

Recursive Decentralized Collaborative Localization for Sparsely Communicating Robots

Lukas Luft*, Tobias Schubert*, Stergios I. Roumeliotis[†], and Wolfram Burgard*

*University of Freiburg, Department of Computer Science, Freiburg, Germany

[†]University of Minnesota, Department of Computer Science and Engineering, Minneapolis, USA

Abstract—This paper provides a new fully-decentralized algorithm for *Collaborative Localization* based on the extended Kalman filter. The major challenge in decentralized collaborative localization is to track inter-robot dependencies – which is particularly difficult in situations where sustained synchronous communication between robots cannot be guaranteed. Current approaches suffer from the need for particular communication schemes, extensive bookkeeping of measurements, overly-conservative assumptions, or the restriction to specific measurement models. To the best of our knowledge, the algorithm we propose in this paper is the first one that tracks inter-robot correlations while fulfilling all of the following relevant conditions: communication is limited to two robots that obtain a relative measurement, the algorithm is recursive in the sense that it does not require storage of measurements and each robot maintains only the latest estimate of its own pose, and it supports generic measurement models. These particularly hard conditions make the approach applicable to a wide range of multi-robot applications. Extensive experiments carried out using real world datasets demonstrate the improved performance of our method compared to several existing approaches.

I. INTRODUCTION

Localization is one of the most fundamental tasks for mobile robots. For a team of multiple robots, a paradigm called *Collaborative Localization* (CL) has been demonstrated to provide significant improvements of the localization performance of the individual teammates [5, 22]. In CL, robots detect each other and communicate their estimates, which correlates the estimates of their individual poses. It is of fundamental importance to keep track of these dependencies to avoid the problem of *double-counting* or *data-incest*, which occurs when two robots treat shared information as uncorrelated. Leading to overly optimistic estimates, double-counting puts the localization reliability at risk.

To fuse information gathered from proprioceptive (e.g., wheel encoders) and exteroceptive sensors (e.g., cameras or laser scanners) researchers have used Particle Filters [5, 21, 7], Maximum-a-Posteriori [19] and Maximum-Likelihood algorithms [6], and – like in our approach – the *Extended Kalman Filter* (EKF). Under the assumption of Gaussian noise for motions and measurements, applying the EKF to the joint system state space containing the poses of all team members yields the optimal estimator apart from linearization errors [24, 8]. Here, double-counting does not occur, because the filter naturally tracks inter-robot correlations in terms of the off-diagonal covariance matrix elements. The disadvantage of this approach lies in the fact that it requires a central processing unit




Fig. 1. This picture shows a five robot localization problem in an indoor environment according to the publicly available data set recorded by Leung et al. [15]. The three robots on the left observe each other, while the rightmost robot observes two landmarks. Accordingly, its error ellipse is small. The fifth (blue) robot has a large error covariance as it did not detect any landmarks, nor was it able to improve its pose estimate with the help of pairwise observations in the most recent steps.

for updating the joint filter. In addition, all robots need to instantaneously send all their motion and measurement data to this unit introducing a substantial communication overhead.

However, in many relevant multi-robot applications – for example underwater, in mines, or in large-scale environments, – communication might be energy consuming, error-prone, slow, or simply not possible at all times. Therefore, decentralized architectures that reduce the need for communication to a minimum are desirable. Particularly appealing are algorithms restricting exchange of information between two robots to points in time where they actually meet. Establishing the correct joint estimate in such a context demands for storage and bookkeeping of measurements [14]. If the memory needed to store measurements is limited and robots are only allowed to communicate at rare occasions, it is impossible to reproduce the centralized EKF. In particular the covariance matrix of the joint system can only be approximated. Here, the double-counting problem can occur if the off-diagonal terms of the matrix are underestimated or completely neglected. Accordingly, in such a situation the estimates of the poses cannot be guaranteed to be consistent.

