hat{\theta}}$ are defined as in (11) and (13). The trick of including the orientation estimates in z is useful for preserving the correlation between the relative position measurements (expressed in the global frame) and the corresponding angular information. As we will see in the following section, such correlation terms play a fundamental role in the outcome of the procedure.

D. Phase 3: Linear Position Estimation

We now want to show that the last phase allows to correct the sub-optimal configuration estimate \hat{p} , leading it towards a minimum of the cost function, i.e.

$$\theta^* = \hat{\theta} + \tilde{\theta}, \quad \rho^* = \hat{\rho} + \tilde{\rho},$$

in which $\tilde{\theta}$ and $\tilde{\rho}$ are first-order correction terms. In order to proceed in the demonstration we need to compute $p^* = [(\rho^*)^\top (\theta^*)^\top]^\top$ as solution of the linear system (12):

$$p^* = \begin{bmatrix} \rho^* \\ \theta^* \end{bmatrix} = (BP_z^{-1}B^\top)^{-1}BP_z^{-1}z. \quad (15)$$

Also in this case, the demonstration of uniqueness of the solution (15) can be easily deduced from Proposition 1. To write in explicit form ρ^* and θ^* , we first compute the information matrix P_z^{-1} :

$$P_z^{-1} = \begin{bmatrix} P_{\Delta^g}^{-1} & -P_{\Delta^g}^{-1}J \\ -J^\top P_{\Delta^g}^{-1} & P_{\hat{\theta}}^{-1} + J^\top P_{\Delta^g}^{-1}J \end{bmatrix}.$$

The previous inverse can be performed using blockwise matrix inversion, see [14]. We can now compute the information matrix $I_{p^*} = (BP_z^{-1}B^\top)$:

$$I_{p^*} = \begin{bmatrix} P_{\hat{\rho}}^{-1} & -A_2 P_{\Delta_g}^{-1} J \\ -J^\top P_{\Delta_g}^{-1} A_2^\top & P_{\hat{\theta}}^{-1} + J^\top P_{\Delta_g}^{-1} J \end{bmatrix},$$

where $P_{\hat{\rho}} \doteq (A_2 P_{\Delta_g}^{-1} A_2^\top)^{-1}$. The inverse of I_{p^*} , namely P_{p^*} , is actually the covariance matrix of our estimated configuration, and it is in the form:

$$P_{p^*} = \begin{bmatrix} P_{\rho^*} & P_{\rho^*, \theta^*} \\ P_{\rho^*, \theta^*}^\top & P_{\theta^*} \end{bmatrix},$$

where:

$$\begin{aligned} P_{\theta^*} &\doteq \left(P_{\hat{\theta}}^{-1} + J^\top P_{\Delta_g}^{-1} J - J^\top P_{\Delta_g}^{-1} A_2^\top P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} J \right)^{-1} \\ P_{\rho^*} &= P_{\hat{\rho}} + P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} J P_{\theta^*} J^\top P_{\Delta_g}^{-1} A_2^\top P_{\hat{\rho}} \\ P_{\rho^*, \theta^*} &= P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} J P_{\theta^*}. \end{aligned}$$

From the expression of P_{p^*} and P_z^{-1} it is straightforward to compute the estimates (15) as:

$$\begin{aligned} \theta^* &= (AP_{\hat{\theta}}^{-1}A^\top)^{-1}AP_{\hat{\theta}}^{-1}\delta + \\ &+ P_{\theta^*} J^\top P_{\Delta_g}^{-1} (A_2^\top P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} - I_{2m}) \hat{R} \Delta^l, \end{aligned} \quad (16)$$

and:

$$\begin{aligned} \rho^* &= \left[A_2 \left(\hat{R} P_{\Delta_l} \hat{R}^\top \right)^{-1} A_2^\top \right]^{-1} A_2 \left(\hat{R} P_{\Delta_l} \hat{R}^\top \right)^{-1} \hat{R} \Delta^l + \\ &+ P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} J P_{\theta^*} J^\top P_{\Delta_g}^{-1} (A_2^\top P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} - I_{2m}) \hat{R} \Delta^l. \end{aligned}$$

