2010 IEEE International Conference on Robotics and Automation
Anchorage Convention District
May 3-8, 2010, Anchorage, Alaska, USA

On the Global Optimum of Planar, Range-based Robot-to-Robot Relative Pose Estimation

Nikolas Trawny and Stergios I. Roumeliotis
Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN 55455
E-mail: {trawny,stergios}@cs.umn.edu

***Abstract*— In this paper, we address the problem of determining the relative position and orientation (pose) of two robots navigating in 2D, based on known egomotion and noisy robot-to-robot distance measurements. We formulate this as a weighted Least Squares (WLS) estimation problem, and determine the exact global optimum by directly solving the multivariate polynomial system resulting from the first-order optimality conditions. Given the poor scalability of the original WLS problem, we propose an alternative formulation of the WLS problem in terms of squared distance measurements (squared distances WLS or SD-WLS). Using a hybrid algebraic-numeric technique, we are able to solve the corresponding first-order optimality conditions of the SD-WLS in 125 ms in Matlab. Both methods solve the minimal (3 distance measurements) as well as the overdetermined problem (more than 3 measurements) in a unified fashion. Simulation and experimental results show that the SD-WLS achieves performance virtually indistinguishable from the maximum likelihood estimator, and significantly outperforms current algebraic methods.**

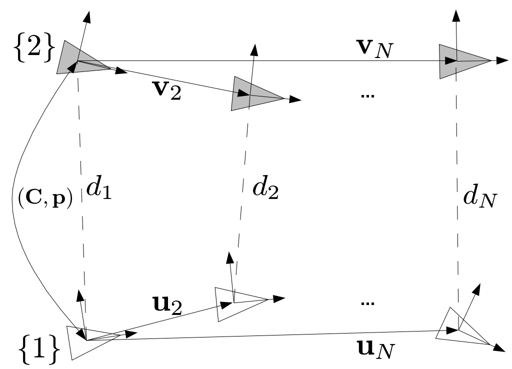

Fig. 1. Robot 1 (white) and 2 (gray) take N relative distance measurements, $d_1, \ldots, d_N$ as they move on a plane. The robots know their positions $\mathbf{u}_i$, $\mathbf{v}_i$, in their respective global frames, $\{1\}$ and $\{2\}$, e.g., from odometry. The objective is to determine the translation, $\mathbf{p}$, and rotation, $\mathbf{C}$, between frames $\{1\}$ and $\{2\}$.

I. Introduction and Related Work

Sensor fusion between multiple mobile robots, enabling tasks such as cooperative localization [1], mapping [2], or tracking [3], critically relies on accurate extrinsic calibration, i.e., the knowledge of the robots' relative position and orientation (pose). In the absence of a common, global frame of reference (e.g., provided by GPS), the robots can determine their relative pose based on robot-to-robot relative sensor measurements and odometry. Arguably the most challenging case of such motion-induced extrinsic calibration occurs when the robots can only measure distance between each other. Such distance measurements can be acquired by various sensors such as sonar, radar, or laser, or indirectly as a function of the received communication signal. Distance-based relative pose estimation in 2D is precisely the focus of this paper (see Fig. 1).

Previous research on leveraging sensor-to-sensor distance measurements to solve the relative localization problem has focussed primarily on *static* sensor networks. These approaches only determine the *position* of sensor nodes, not their relative orientation. Localization algorithms for static sensor networks infer the node positions using measurements to so-called anchor nodes with known global position. The position of the remaining sensors in the network can be uniquely determined if certain graph-rigidity constraints are satisfied [4]. A number of algorithms for 2D node localization have been proposed based on convex optimization [5], [6], multidimensional scaling [7], sum-of-squares relaxation [8], or graph connectivity [9]. There also exist methods for simultaneous localization and mapping of a *single* mobile robot using range measurements to static beacons [10].

For many practical applications, such as cooperative tracking [3] or sensor fusion [2], knowledge of both relative sensor position *and* orientation between multiple robots is required. Recent work by Zhou and Roumeliotis [11] has addressed the problem of relative pose estimation based on distance measurements for two mobile robots moving in a planar 2D environment. The authors propose computing the maximum likelihood estimate (MLE) of the relative pose in a two-step process: First, they derive methods to compute a coarse estimate for the relative pose, based on solving an overdetermined system of polynomial constraints using 4 or 5 distance measurements. In a second step, they use the result as an initial guess in an iterative weighted Least Squares (i-WLS) optimization. While such a two-step process is a standard way of solving nonlinear WLS problems, it is well known that the quality of an iterative WLS optimization depends critically on the accuracy of the initial guess.

Unfortunately, while exact in the noise-free case, in the presence of noise, the methods of [11] can produce an initial guess that is far from the MLE. Moreover, these methods are incapable of accounting for uncertainty in the measurements. As a result, the i-WLS refinement can converge to local minima, can take an excessive number of iterations to converge, or even diverge completely (see Sec. V).

Instead, in this paper we present a fundamentally different way to compute the MLE, by solving the nonlinear WLS directly for the guaranteed global optimum, in a non-iterative fashion. The approach is to first determine all stationary points by solving the first-order optimality conditions, and then to retain the one with minimum cost function value as the global minimum. The key difficulty – solving the first-order optimality conditions – can be overcome if the optimality conditions can be expressed as (or transformed into) a multivariate system of polynomials, which can be solved efficiently due to recent advances in polynomial system solving [12], [13], [14], [15].