Our contribution is an EKF-based localization algorithm that is able to approximate the covariance under the above-mentioned constraints on communication and memory:

- Exchange of information takes place only between two robots obtaining a relative measurement. Our approach does not require communication with further teammates. The communication complexity¹ is thus $\mathcal{O}(1)$ per relative measurement.
- The algorithm is recursive in the sense that no measurements have to be stored. For a team of N robots, each robot only updates the belief about its own pose and $N - 1$ matrices that carry information about the correlations with the estimates of the teammates.
- In addition to relative observations, the algorithm also supports absolute positioning of individual robots. It furthermore allows for generic measurement models.

II. RELATED WORK

The problem of multi-robot localization has received considerable attention in the past. Roumeliotis and Bekey [22] present a distributed algorithm that is equivalent to the joint EKF. In their approach, each robot processes its own odometry and measurements are communicated to all teammates. For a team of N robots, this results in communication complexity of $\mathcal{O}(N)$ per robot-robot perception. In contrast to that, our approach is of communication complexity $\mathcal{O}(1)$ and is applicable even if the exchange of information is limited to robots performing a relative measurement. Similar communication networks without centralized entity are used in gossip protocols [3]. As opposed to our method, all centralized-equivalent approaches share the need for either a particular structure of the communication graph [2, 11, 17, 26] or the storage of measurement information [14].

In the following, we concentrate on EKF-based CL algorithms that meet our assumptions – namely algorithms that forgo storage of measurement information and rely only on pairwise communication in the context of a relative measurement. There are approaches in which the robots treat the poses of teammates as deterministic parameters [20] and cross-correlations are simply neglected. This entails the risk of becoming over-confident, or forces the robots to follow certain motion patterns [12].

A more elaborate approach is to treat incorporation of relative measurements as fusion of estimates with unknown correlations. To avoid over-confidence, *Covariance Intersection* (CI) or *Split Covariance Intersection* (SCI) techniques constitute a natural choice. Approaches based on CI [4] treat estimates as if they were maximally correlated. Approaches based on SCI [16, 25] split the covariances into a dependent and an independent component. The fusion of consistent estimates by CI or SCI provably preserves consistency. However, as a price, they are in general overly pessimistic. Moreover – in contrast to our approach – the implementations as presented in the literature, have to assume perception systems capable of identifying the relative poses of neighboring robots. More precisely, they require relative range and orientation measurements and cannot be applied to systems that, for example, only

provide relative range or bearing measurements. For details on different measurement models, we refer the reader to the work of Martinelli et al. [18].

Karam et al. [10] present an approach based on state exchange. It only allows independent states to be fused and is thus not optimal. Moreover, it requires relative pose detection. Unlike the already mentioned approaches, the interleaved update algorithm by Bahr et al. [1] can handle general measurement models. In this approach, each vehicle of a team of N robots has to maintain and update up to 2^N estimates and covariances and has to keep track of their dependencies with all the team members. Robot i can only use those estimates of robot j that are not (directly or indirectly via other robots) correlated to robot i . Compared to this algorithm, our method has the advantage that each robot has to maintain only one estimate and N matrices (as opposed to 2^N estimates and matrices). Moreover, our algorithm is recursive in the sense that only the most recent estimates are used, which separates it from the approaches of Bahr et al. [1] and Karam et al. [10].

III. THE PROPOSED APPROACH

In a team of N robots, let X_i be the pose of robot i in a fixed reference frame. Throughout this paper, we consider the problem of localizing N vehicles navigating on a plane, although our approach is not limited to this case. Under this assumption, a robot pose is a triple $X_i = [x_i, y_i, \theta_i]^T$ with Cartesian coordinates $[x_i, y_i]$ and orientation θ_i . The number N can change over time and does not need to be known to the robots beforehand. Each robot is provided with an estimate of its own pose \hat{X}_i^t and the corresponding covariance matrix Σ_{ii}^t at each time step t . We call $\text{bel}_i^t = \{\hat{X}_i^t, \Sigma_{ii}^t\}$ the *belief* of robot i at time t . The belief of the joint system is

$$\text{bel}^t = \{\hat{X}^t, \Sigma^t\}, \quad (1)$$

with $\hat{X}^t = [\hat{X}_1^t; \dots; \hat{X}_N^t]$ and $\Sigma^t = [\Sigma_{ij}^t]_{1 \leq i, j \leq N}$.