By simple inspection it is possible to verify that the obtained estimate θ^* is already in the form $\theta^* = \hat{\theta} + \tilde{\theta}$, since the first summand in (16) coincides with (13). The same consideration holds for ρ^* with respect to equation (14). Therefore we can write $\hat{\theta}$ and $\hat{\rho}$ as:

$$\begin{aligned} \hat{\theta} &= P_{\theta^*} J^\top P_{\Delta_g}^{-1} (A_2^\top P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} - I_{2m}) \hat{R} \Delta^l \\ \hat{\rho} &= P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} J P_{\theta^*} J^\top P_{\Delta_g}^{-1} (A_2^\top P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} - I_{2m}) \hat{R} \Delta^l. \end{aligned} \quad (17)$$

We are now ready to state the following key result.

Theorem 1: The outcome of the three-phases procedure approximates the solution of the nonlinear optimization problem (6). In particular, the third phase produces a local correction of a sub-optimal configuration estimate computed in the first phase, leading it towards a minimum of the cost function.

Proof. We have already shown that the final solution is composed by the suboptimal solution plus a correction term. Now the demonstration reduces to verify that $\tilde{p} = [\tilde{\rho}^\top \tilde{\theta}^\top]^\top$ is a local solution of our optimization problem around $\hat{p} = [\hat{\rho}^\top \hat{\theta}^\top]^\top$. In order to prove this point we compute a first-order approximation of the residual errors in (6) around the suboptimal solution \hat{p} :

$$\begin{aligned} f(\mathcal{P}) &\approx \left(A_2^\top \hat{\rho} + A_2^\top \delta_\rho - \hat{R} \Delta^l - J \delta_\theta \right)^\top \\ &\left(\hat{R} P_{\Delta_l} \hat{R}^\top \right)^{-1} \left(A_2^\top \hat{\rho} + A_2^\top \delta_\rho - \hat{R} \Delta^l - J \delta_\theta \right) + \\ &+ \left(A^\top \hat{\theta} + A^\top \delta_\theta - \delta \right)^\top P_\delta^{-1} \left(A^\top \hat{\theta} + A^\top \delta_\theta - \delta \right), \end{aligned} \quad (18)$$

where δ_θ and δ_ρ are the displacements from the linearization point. This convex approximation can now be easily minimized by taking the first derivative with respect to the

optimization variables δ_θ and δ_ρ and imposing it to be zero. The minimum of (18) corresponds to:

$$\begin{aligned} \delta_\theta &= P_{\theta^*} J^\top P_{\Delta_g}^{-1} \left(A_2^\top P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} - I_{2m} \right) \hat{R} \Delta^l \\ \delta_\rho &= P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} J P_{\theta^*} J^\top P_{\Delta_g}^{-1} \left(A_2^\top P_{\hat{\rho}} A_2 P_{\Delta_g}^{-1} - I_{2m} \right) \hat{R} \Delta^l, \end{aligned}$$

which can be easily seen to coincide, respectively, with $\tilde{\theta}$ and $\tilde{\rho}$, see (17), thus proving our thesis. \square

Remark 1: A direct (iterative) method (see [7] and the references therein) would require solving N linear problems of size $3n$ (in which the iterations N may increase arbitrarily, depending on the initial guess); our approach solves a smaller problem (on robot orientations - size n) and a problem of the same size of a single step of the direct method. The advantage, now, is that the first-order correction provided by the last phase of the proposed approach, refines a sub-optimal solution which in practice is already close to a global minimum. Therefore the approach is expected to be accurate, while reducing the risk of being trapped in a local minimum. Furthermore, the sub-optimal solution needs not be computed explicitly: for the purpose of proving Theorem 1 we reported $\hat{\rho}$, but it is not actually computed in the three-phases procedure.