The contributions of this paper are twofold: (i) We present the nonlinear WLS cost function for the relative pose problem, and a method to find its global minimum (Section III). Under the assumption of Gaussian noise, the optimal solution to the WLS will yield the MLE. We derive the corresponding first-order optimality conditions and show how to transform them into a system of polynomials, which can be solved for all stationary points using homotopy continuation. Our results show that this method is a feasible approach to find the MLE for a small number of measurements ($N < 5$). (ii) In order to efficiently address the case of more measurements, we present the WLS cost function and the corresponding first-order optimality conditions for the relative pose problem using squared distance measurements (SD-WLS, Section IV). The resulting polynomial system is shown to have constant solution complexity, independent of the number of measurements. We describe how to compute all stationary points efficiently using a hybrid algebraic-numeric technique based on the eigendecomposition of a generalized companion matrix [13].

Results from simulations and experiments described in Sections V and VI show that the SD-WLS estimator significantly outperforms the linear algorithm of [11], and yields performance almost indistinguishable from the original WLS estimator. While computing the global optimum of the WLS numerically takes several minutes even for as little as 4 measurements, our Matlab implementation can solve the SD-WLS in 125 ms, irrespective of the number of measurements.

II. Geometric Problem Formulation

As illustrated in Fig. 1, consider two robots moving in 2D and acquiring robot-to-robot distance measurements $d_i,\ i = 1, \dots, N$. We assume without loss of generality that the global frames of each robot, $\{1\}$ and $\{2\}$, are attached to the points where the first mutual measurement occurs. Further, we assume that each robot can localize with respect to its own global frame of reference, for example, using wheel odometry. Therefore, we assume that at the time the robots acquire the distance measurements, the coordinates of the first robot, $\mathbf{u}_i := {}^1\mathbf{u}_i$, and the second robot, $\mathbf{v}_i := {}^2\mathbf{v}_i$, are known.[1] The objective is to find the 3 degree-of-freedom transformation between frames $\{1\}$ and $\{2\}$, i.e., the translation $\mathbf{p} := {}^1\mathbf{p}_2 = \begin{bmatrix} x & y \end{bmatrix}^T$ and rotation $\mathbf{C} := {}^1_2\mathbf{C}(\phi) = \begin{bmatrix} \cos(\phi) & -\sin(\phi) \\ \sin(\phi) & \cos(\phi) \end{bmatrix}$.

The robot-to-robot distance d_i can be expressed as the length of vector $\mathbf{w}_i,\ i = 1, \dots, N$, connecting the two robots at the time of measurement (see Fig. 1)

$$d_i = ||\mathbf{w}_i||_2 = \sqrt{\mathbf{w}_i^T \mathbf{w}_i}\,, \qquad \mathbf{w}_i := \mathbf{p} + \mathbf{C}\mathbf{v}_i - \mathbf{u}_i \quad (1)$$

The approach of [11] is to assume noise-free measurements, and to stack the constraints provided by each measurement into the following system of polynomial equations in the four unknowns x, y, $s\phi$, and $c\phi$:

Deterministic System:

$$\mathbf{w}_i^T \mathbf{w}_i - d_i^2 = 0, \quad i = 1, \dots, N \qquad (2a)$$

$$s^2\phi + c^2\phi - 1 = 0 \qquad (2b)$$

From [11], it is known that $N = 3$ measurements are necessary for this problem to have a discrete set of 6 possibly complex solutions (minimal problem). Moreover, $N > 3$ measurements uniquely determine the relative pose.

In practice, however, the robot-to-robot distances d_i will have to be replaced by noisy measurements

$$z_i = d_i + n_i, \quad i = 1, \dots, N \qquad (3)$$

$$\mathbf{n} = \begin{bmatrix} n_1 & \dots & n_N \end{bmatrix}^T \sim \mathcal{N}(\mathbf{n}; \mathbf{0}, \mathbf{R}) \qquad (4)$$

where we have modeled the measurement noise $\mathbf{n}$ as zero-mean Gaussian, with covariance matrix $\mathbf{R}$. In case of independent measurements, $\mathbf{R} = \mathrm{diag}(\sigma_1^2, \dots, \sigma_{\mathrm{N}}^2)$, but in the presence of correlated noise, $\mathbf{R}$ can be full. Replacing d_i by the noisy measurements z_i, the overdetermined system (2) with $N > 3$ will be inconsistent, and no guarantees can be given for the solutions computed by the polynomial 4-point or the linear 5-point algorithm in [11]. In particular, these algorithms will not be able to account for the measurement noise in a statistically sound fashion. Therefore, they can only be used to provide an initial guess for a subsequent i-WLS refinement.

As an alternative, we propose to combine the constraints (2a) arising from the measurements in a WLS fashion, thus correctly accounting for the measurement noise. In what follows, we present two methods to solve the WLS problem directly for the *globally* optimal estimate of the relative pose. The approach is to formulate the first-order optimality conditions of the WLS cost function as a square system of polynomials, which can be solved directly to obtain all stationary points. A desirable property of this formulation is that without any modification it will yield the global optimum for the overdetermined system ($N > 3$), as well as all 6 solutions to the minimal problem ($N = 3$), which will be global minima of the WLS cost function with cost function value identically equal to zero.

III. Weighted Least-Squares (WLS) Formulation and Solution

Ideally, given (1) and (3), we would like to solve the following WLS problem

[1] Throughout the paper the notation ${}^i\mathbf{p}$ denotes a vector expressed with respect to frame $\{i\}$, and ${}^i_j\mathbf{C}$ denotes the rotation matrix that transforms vectors from frame $\{j\}$ to frame $\{i\}$. Also, in polynomial formulations we substitute $\sin(\phi)$ and $\cos(\phi)$ by the variables $s\phi$ and $c\phi$.