At the beginning of their mission the robots might be uncorrelated, which means $\Sigma_{ij}^t = 0$ for all $i \neq j$, and the individual beliefs are sufficient to represent bel^t . When robot i detects robot j at time t , we have in general $\Sigma_{ij}^{t+1} \neq 0$. Inspired by Roumeliotis and Bekey [22], we decompose

$$\Sigma_{ij}^{t+1} = \sigma_{ij}^{t+1} (\sigma_{ji}^{t+1})^T,$$

where we can choose any possible decomposition. In Algorithm 1 given below, for example, we set $\sigma_{ij}^{t+1} = \Sigma_{ij}^{t+1}$ and $\sigma_{ji}^{t+1} = \mathbb{I}$. In the following, robots i and j independently update σ_{ij}^{t+1} respectively σ_{ji}^{t+1} . The details are presented in Section III-C.

We assume the belief (1) to be distributed among the team such that each robot i carries bel_i^t and $\{\sigma_{ij}^t\}_{\forall j \neq i}$. During their mission, the robots perform

- **Motion:** Each robot has access to a control, which is a velocity command or odometry information. We allow for generic motion models.
- **Private Measurements:** Any of the robots might or might not perform measurements that are not shared with

¹We define *communication complexity* as the number of edges required in the communication graph [11].

the group, e.g., detection of known landmarks or GPS. We allow for a heterogeneous team and generic measurement models.

- **Relative Measurements:** Any of the robots might or might not perform measurements that include itself and a teammate. Again, we allow for varying and generic detection models. In particular, we do not demand access to relative poses.
- **Communication:** Relative measurements are the only occasions where communication takes place. When robot i detects robot j , both share $\text{bel}_{i/j}$ and $\sigma_{ij/ji}$. We do not require communication to further teammates.

Algorithm 1 provides an overview of how the robots perform the corresponding updates. We derive the occurring matrices G_i^t , M^{private} , and M^{rel} as well as the functions predictBelief , correctBelief , $\text{relCorrectBelief}^{a/b}$, and $\text{relCorrectCorrelation}$ in III-A to III-C.

Algorithm 1 Update Procedure for Robot i

- 1: **Input:** $\text{bel}_i^t, \{\sigma_{ij}^t\}_{\forall j \neq i}, A$
 - 2: **Output:** $\text{bel}_i^{t+1}, \{\sigma_{ij}^{t+1}\}_{\forall j \neq i}$
 - 3: **if** A is a control **then**
 - 4: $\text{bel}_i^{t+1} \leftarrow \text{predictBelief}(\text{bel}_i^t, A)$
 - 5: **for** $1 \leq j \leq N, j \neq i$ **do**
 - 6: $\sigma_{ij}^{t+1} \leftarrow G_i^t(\text{bel}_i^t, A) \cdot \sigma_{ij}^t$
 - 7: **if** A is a private measurement **then**
 - 8: $\text{bel}_i^{t+1} \leftarrow \text{correctBelief}(\text{bel}_i^t, A)$
 - 9: **for** $1 \leq j \leq N, j \neq i$ **do**
 - 10: $\sigma_{ij}^{t+1} \leftarrow M^{\text{private}}(\text{bel}_i^t) \cdot \sigma_{ij}^t$
 - 11: **if** A is a relative measurement to robot j **then**
 - 12: send to robot j : $\text{bel}_i^t, \sigma_{ij}^t, A$
 - 13: receive from robot j : $\text{bel}_j^t, \sigma_{ji}^t$
 - 14: $\Sigma_{ij}^t \leftarrow \sigma_{ij}^t (\sigma_{ji}^t)^T$
 - 15: $\text{bel}_i^{t+1} \leftarrow \text{relCorrectBelief}^a(\text{bel}_i^t, \text{bel}_j^t, \Sigma_{ij}^t, A)$
 - 16: $\Sigma_{ij}^{t+1} \leftarrow \text{relCorrectCorrelation}(\text{bel}_i^t, \text{bel}_j^t, \Sigma_{ij}^t, A)$
 - 17: **for** $1 \leq k \leq N, k \notin \{i, j\}$ **do**
 - 18: $\sigma_{ik}^{t+1} \leftarrow M^{\text{rel}}(\text{bel}_i^t, \text{bel}_i^{t+1}) \cdot \sigma_{ik}^t$
 - 19: $\sigma_{ij}^{t+1} \leftarrow \Sigma_{ij}^{t+1}$
 - 20: **if** A is a relative measurement received from robot j **then**
 - 21: receive from robot j : $\text{bel}_j^t, \sigma_{ji}^t$
 - 22: send to robot j : $\text{bel}_i^t, \sigma_{ij}^t$
 - 23: $\Sigma_{ij}^t \leftarrow \sigma_{ij}^t (\sigma_{ji}^t)^T$
 - 24: $\text{bel}_i^{t+1} \leftarrow \text{relCorrectBelief}^b(\text{bel}_i^t, \text{bel}_j^t, \Sigma_{ij}^t, A)$
 - 25: **for** $1 \leq k \leq N, k \notin \{i, j\}$ **do**
 - 26: $\sigma_{ik}^{t+1} \leftarrow M^{\text{rel}}(\text{bel}_i^t, \text{bel}_i^{t+1}) \cdot \sigma_{ik}^t$
 - 27: $\sigma_{ij}^{t+1} \leftarrow \mathbb{I}$
-