E. Regularization and Existence of the Correction Factors

We now discuss a crucial point of the proposed approach, which is connected with the periodic nature of the angular information, i.e., robot orientations are defined up to $2k\pi$, $k \in \mathbb{Z}$. Let us introduce the discussion with an example: consider a simple scenario, in which a robot travels along a circumference (in anticlockwise direction) coming back to the starting point. In a noiseless case, summing up all the relative orientation measurements from the one taken with respect to the first node, to the loop closing constraint, referred to the last node, we obtain 2π . This is because we sum small angular measurements which are defined in $(-\pi, \pi]$. However, the loop closing constraint is expected to link the last pose with the initial pose, whose orientation was set by convention to zero. The linear estimation framework presented so far cannot recognize that the angles 0 and 2π actually correspond to the same orientation, hence tries to impose contrasting constraints, producing meaningless results. An easy solution to the previous problem consists in adding a correction factor, in the form $2k\pi$, $k \in \mathbb{Z}$, to some orientation measurements. These correction factors do not alter measurement content, because of the periodicity of the angular information, but let the relative orientation measurements sum up to zero (this property will be lately referred to as *zero-sum* property). Hence the input data provided to the problem solver are consistent and the estimated configuration is correct.

In the rest of this section we will prove the existence of suitable correction factors for any connected graph and we will describe a methodology for retrieving such correction terms. Before presenting the main result (Theorem 3) let us introduce some specific concepts from graph theory.

A *cycle* is a subgraph in which every node appears in an even number of edges. A *circuit* is a cycle in which every node appears exactly in two edges. We can represent a (*directed*) circuit as a vector c_i of m elements in which the k -th element is $+1$ or -1 if edge k belongs to the circuit and it is traversed respectively forwards (from tail to head) or backwards, and 0 if it does not appear in the circuit (notice that the ordering of the edges in c_i is arbitrary).

Definition 1: Given a directed graph \mathcal{G} and a spanning tree T of \mathcal{G} , a *fundamental circuit* is a circuit composed by a

chord (i, j) of \mathcal{G} with respect to T and the unique path in T connecting i and j .

A *cycle basis* of \mathcal{G} is the smallest set of circuits such that any cycle in the graph can be written as a combination of the circuits in the basis. The space spanned by the vectors in the basis is called *cycle space*.

Theorem 2: The set of the fundamental circuits of a directed graph constitutes a cycle basis for \mathcal{G} .

The proof of the previous theorem can be found in several books of graph theory, see [4]. We already mentioned that a spanning tree T of a connected graph \mathcal{G} contains exactly n edges. Accordingly, the number of chords, hence of fundamental circuits in \mathcal{G} , is $m - n$.

Corollary 1: The dimension of the cycle space of a connected graph G is $d = m - n$, and it is usually called *cyclomatic number* of the graph [16].

Corollary 2: Ordering the edges of a connected directed graph \mathcal{G} from 1 to m , so that the first $m - n$ elements are the chords of \mathcal{G} with respect to a given spanning tree T and the last n elements are the edges of T , the matrix containing all the vectors c_i corresponding to the fundamental circuits can be written as:

$$\mathcal{C} = [c_1 \ c_2 \ \dots \ c_d]^\top = [I_d \ \mathcal{B}], \quad (19)$$

where I_d is the identity matrix of dimension d and \mathcal{B} is a matrix with elements in $\{-1, 0, 1\}$. \mathcal{C} is usually referred to as *cycle basis matrix*.

The previous result is a direct consequence of the structure of the fundamental circuits, each one containing a single chord and a collection of edges in the spanning tree [4]. With slight abuse of notation, in the following, \mathcal{C} will denote both the cycle basis and the cycle basis matrix.

According to the machinery introduced so far, we notice that the zero-sum property essentially requires that the sum of the relative orientation measurements along every cycle in the graph is zero, instead of $2k\pi, k \in \mathbb{Z} \setminus \{0\}$. Hence we can state the following theorem that holds under the assumption of noiseless angular measurements.