WLS Cost Function:

$$\min_{x,y,\phi} \frac{1}{2}\, \mathbf{e}_d^T \mathbf{R}^{-1} \mathbf{e}_d \tag{5}$$

where

$$\mathbf{e}_d := \begin{bmatrix}\sqrt{\mathbf{w}_1^T\mathbf{w}_1} - z_1 & \sqrt{\mathbf{w}_2^T\mathbf{w}_2} - z_2 & \dots & \sqrt{\mathbf{w}_N^T\mathbf{w}_N} - z_N\end{bmatrix}^T$$

Taking the gradient with respect to the unknowns and setting it to zero, we obtain the following first-order optimality conditions

$$\nabla_{\mathbf{p}} = \mathbf{e}_d^T \mathbf{R}^{-1} \begin{bmatrix}\frac{\mathbf{w}_1}{\sqrt{\mathbf{w}_1^T\mathbf{w}_1}} & \frac{\mathbf{w}_2}{\sqrt{\mathbf{w}_2^T\mathbf{w}_2}} & \cdots & \frac{\mathbf{w}_N}{\sqrt{\mathbf{w}_N^T\mathbf{w}_N}}\end{bmatrix}^T = \mathbf{0} \tag{6a}$$

$$\nabla_{\phi} = \mathbf{e}_d^T \mathbf{R}^{-1} \begin{bmatrix}\frac{\mathbf{w}_1^T\mathbf{J}\mathbf{C}\mathbf{v}_1}{\sqrt{\mathbf{w}_1^T\mathbf{w}_1}} & \frac{\mathbf{w}_2^T\mathbf{J}\mathbf{C}\mathbf{v}_2}{\sqrt{\mathbf{w}_2^T\mathbf{w}_2}} & \cdots & \frac{\mathbf{w}_N^T\mathbf{J}\mathbf{C}\mathbf{v}_N}{\sqrt{\mathbf{w}_N^T\mathbf{w}_N}}\end{bmatrix}^T = 0 \tag{6b}$$

where $\mathbf{J} = \left[\begin{smallmatrix}0 & -1\\ 1 & 0\end{smallmatrix}\right]$.

In order to solve this system of equations and thus to determine all stationary points, we transform it into a system of polynomials by introducing auxiliary variables. Our objective is then to apply efficient polynomial system solving techniques that can find all roots simultaneously. One possible choice is to introduce the auxiliary variables

$$a_i = \sqrt{\mathbf{w}_i^T\mathbf{w}_i}, \quad i = 1, \dots, N \tag{7}$$

Further, instead of optimizing over ϕ, we optimize over the two variables $s\phi$ and $c\phi$, related by the trigonometric constraint

$$s^2\phi + c^2\phi = 1 \tag{8}$$

We then rewrite (5) as a constrained optimization problem, with Lagrangian

$$L = \frac{1}{2}\begin{bmatrix}a_1 - z_1\\ a_2 - z_2\\ \vdots\\ a_N - z_N\end{bmatrix}^T \mathbf{R}^{-1} \begin{bmatrix}a_1 - z_1\\ a_2 - z_2\\ \vdots\\ a_N - z_N\end{bmatrix} + \sum_{i=1}^{N} \lambda_i/2(\mathbf{w}_i^T\mathbf{w}_i - a_i^2) + \mu(s^2\phi + c^2\phi - 1) \tag{9}$$

Taking the gradient with respect to the unknowns and eliminating μ, we obtain the following system of $2N+4$ polynomials in the $2N+4$ unknowns $x, y, s\phi, c\phi$, and $\mathbf{a} := \begin{bmatrix}a_1 & \dots & a_N\end{bmatrix}^T$, $\boldsymbol{\lambda} := \begin{bmatrix}\lambda_1 & \dots & \lambda_N\end{bmatrix}^T$.

WLS First-Order Optimality Conditions:

$$\nabla_{\mathbf{p}} = \sum_{i=1}^{N} \lambda_i \mathbf{w}_i = \mathbf{0} \tag{10a}$$

$$c\phi \cdot \nabla_{s\phi} L - s\phi \cdot \nabla_{c\phi} L = \sum_{i=1}^{N} \lambda_i \mathbf{w}_i^T \mathbf{J}\mathbf{C}\mathbf{v}_i = 0 \tag{10b}$$

$$\nabla_{\mathbf{a}} L = \mathbf{R}^{-1}\begin{bmatrix}a_1 - z_1\\ \vdots\\ a_N - z_N\end{bmatrix} - \begin{bmatrix}a_1\lambda_1\\ \vdots\\ a_N\lambda_N\end{bmatrix} = \mathbf{0} \tag{10c}$$

$$2\nabla_{\lambda_i} L = \mathbf{w}_i^T\mathbf{w}_i - a_i^2 = 0, \quad i = 1, \dots, N \tag{10d}$$

$$\nabla_{\mu} L = s^2\phi + c^2\phi - 1 = 0 \tag{10e}$$

All polynomials are quadratic, except (10b) which is cubic.

Notice that squaring the constraints on the a_i (cf. (7) vs. (10d)) introduces spurious solutions with erroneous signs. Therefore, we will only accept real solutions that fulfill the original constraint (7) as candidate stationary points.

Once all the stationary points of (5) are found by directly solving system (10), the final step is to evaluate the cost function on the candidate points, and to choose the one with minimum value as the final, guaranteed globally optimal WLS estimate of the relative pose. The numerical solution of (10) can be obtained, for example, using homotopy continuation [16]. System (10) has 70 solutions for $N = 3$ measurements, 240 solutions for $N = 4$, and 784 solutions for $N = 5$, as determined both through analysis of the Gröbner basis using Macaulay 2 [17] and numerically using PHCPack [16]. In our experiments, we solved example systems with $N = 4$ measurements in about 2 to 4 minutes (diagonal vs. dense covariance matrix), and $N = 5$ measurements in 12 to 20 minutes, using PHCPack on an Intel T9400 2.53 GHz laptop with 2GB of RAM. The reason for the increase in time is the introduction of additional monomials in the polynomial system due to cross-coupling induced by the off-diagonals of $\mathbf{R}$.