A. Motion

We assume the motion of different robots to be independent and corrupted by Gaussian noise ν_x, ν_u ,

$$X_i^{t+1} = g(X_i^t, U_i^t + \nu_u) + \nu_x,$$

with motion model g , where U_i^t can either be a velocity command or an odometry measurement. If robot i performs a motion, the correct prediction step for the joint system can be derived from the standard EKF and is given by

$$\hat{X}_i^{t+1} = g(\hat{X}_i^t, U_i^t) \quad (2)$$

$$\Sigma_{ii}^{t+1} = G_i^t \Sigma_{ii}^t (G_i^t)^T + R_i^t \quad (3)$$

$$\Sigma_{ij}^{t+1} = G_i^t \Sigma_{ij}^t \quad (4)$$

$$\hat{X}_j^{t+1} = \hat{X}_j^t \quad (5)$$

$$\Sigma_{jj}^{t+1} = \Sigma_{jj}^t, \quad (6)$$

for all $j \neq i$ with linearization $G_i^t = \frac{\partial g(X, U)}{\partial X}(\hat{X}_i^t, U_i^t)$ and noise covariance R_i^t . These equations correspond to (24) in [22]. In the case of odometry measurements with noise distributed as $\nu_u \sim \mathcal{N}(0; M_i^t)$ and $\nu_x = 0$, for example, we have $R_i^t = V_i^t M_i^t V_i^{tT}$, with $V_i^t = \frac{\partial g(X, U)}{\partial U}(\hat{X}_i^t, U_i^t)$. The crucial insight here – as already pointed out by Roumeliotis and Bekey [22] – is that robot i can correctly update the cross-correlations to all teammates by updating

$$\sigma_{ij}^{t+1} = G_i^t \sigma_{ij}^t.$$

This is because, if robots i and j meet, they can reproduce their cross-correlation by the matrix multiplication

$$\sigma_{ij}^{t+1} (\sigma_{ji}^{t+1})^T = G_i^t \sigma_{ij}^t (\sigma_{ji}^t)^T = \Sigma_{ij}^{t+1}.$$

The function predictBelief in Algorithm 1 is defined by (2) and (3).

B. Private Measurements

We define a private measurement to be a function of the observing robot's pose, corrupted by Gaussian noise

$$z_i^t = h(X_i^t) + \nu_{pm},$$

with $\nu_{pm} \sim \mathcal{N}(0, Q_i^t)$ and measurement model h . In the following paragraphs, we will derive the equations used in our approach for the private measurement update. The derivations prove that this component of our method is identical to the well-known Schmidt-Kalman filter [23]. Let i receive a private measurement and $j \in 2^{\{1, \dots, N\}} \setminus \{i\}$ indicate any subset of the remaining robots. Then, as we show in Section VI-A, the exact correction step for the joint system is