Theorem 3 (Existence of correction factors): Given the relative orientation measurements $\delta = [\delta_1, \delta_2, \dots, \delta_m]^\top$ there exists a correction vector $\nu = [\nu_1, \nu_2, \dots, \nu_m]^\top$ so that the corrected measurements $\bar{\delta} = [\delta_1 + \nu_1, \delta_2 + \nu_2, \dots, \delta_m + \nu_m]^\top$ satisfy the zero-sum property.

The elements of $\bar{\delta}$ are referred to as *regularized relative orientation measurements*. The process of compensating the relative orientation measurements is named *regularization*.

Proof. Let us start by formulating the zero-sum property in a more familiar way. A necessary and sufficient condition for the zero-sum property to be satisfied for all the cycles in the graph is that it is satisfied for the cycles in the cycle basis, see [16]. Let us consider the cycle basis composed by the fundamental circuits; for the zero-sum property to hold true, the corrected measurements $\bar{\delta}$ have to satisfy:

$$c_i^\top \bar{\delta} = 0 \quad \forall c_i \in \mathcal{C}. \quad (20)$$

Roughly speaking, if the sum of the relative orientation measurements is zero for the edges in the fundamental circuits, this property is true for every cycle in the graph. Equation (20) can be written in compact form using the cycle basis matrix:

$$\mathcal{C}\bar{\delta} = \mathbf{0}_d. \quad (21)$$

According to the definition of regularized measurements we can rewrite (21) as:

$$\mathcal{C}(\delta + \nu) = \mathbf{0}_d \implies \mathcal{C}\nu = -\mathcal{C}\delta. \quad (22)$$

The right hand side will contain the sum of the original measurements for each fundamental circuit. In a noiseless case, the vector $\mathcal{C}\delta$ contains terms in the form $2k\pi, k \in \mathbb{Z}$. Since the cycle basis matrix \mathcal{C} can be computed from the graph and δ is a given of the problem, the only unknown of (22) is ν and the existence of a proper regularization is reduced to the demonstration of existence of a solution to system (22). Recalling (19), it is easy to show that a solution to system (22), is $\nu = [-(\mathcal{C}\delta)^\top \ \mathbf{0}_n^\top]^\top$:

$$\mathcal{C}\nu = [I_d \ \mathcal{B}] [-(\mathcal{C}\delta)^\top \ \mathbf{0}_n^\top]^\top = -\mathcal{C}\delta, \quad (23)$$

hence proving our thesis. \square

We notice that the aforementioned solution only requires to correct the angular measurements corresponding to the chords, without any modification to the edges in the spanning tree. We remark that the correction terms are in the form $2k\pi, k \in \mathbb{Z}$, hence the regularization procedure does not alter the information content of the orientation measurement. It is now evident that, in the case of noisy relative measurements, condition (20) cannot be met exactly: our approach, in fact, will be in charge of compensating the measurement errors by minimizing a suitable cost function. Accordingly, the term $\mathcal{C}\delta$ in (22) will not contain exact multiples of 2π . However, a simple rounding to the closest multiple of 2π allows to retrieve the desired correction factors. One may argue that, if the noise is large, it is not possible to discern the desired correction factors, since the rounding cannot compensate measurement errors; however this issue was not found to be relevant in common applications (see experimental section): as it will be clear in a while, the impossibility to determine the correct multiple of 2π means that the amount of noise is so high that the robot, revisiting a past pose, is not able to discern how many times the robot turned around itself (i.e., completed 2π turns) since the previous visit. Note that this result also sheds some light on the so called *orientation wraparound* problem [11], which is known to prevent convergence in iterative approaches.

We now tailor the previous formulation to the SLAM setup. We state the following facts, whose demonstration is omitted, since it can be easily inferred from the basic definitions.

Proposition 2: The edges corresponding to odometric constraints in the pose graph constitute a spanning tree, T , for the connected graph \mathcal{G} , describing the SLAM problem.

Corollary 3: The edges corresponding to loop closing constraints are chords of the pose graph \mathcal{G} , with respect to the spanning tree T .