Unfortunately, as evident from these results, this formulation scales quite poorly, because the number of unknowns (i.e., a_i, λ_i) grows with the number of measurements. The complexity of solving polynomial systems is exponential in the number of unknowns [18], rendering this approach feasible only for small-scale problems. This leads us to the introduction of an alternative formulation using squared distance measurements.

IV. Squared Distances WLS (SD-WLS) Formulation and Solution

A much preferable approach from the point of scalability is to solve the WLS based on noisy measurements of the squared distance

$$z_i' = d_i^2 + n_i', \quad i = 1, \dots, N \tag{11}$$

$$\mathbf{n}' = \begin{bmatrix}n_1' & \dots & n_N'\end{bmatrix} \sim \mathcal{N}(\mathbf{n}'; \mathbf{0}, \mathbf{R}') \tag{12}$$

where for simplicity we temporarily assume that the noise in the squared measurements is zero-mean Gaussian. Under this assumption, the corresponding WLS will actually yield the MLE. We will relax this assumption in Section IV-A.

The idea of applying the LS methodology to squared distance measurements has previously been used in source localization, e.g., [6], [8], where the authors solve an unweighted LS approximately using convex or sum-of-squares relaxation.

The cost function of the WLS problem using squared distances (SD-WLS) is defined as

SD-WLS Cost Function:

$$\min_{x,y,\phi} \frac{1}{2}\, \mathbf{e}_{sd}^T \mathbf{R}'^{-1} \mathbf{e}_{sd} \tag{13}$$

where

$$\mathbf{e}_{sd} := \begin{bmatrix}\mathbf{w}_1^T\mathbf{w}_1 - z_1' & \mathbf{w}_2^T\mathbf{w}_2 - z_2' & \dots & \mathbf{w}_N^T\mathbf{w}_N - z_N'\end{bmatrix}^T$$

Setting the gradient to zero and dividing by constant factors, we obtain the following polynomial system of 4 equations in the 4 unknowns $x, y, s\phi, c\phi$.

SD-WLS First-Order Optimality Conditions:

$$\mathbf{e}_{sd}^T \mathbf{R}'^{-1} \begin{bmatrix} \mathbf{w}_1 & \mathbf{w}_2 & \dots & \mathbf{w}_N \end{bmatrix}^T = \mathbf{0} \tag{14a}$$

$$\mathbf{e}_{sd}^T \mathbf{R}'^{-1} \begin{bmatrix} \mathbf{w}_1^T \mathbf{J}\mathbf{C}\mathbf{v}_1 & \mathbf{w}_2^T \mathbf{J}\mathbf{C}\mathbf{v}_2 & \dots & \mathbf{w}_N^T \mathbf{J}\mathbf{C}\mathbf{v}_N \end{bmatrix}^T = 0 \tag{14b}$$

$$s^2\phi + c^2\phi - 1 = 0 \tag{14c}$$

The two polynomials (14a) have total degree 3, (14b) has total degree 4, and the trigonometric constraint (14c) is quadratic.

Analysis of this system's Gröbner basis for different numerical instances using coefficients from a finite field shows that this system has 28 possibly complex solutions. This is a general result that holds for all cases of N and $\mathbf{R}'$. In particular, we prove the following proposition:

Proposition 1: The number of solutions of the system (14), and its solution complexity, is independent of the number of measurements, N, for $N \geq 3$, and of the covariance matrix, $\mathbf{R}'$.

Proof: We base our proof on the following observation: If for a system of polynomials arising from a specific problem class, different (non-singular) instantiations of the problem differ only in the numeric coefficients but not in the structure of the polynomial system (i.e., the monomials comprising each polynomial remain the same), the leading monomials of the corresponding Gröbner basis will generally also be the same, with the consequence that the standard basis of the quotient ring and the number of solutions of the system is constant [19], [14]. We therefore have to show that the structure of system (14) and the monomials in each polynomial (14a)-(14c) are independent of N and $\mathbf{R}'$ generically. Clearly (but in contrast to the polynomial system (10)), the number of equations and the number of unknowns of (14) is constant and independent of N and $\mathbf{R}'$. Also, (14c) remains unchanged for different problem instantiations. To see that the monomials in (14a) and (14b) are independent of N and $\mathbf{R}'$, consider rewriting these equations as $\sum_{i,j=1}^N r_{i,j}'^{-1}(z_i' - \mathbf{w}_i^T\mathbf{w}_i)\mathbf{w}_j^T$ and $\sum_{i,j=1}^N r_{i,j}'^{-1}(z_i' - \mathbf{w}_i^T\mathbf{w}_i)(\mathbf{w}_j^T\mathbf{J}\mathbf{C}\mathbf{v}_j)$, where $r_{i,j}'^{-1}$ denotes the (i,j)-th entry of $\mathbf{R}'^{-1}$. Generically, the polynomials $z_i' - \mathbf{w}_i^T\mathbf{w}_i$, $\mathbf{w}_i$, and $\mathbf{w}_i^T\mathbf{J}\mathbf{C}\mathbf{v}_i$ each contain the same monomials regardless of i. Analogously, the products $(z_i' - \mathbf{w}_i^T\mathbf{w}_i)\mathbf{w}_j^T$ and $(z_i' - \mathbf{w}_i^T\mathbf{w}_i)(\mathbf{w}_j^T\mathbf{J}\mathbf{C}\mathbf{v}_j)$ each contain the same monomials generically, regardless of i and j. Finally, linear combinations of polynomials do not introduce new monomials, which concludes the proof. ■

Given the constant structure of the polynomial system (14), we can avoid using homotopy continuation methods, such as PHCPack, and instead apply a hybrid algebraic-numeric method based on the eigendecomposition of a generalized companion matrix [13], [15], [12]. In particular, we employed the normal set-based method, described in detail in [13], which exploits the specific system structure to provide all 28 solutions simultaneously, in a fast, non-iterative fashion.