Therefore, according to Corollary 3 and Theorem 3, we can regularize the orientation measurements by simply correcting loop closing constraints; in particular, a loop closing relation, constraining two robot poses, has to be corrected taking into account the number of complete turns (2π turns) the robot did when traveling from the first to the second pose.

V. EXPERIMENTAL VALIDATION

In this section we will show an application of the proposed methodology to the Intel Research Lab dataset [18]. The solution of the full SLAM problem from real data requires three steps:

- 1) Odometric and loop closing constraints are extracted from real sensor data;
- 2) Regularization is performed for making the orientation measurements consistent;
- 3) The three-phases procedure is applied to solve the graph embedding problem.

Number of nodes	Number of odometric constraints	Number of loop closing constraints	Number of chords needing regularization
1228	1227	278	198

TABLE I
NUMBER OF ODOMETRIC CONSTRAINTS, LOOP CLOSING CONSTRAINTS, AND REGULARIZED ORIENTATION MEASUREMENTS FOR INTEL DATASET.




Fig. 1. Intel Research Lab dataset: odometric constraints (solid lines), corresponding to the odometric trajectory corrected with scan matching, and loop closing constraints (dotted lines).

In the following, we will tailor the formulation to the particular application scenario, in which a mobile robot is equipped with wheel encoders and a laser range finder. The odometric constraints are obtained by refining the wheel odometry measurements with a scan matching procedure, see [19]. The loop closing constraints are selected from the relations available at [18]. The number of odometric and loop closing constraints is reported in Table I.

In Figure 1 we show the odometric trajectory of the robot (solid line) obtained from wheel encoders estimates and scan matching. The figure also shows, as dotted lines, the edges corresponding to loop closing constraints. The scan matching algorithm is only able to enforce local consistency by aligning the laser readings acquired at subsequent poses, thus failing in producing a globally consistent map. Once the relative pose information (δ, Δ^l) is available, it is possible to perform regularization for orientation measurements. Computing the cycle basis \mathcal{C} and the correction terms $\nu = [-(\mathcal{C}\delta)^T \mathbf{0}_p^T]^T$, it is possible to obtain the regularized measurements $\hat{\delta}$. In Table I we show the number of loop closing constraints for which a correction factor $2k\pi, k \in \mathbb{Z} \setminus \{0\}$ was needed. The configuration estimated with the three-phases procedure and the corresponding occupancy grid map are reported in




Fig. 2. Intel Research Lab dataset: (a) estimated node configuration, (b) occupancy grid map obtained by associating the corresponding laser scan to each node of the estimated configuration.




Fig. 3. Errors versus iterations for TORO (solid lines) and errors for the three-phases procedure (dashed line). (a) translation error squared (η_c); (b) angular error squared (η_a).

	$P_{ij} = \text{diag}(P_{\Delta_{ij}^l}, P_{\delta_{ij}})$		$P_{ij} = I_3$	
	η_c [m ²]	η_a [rad ²]	η_c [m ²]	η_a [rad ²]
TORO (100 iter.)	$2.60 \cdot 10^{-3}$	$2.30 \cdot 10^{-4}$	$6.16 \cdot 10^{-4}$	$1.62 \cdot 10^{-4}$
TORO (1000 iter.)	$1.90 \cdot 10^{-3}$	$2.18 \cdot 10^{-4}$	$5.83 \cdot 10^{-4}$	$1.62 \cdot 10^{-4}$
Proposed approach	$0.91 \cdot 10^{-3}$	$2.25 \cdot 10^{-4}$	$3.36 \cdot 10^{-4}$	$2.00 \cdot 10^{-4}$

TABLE II
BENCHMARK METRICS [19] FOR THE INTEL DATASET: TRANSLATION ERROR SQUARED (η_c) AND ANGULAR ERROR SQUARED (η_a).

Figure 2. For a quantitative evaluation we report the values of the SLAM benchmark metrics proposed in [19]. Without entering into details we mention that such metrics provide a tool for comparing the SLAM approaches in terms of accuracy. In Table II we show the values of the *constraint satisfaction metrics*, comparing the proposed solution with the *Tree-based netwORk Optimizer* (TORO), which is available online, see [11]. We consider two scenarios, in which different measurements covariance matrices are considered: in one case the covariance matrix is chosen to be the identity matrix ($P_{ij} = I_3$), whereas, in the other scenario, measurements covariance is in the form $P_{ij} = \text{diag}(P_{\Delta_{ij}^l}, P_{\delta_{ij}})$, and the corresponding uncertainties are assumed to be proportional to the respective measurements, e.g. bigger displacements correspond to higher uncertainty. For this last case, in Figure 3 we plot the errors versus iterations for TORO, compared with the corresponding statistics obtained with the three-phases procedure. It is now clear that the proposed approach is accurate in practice. Further experiments and simulations can be found in [2], confirming the results presented so far.

Remark 2: The assumption of independent position and orientation measurements holds true when dealing with holonomic platforms. For non-holonomic platforms it constitutes an approximation, but several state-of-the-art techniques have been demonstrated to produce excellent results, even under stricter assumptions on the covariance structure (e.g. spherical covariances in [11]).

VI. CONCLUSION

The contribution of this work is twofold: *in primis* we combine tools of linear estimation and graph theory, to gain a deep insight on SLAM with graphical models; then we apply this theoretical analysis for retrieving an approximate solution to the full SLAM problem, under mild assumptions on the structure of the involved covariance matrices. The proposed estimation process requires no initial guess and is formally demonstrated to admit solution when applied to the

embedding of the pose graph. Experiments on a real dataset confirm the validity of the theoretical analysis. It is possible to consider the proposed approach as a linear initialization method for iterative optimization or as a stand-alone tool. The impact of the proposed methodology concerns different aspects of the SLAM problem: (i) the solution only requires basic linear algebra machinery hence it can be envisioned to apply complexity reduction techniques (Cholesky decomposition, QR factorization, blockwise inversion, etc.) or to use parallel computational architectures (e.g., FPGA) making the approach suitable for large scale mapping; (ii) the paper provides an insight on the orientation wraparound problem: in large-scale applications, one cannot expect to let the robot travel for a long time without incurring in this issue; (iii) the linearity of the framework provides a chance for devising an incremental solution to graph-based SLAM.

ACKNOWLEDGMENTS

The authors wish to thank Frank Dellaert and four anonymous reviewers for comments that helped to improve the final version of this article. This work was partially funded by Ministero dell'Istruzione, dell'Università e della Ricerca (MIUR) under MEMONET National Research Project, by Regione Piemonte under MACP4log Grant (RU/02/26), by projects MICINN-FEDER DPI2009-08126 and DPI2009-13710, and grant MEC BES-2007-14772.