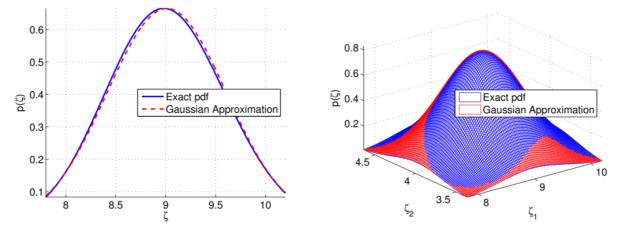

Fig. 2. Examples of approximating the pdf of squared distance measurements by a Gaussian. (a) Scalar measurement case, with $d_1 = 3$m and $\sigma_1 = 0.1$m. (b) Two distance measurements with $\mathbf{d} = \begin{bmatrix} 3\text{m} & 2\text{m} \end{bmatrix}^T$, and correlated noise with covariance $\mathbf{R} = \begin{bmatrix} 0.1^2 & 0.01^2 \\ 0.01^2 & 0.08^2 \end{bmatrix}$. For small variances, the approximation is reasonably accurate: for the examples shown here, the KL divergence is $8.34 \cdot 10^{-4}$ and $2.0 \cdot 10^{-3}$, respectively.

The idea behind the normal set-based method is to first expand the polynomial system (14), compactly written as

$$\mathbf{\Psi}(x, y, s\phi, c\phi) = \mathbf{0} \tag{15}$$

by adding new polynomials of the form

$$\psi_i' = \psi_i(x, y, s\phi, c\phi) \cdot x^{\alpha_x} y^{\alpha_y} (s\phi)^{\alpha_s} (c\phi)^{\alpha_c} = 0 \tag{16}$$

The new polynomials, ψ_i', are products of the original polynomials, ψ_i, with monomials in the unknowns raised to some power α. For the SD-WLS problem, we chose α so as to create new polynomials up to total degree 8. As an example, we created 70 polynomials by multiplying (14b) with all monomials of total degree 0 to 4. Notice that adding these new polynomials does not change the solution set, since all ψ_i' vanish on the roots of the original system, and do not add new solutions. The next step is to write the expanded polynomial system in matrix form

$$\mathbf{C}_e \mathbf{x} = \mathbf{0}, \tag{17}$$

gathering the monomials in the unknowns in the vector $\mathbf{x}$, and the numeric coefficients in the expanded coefficient matrix $\mathbf{C}_e$. From this matrix one can extract a 28×28 generalized companion matrix that defines multiplication by a function (which we chose to be x) within the so-called quotient ring. An eigendecomposition of this matrix yields all 28 solutions simultaneously [15], [13].

Our current Matlab implementation requires approximately 125 ms (as determined by Matlab's profiler) to solve an instance of the problem, of which approximately 16 ms are spent creating the expanded coefficient matrix $\mathbf{C}_e$, which for this problem is of dimension 532×495, with 4% non-zero entries. Extracting the multiplication matrix and all 28 solutions requires 109 ms.

A. *Gaussian Approximation of Squared Noisy Measurements*

Notice that if only noisy measurements of the (non-squared) distance are available, replacing (3) by (11) cannot be achieved by simply squaring the noisy measurements, i.e., $z_i' \neq z_i^2$. The reason is that $z_i^2 = d_i^2 + 2d_i n_i + n_i^2$, and the corresponding noise term $s_i := 2d_i n_i + n_i^2$ is *not* zero-mean Gaussian. Indeed, following the standard formulas to compute the pdf of functions of random variables [20], the pdf of the vector $\boldsymbol{\zeta} := \begin{bmatrix} z_1^2 & \dots & z_N^2 \end{bmatrix}$ is given by

$$p(\boldsymbol{\zeta}) = \sum_{j=1}^{2^N} \frac{1}{2^N \prod_i \sqrt{\zeta_i}} \mathcal{N}(\boldsymbol{\gamma}_j; \mathbf{d}, \mathbf{R}) \tag{18}$$

where each $\boldsymbol{\gamma}_j,\ j = 1, \dots, 2^N$, is a vector of the form $\begin{bmatrix} \pm\sqrt{\zeta_1} & \dots & \pm\sqrt{\zeta_N} \end{bmatrix}^T$ with one of the 2^N possible different sign assignments for its individual elements.

However, the non-Gaussian pdf resulting from squaring a Gaussian random variable can be well approximated by a Gaussian pdf with matching first and second order moments. Specifically, computing the mean, $\bar{s}_i$, and covariance, $\boldsymbol{\Sigma}$, of s_i yields

$$\bar{s}_i := E[s_i] = \mathbf{R}_{ii} \tag{19}$$

$$\boldsymbol{\Sigma}_{ii} := E[(s_i - \bar{s}_i)^2] = \mathbf{R}_{ii}(4d_i^2 + 2\mathbf{R}_{ii}) \tag{20}$$

$$\boldsymbol{\Sigma}_{ij} := E[(s_i - \bar{s}_i)(s_j - \bar{s}_j)] = \mathbf{R}_{ij}(4d_i d_j + 2\mathbf{R}_{ij})$$

We can now approximate (11) by setting

$$z_i' \simeq z_i^2 - \bar{s}_i \tag{21}$$

$$\mathbf{R}' \simeq \boldsymbol{\Sigma} \tag{22}$$

where we replace the distances d_i in the expression for $\boldsymbol{\Sigma}$ by their noisy measurements z_i.

As illustrated in Fig. 2, using (21), (22) to approximate the non-Gaussian pdf (18) by a Gaussian is reasonably accurate, particularly for high signal-to-noise ratios.