REFERENCES

- [1] P. Barooah and J.P. Hespanha. Estimation on graphs from relative measurements. *IEEE Control Systems Magazine*, 27(4):57–74, 2007.
- [2] L. Carlone, R. Aragues, J.A. Castellanos, and B. Bona. A first-order solution to simultaneous localization and mapping with graphical models. In *Proc. of the IEEE International Conf. on Robotics and Automation*, 2011.
- [3] J.A. Castellanos, J. Neira, and J.D. Tardós. Limits to the consistency of EKF-based SLAM. In *5th IFAC Symp. on Intelligent Autonomous Vehicles*, 2004.
- [4] W. Chen. *Graph Theory and Its Engineering Applications*. Advanced Series in Electrical and Computer Engineering, 1997.
- [5] F. Dellaert and M. Kaess. Square root SAM: Simultaneous localization and mapping via square root information smoothing. *Int. J. Robot. Res.*, 25(12):1181–1203, 2006.
- [6] F. Dellaert and A. Stoupe. Linear 2D localization and mapping for single and multiple robots. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 2002.
- [7] F. Dellaert, J. Carlson, V. Ila, K. Ni, and C. Thorpe. Subgraph-preconditioned conjugate gradients for large scale SLAM. In *Proc. of the IEEE-RSJ Int. Conf. on Intelligent Robots and Systems*, 2010.
- [8] A. Doucet, N. de Freitas, K. Murphy, and S. Russel. Rao-Blackwellized particle filtering for dynamic bayesian networks. In *Proc. of the Conf. on Uncertainty in Artificial Intelligence*, pages 176–183, 2000.
- [9] R.M. Eustice, H. Singh, and J.J. Leonard. Exactly sparse delayed-state filters for view-based SLAM. *Int. J. Robot. Res.*, 22(6):1100–1114, 2006.
- [10] U. Frese, P. Larsson, and T. Duckett. A multilevel relaxation algorithm for simultaneous localization and mapping. *IEEE Trans. on Robotics*, 21(2):196–207, 2005.
- [11] G. Grisetti, C. Stachniss, and W. Burgard. Non-linear constraint network optimization for efficient map learning. *IEEE Trans. on Intelligent Transportation Systems*, 10(3):428–439, 2009.
- [12] J.L. Gross and T.W. Tucker. *Topological graph theory*. Wiley Interscience, 1987.
- [13] R.I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521623049, 2000.
- [14] R.A. Horn and C.R. Johnson. *Matrix Analysis*. Cambridge University Press, UK, 1985.
- [15] G. Huang, A.I. Mourikis, and S.I. Roumeliotis. Observability-based rules for designing consistent EKF-SLAM estimators. *Int. J. Robot. Res.*, 29(5):502–528, 2010.
- [16] T. Kavitha, C. Liebchen, K. Mehlhorn, D. Michail, R. Rizzi, T. Ueckerdt, and K. Zweig. Cycle bases in graphs: Characterization, algorithms, complexity, and applications. *Computer Science Rev.*, 3(4):199–243, 2009.
- [17] K. Konolige. Large-scale map-making. In *Proc. of the AAAI National Conf. on Artificial Intelligence*, 2004.
- [18] R. Kümmerle, B. Steder, C. Dornhege, M. Ruhnke, G. Grisetti, C. Stachniss, and A. Kleiner. Slam benchmarking webpage. <http://ais.informatik.uni-freiburg.de/slamevaluation>, 2009.
- [19] R. Kümmerle, B. Steder, C. Dornhege, M. Ruhnke, G. Grisetti, C. Stachniss, and A. Kleiner. On measuring the accuracy of SLAM algorithms. *Autonomous Robots*, 27(4):387–407, 2009.
- [20] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4: 333–349, 1997.
- [21] R. Martinez-Cantin, N. de Freitas, and J. Castellanos. Analysis of particle methods for simultaneous robot localization and mapping and a new algorithm: Marginal-SLAM. In *Proc. of the IEEE International Conf. on Robotics and Automation*, 2007.
- [22] J.M. Mendel. *Lessons in estimation theory for signal processing, communications, and control*. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [23] E. Olson, J.J. Leonard, and S. Teller. Fast iterative optimization of pose graphs with poor initial estimates. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 2262–2269, 2006.
- [24] L. Quan and T. Kanade. Affine structure from line correspondences with uncalibrated affine cameras. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 19: 834–845, 1997.
- [25] D. Sabatta, D. Scaramuzza, and R. Siegwart. Improved appearance-based matching in similar and dynamic environments using a vocabulary tree. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 2262–2269, 2010.
- [26] R. Smith and P. Cheesman. On the representation of spatial uncertainty. *Int. J. Robot. Res.*, 5(4):56–68, 1987.
- [27] S. Thrun and M. Montemerlo. The GraphSLAM algorithm with applications to large-scale mapping of urban structures. *Int. J. Robot. Res.*, 25:403–429, 2006.
- [28] S. Thrun, D. Koller, Z. Ghahramani, H. Durrant-Whyte, and A.Y. Ng. Simultaneous mapping and localization with sparse extended information filters. In *Proc. of the 5th Int. Workshop on Algorithmic Foundations of Robotics*, 2002.
- [29] S. Thrun, W. Burgard, and D. Fox. *Probabilistic robotics*. MIT press, 2005.