V. Simulation Results

We compared the performance of the WLS, the SD-WLS, and the linear algorithm of [11] for the overdetermined case where $N > 3$ noisy (but outlier-free) measurements are available.[2] Specifically, we conducted Monte Carlo simulations using $N = 5$ measurements, measurement noise with diagonal covariance matrix $\mathbf{R} = \sigma_d^2 \mathbf{I}$ for different values of σ_d, and 1000 trials per setting. The trajectories were chosen so that the robot-to-robot distances varied randomly between 1-2 m, and the displacement between measurements varied between 3-6 m. The parameters z_i' and $\mathbf{R}'$ were determined based on the approximations (21) and (22).

The results are shown in Fig. 3. As evident, SD-WLS consistently and significantly outperforms the linear algorithm. From the doubly logarithmic plot one can deduce that the median position error of the linear algorithm in [11] with respect to the true solution is approximately twice as large as the WLS and the SD-WLS error, and the orientation error approximately 3.5 times as large, independent of the measurement noise standard deviation. Even more remarkable is the fact that the performance of the SD-WLS is virtually indistinguishable from that of the true WLS, despite it not being optimal in the maximum likelihood sense (due to the approximation of Section IV-A). Only for large noise levels does the performance start to degrade slightly.

[2]If outliers are present, one should first perform RANSAC with the solutions of the minimal problem for outlier rejection.

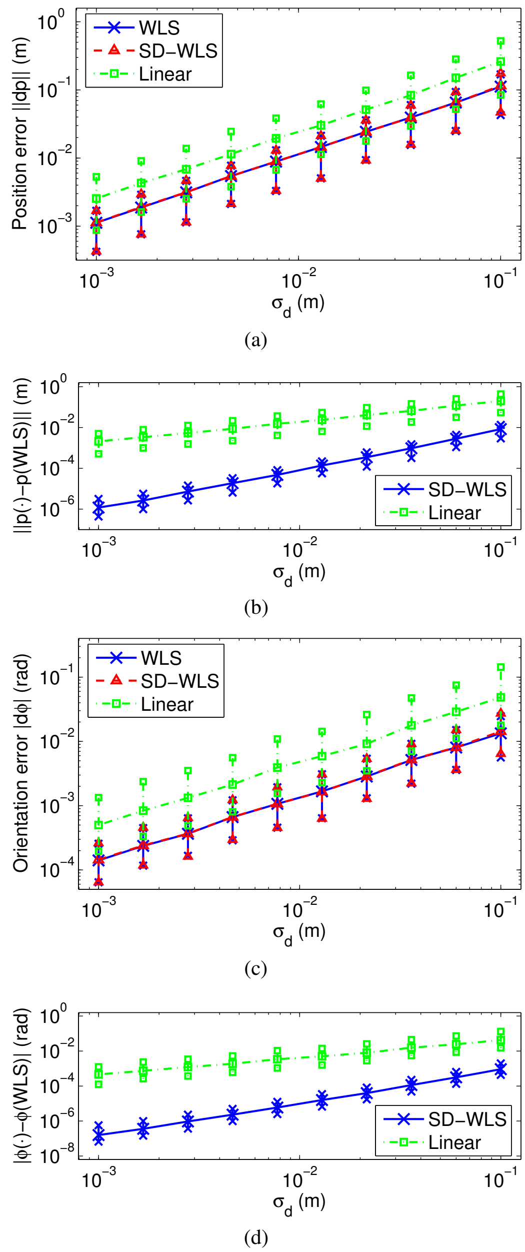

Fig. 3. Simulation results from 1000 trials for each noise setting, comparing WLS, SD-WLS and the linear algorithm of [11], using 5 distance measurements corrupted by zero-mean Gaussian measurement noise with covariance $\mathbf{R} = \sigma_d^2 \mathbf{I}$. Plotted are the median and the 25% and 75% quantiles. The accuracy of the WLS and the SD-WLS formulation is almost identical, and higher than that of the linear method by a constant factor. The SD-WLS solution is close to that of WLS particularly for small noise values. (a) Norm of position error with respect to ground truth. (b) Difference in position with respect to WLS. (c) Orientation error with respect to ground truth. (d) Difference in orientation with respect to WLS.

If despite the demonstrated performance of SD-WLS the actual MLE were required, one could use the result from SD-WLS as an extremely accurate initial guess for i-WLS refinement. We tested the performance of initializing Gauss-Newton-based i-WLS with the linear solution, the SD-WLS solution, and with ground-truth. The results shown in Fig. 4 confirm that SD-WLS-based initialization outperforms linear initialization in terms of required iterations, instances of divergence, and number of inconsistent estimates. Particularly

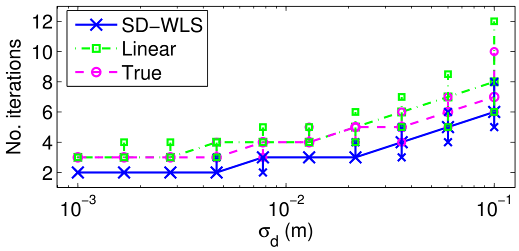

(a)

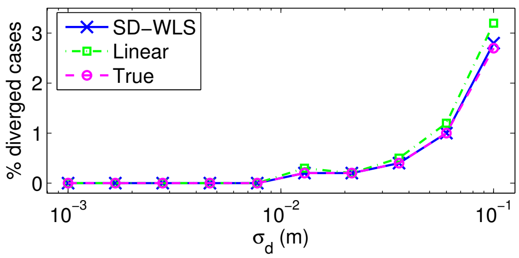

(b)

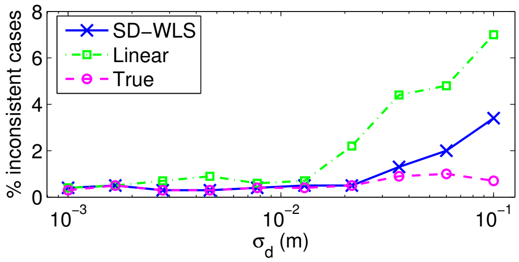

(c)

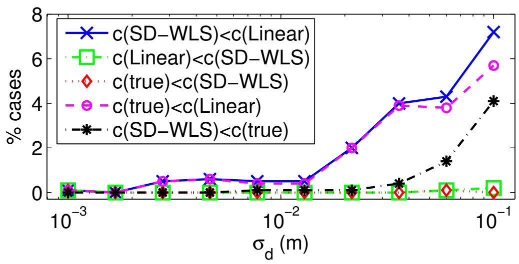

(d)

Fig. 4. Comparison of i-WLS initialized with the result of the linear algorithm of [11], that of SD-WLS, and with ground truth. (a) Number of iterations to convergence (median and 25% / 75% quantiles). (b) Percentage of diverged cases ($>$ 1000 iterations). (c) Percentage of inconsistent cases (error $> 3\sigma$). (d) Percentage of cases where i-WLS converges to an optimum with lower WLS cost function value, depending on initialization. The notation c(SD-WLS) refers to the WLS cost function after i-WLS convergence when initialized with the method inside the parentheses, here the SD-WLS algorithm.

interesting is the fact that initializing i-WLS with the less accurate solution of the linear method can indeed lead to convergence to a local minimum, as illustrated in Fig. 4(d). In contrast, this happens only extremely rarely when initializing using the SD-WLS solution. In fact, in the presence of large noise, even the ground truth might not reside within the basin of attraction of the global WLS optimum, and i-WLS initialized with the SD-WLS solution can achieve a lower cost function value than i-WLS initialized with ground truth.

In conclusion, the simulation results confirm the superior accuracy of SD-WLS over the linear method. However, particularly for low measurement noise, the linear method yields surprisingly accurate results that may very well be acceptable for some applications, especially given its considerably lower computational complexity.[3]

VI. Experimental Results

We have also tested the algorithms experimentally, using two Pioneer-II robots moving randomly in an area of 4 m $\times$ 5 m (see Fig. 5(a)). Each robot estimated its position from 10 Hz wheel odometry, using a differential drive kinematic model with noisy wheel velocity measurements having standard deviation of $\sigma_v = 8$ mm/s. The ground truth was established from observations using a calibrated ceiling-mounted camera. These data also provided synthetic relative distance measurements by adding white, zero-mean Gaussian noise with standard deviation $\sigma_d = 0.05$ m.

We compared the estimates of the initial relative pose obtained using the SD-WLS, the linear algorithm with 5 measurements, and the i-WLS initialized with SD-WLS. For the SD-WLS solution we processed two iterations. The first accounted only for noise in the distance measurements, and its optimum was used as linearization point for the Jacobians with respect to the robot positions, in order to correctly account for uncertainty due to odometry in the second iteration. We see that SD-WLS is significantly more accurate than the linear method, almost independently of the number of measurements, since the growing position uncertainty due to integrated odometry errors cancels gains from processing additional measurements. The errors after i-WLS refinement are in the order of 10 cm for position and 3 degrees for orientation.

VII. Conclusions

In this paper, we have presented two methods to estimate the relative pose between two robots based on robot-to-robot distance measurements and knowledge of the robots' egomotion. In particular, we have solved for the global optimum of the corresponding WLS problem by finding all stationary points as the roots of a square multivariate polynomial system. We have shown how to construct this polynomial system for the WLS of the original problem, which is also the MLE for distance measurements corrupted by Gaussian noise. The complexity of this system was shown to scale very poorly, since the number of variables grows with the number of measurements. On the other hand, the alternative SD-WLS formulation using squared distance measurements was shown to have constant solution complexity, independent of the number of measurements or the structure of the covariance matrix. Solving the problem is carried out very efficiently using recent hybrid algebraic-numeric techniques to solve multivariate polynomial systems based on the eigendecomposition of a generalized companion matrix [13]. Our algorithm can find the six solutions of the minimal problem ($N = 3$ measurements) or the unique

[3]In our Matlab implementations, the linear method takes approx. 0.11 ms, and a single iteration of i-WLS requires approx. 0.15 ms.

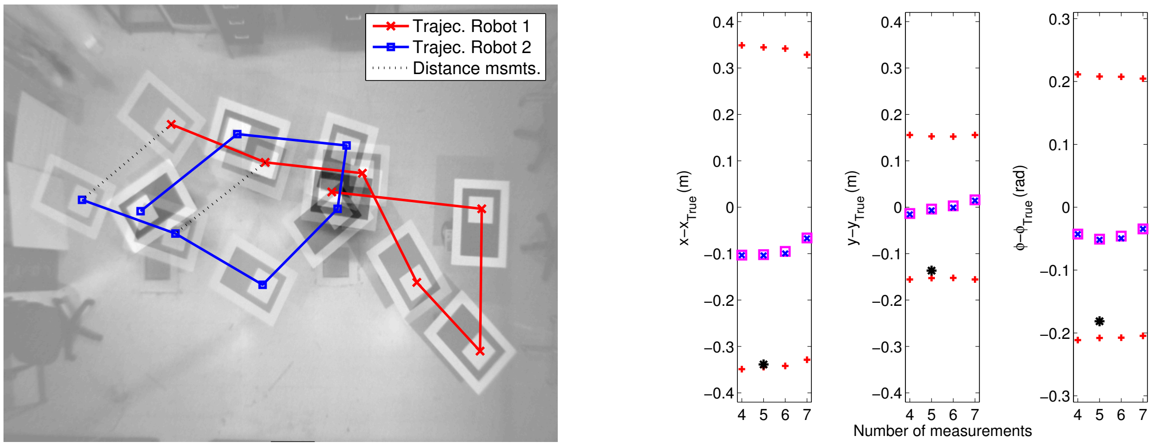

(a) (b)

Fig. 5. Experiment with two Pioneer robots equipped with markers for optical tracking using an overhead camera. (a) The camera images provided ground truth for relative pose, and served to create 7 synthetic relative distance measurements at the marked points (for clarity, only the first two measurements are drawn as dotted lines). (b) Error in relative pose as a function of the number of measurements. Shown are i-WLS refinement of SD-WLS (blue x) and corresponding 3σ-bounds (red +), as well as the output of pure SD-WLS (magenta square) and the linear algorithm (black star, for 5 measurements only).

solution in the overdetermined case ($N > 3$ measurements) in a unified framework.

Despite losing optimality in the MLE sense, the squared WLS method performed almost indistinguishably from the WLS estimate in simulation, and was shown to be approximately two (for position) or three (for orientation) times more accurate than the linear method presented in [11] when using 5 measurements, over a wide range of noise levels.

In a broader context, this paper demonstrates that the newly available tools for polynomial system solving allow moving away from iterative optimization schemes towards global optimization, if the first-order optimality conditions can be transformed into polynomial form. Moreover, we have provided an example that for overdetermined systems, combining polynomial constraints in a Least Squares formulation is far more accurate than solving the system of polynomial constraints deterministically.

Acknowledgements

This work was supported by the University of Minnesota (DTC), and the National Science Foundation (IIS-0643680, IIS-0811946, IIS-0835637).

References

[1] I. M. Rekleitis, G. Dudek, and E. E. Milios, "Multi-robot cooperative localization: a study of trade-offs between efficiency and accuracy," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Lausanne, Switzerland, Oct. 2002, pp. 2690–2695.

[2] X. S. Zhou and S. I. Roumeliotis, "Multi-robot SLAM with unknown initial correspondence: The robot rendezvous case," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Beijing, China, Oct. 2006, pp. 1785–1792.

[3] K. Zhou and S. I. Roumeliotis, "Optimal motion strategies for range-only constrained multi-sensor target tracking," *IEEE Transactions on Robotics and Automation*, vol. 24, no. 5, pp. 1168–1185, Oct. 2008.

[4] J. Aspnes, T. Eren, D. K. Goldenberg, A. S. Morse, W. Whiteley, Y. R. Yang, B. D. O. Anderson, and P. N. Belhumeur, "A theory of network localization," *IEEE Transactions on Mobile Computing*, vol. 5, no. 12, pp. 1663–1678, Dec. 2006.

[5] L. Doherty, K. S. J. Pister, and L. E. Ghaoui, "Convex position estimation in wireless sensor networks," in *Proc. of the 20th Annual Joint Conf. of the IEEE Computer and Communications Societies (IEEE INFOCOM)*, Anchorage, AK, Apr. 2001, pp. 1655–1663.

[6] A. Beck, P. Stoica, and J. Li, "Exact and approximate solutions of source localization problems," *IEEE Transactions on Signal Processing*, vol. 56, no. 5, pp. 1770–1778, May 2008.

[7] J. A. Costa, N. Patwari, and I. Alfred O. Hero, "Distributed weighted-multidimensional scaling for node localization in sensor networks," *ACM Transactions on Sensor Networks*, vol. 2, no. 1, pp. 39–64, Feb. 2006.

[8] I. Shames, P. Bibalan, B. Fidan, and B. Anderson, "Polynomial methods in noisy network localization," in *Proc. of the 17th Mediterranean Conference on Control and Automation*, Thessaloniki, Greece, Jun. 24–26, 2009, pp. 1307–1312.

[9] Y. Shang, W. Ruml, Y. Zhang, and M. P. J. Fromherz, "Localization from mere connectivity," in *Proc. of the 4th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, Annapolis, MD, Jun. 2003, pp. 201–212.

[10] J. Djugash and S. Singh, "A robust method of localization and mapping using only range," in *International Symposium on Experimental Robotics*, Athens, Greece, Jul. 14–17, 2008.

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

[12] M. Byröd, Z. Kukelova, K. Josephson, T. Pajdla, and K. Åström, "Fast and robust numerical solutions to minimal problems for cameras with radial distortion," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Anchorage, AK, Jun. 23–28, 2008, pp. 1–8.

[13] N. Trawny, X. S. Zhou, and S. I. Roumeliotis, "3D relative pose estimation from six distances," in *Proceedings of Robotics: Science and Systems*, Seattle, WA, Jun. 28– Jul. 1, 2009.

[14] Z. Kukelova and T. Pajdla, "A minimal solution to the autocalibration of radial distortion," in *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, Minneapolis, MN, Jun. 18–23, 2007, pp. 1–7.

[15] D. Cox, J. Little, and D. O'Shea, *Ideals, Varieties, and Algorithms*, 3rd ed. Springer, 2008.

[16] J. Verschelde, "Algorithm 795: PHCpack: A general-purpose solver for polynomial systems by homotopy continuation," *ACM Transactions on Mathematical Software*, vol. 25, no. 2, pp. 251–276, 1999.

[17] D. R. Grayson and M. E. Stillman, "Macaulay 2, a software system for research in algebraic geometry," Available at http://www.math.uiuc.edu/Macaulay2/.

[18] I. Z. Emiris and A. Rege, "Monomial bases and polynomial system solving," in *Proc. of the ACM International Symposium on Symbolic and Algebraic Computation*, Oxford, UK, Jul. 20–22, 1994, pp. 114–122.

[19] D. Cox, J. Little, and D. O'Shea, *Using Algebraic Geometry*, 2nd ed. Springer, 2005.

[20] H. Stark and J. W. Woods, *Probability and Random Processes with Applications to Signal Processing*. Prentice Hall, 2001.