$$\hat{X}_i^{t+1} = \hat{X}_i^t + K_i^t \left[z_i^t - h(\hat{X}_i^t) \right] \quad (7)$$

$$\Sigma_{ii}^{t+1} = (\mathbb{I} - K_i^t H_i^t) \Sigma_{ii}^t \quad (8)$$

$$\Sigma_{ij}^{t+1} = (\mathbb{I} - K_i^t H_i^t) \Sigma_{ij}^t \quad (9)$$

$$\hat{X}_j^{t+1} = \hat{X}_j^t + K_j^t \left[z_i^t - h(\hat{X}_i^t) \right] \quad (10)$$

$$\Sigma_{jj}^{t+1} = \Sigma_{jj}^t - K_j^t H_i^t \Sigma_{ij}^t, \quad (11)$$

with the linearization $H_i^t = \frac{\partial h(X)}{\partial X}(\hat{X}_i^t)$, the Kalman gains $K_{i/j}^t = \Sigma_{ii/ji}^t (H_i^t)^T S^{-1}$, and the invertible residual covariance $S = H_i^t \Sigma_{ii}^t H_i^{tT} + Q_i^t$.

We can execute the correction updates (7)-(9) without communication because all necessary terms are known to robot i .

The function `correctBelief` in Algorithm 1 is defined by (7) and (8). The cross-correlation update (9) is realized by the matrix multiplication

$$\sigma_{ij}^{t+1} = (\mathbb{I} - K_i^t H_i^t) \sigma_{ij}^t = M^{private} \sigma_{ij}^t,$$

similar to the prediction step in Section III-A.

To avoid communication, we approximate updates (10) and (11) to leave the corresponding beliefs unchanged, i.e.,

$$\begin{aligned} \hat{X}_j^{t+1} &\approx \hat{X}_j^t =: \tilde{X}_j^{t+1} \\ \Sigma_{jj}^{t+1} &\approx \Sigma_{jj}^t =: \tilde{\Sigma}_{jj}^{t+1}. \end{aligned}$$

In summary we arrive at a set of equations equivalent to the Schmidt-Kalman Filter². In expectation, this approximation is consistent with respect to the joint system. This is due to the following two observations: Given the correct belief bel^t , we have expectation $E_{z_i^t}(\hat{X}_j^{t+1}) = \hat{X}_j^t$; the difference of the approximated $\tilde{\Sigma}^{t+1}$ and the exact covariance Σ^{t+1} is a positive semi-definite matrix, i.e.,

$$\tilde{\Sigma}^{t+1} - \Sigma^{t+1} = \begin{pmatrix} 0 & 0 \\ 0 & \Sigma_{ji}^t (H_i^t)^T S^{-1} H_i^t \Sigma_{ij}^t \end{pmatrix} \succeq 0.$$

C. Relative Measurements

We define relative measurements – most generally – as noisy observations depending on the state of two robots

$$r_{ij}^t = f(X_i^t, X_j^t) + \nu_{rm},$$

with measurement model f . To meet the conditions for applying the EKF, we assume the noise term to be normally distributed $\nu_{rm} \sim \mathcal{N}(0; Q_{ij}^t)$. We emphasize again, that we neither require the relative measurements to be a function of pose differences of two robots $r_{ij}^t = f(X_i^t - X_j^t) + \nu_{rm}$, nor do we assume that the relative measurements provides the observing robot with an estimation of the observed robot's pose. Both are special cases of our general model. In particular we can handle the practically relevant case of range and/or bearing measurements. When a robot obtains a relative measurement – w.l.o.g. i observes j – the exact correction step for the joint system can be written as follows. We present the derivation in Section VI-A. Here k represents the indices of all $N - 2$ robots that do not participate at the relative measurement:

$$\hat{X}_{i/j}^{t+1} = \hat{X}_{i/j}^t + K_{i/j} \left[r_{ij}^t - f(\hat{X}_i^t, \hat{X}_j^t) \right] \quad (12)$$

$$\Sigma_{ii/jj}^{t+1} = (\mathbb{I} - K_{i/j} F_{i/j}) \Sigma_{ii/jj}^t - K_{i/j} F_{j/i} \Sigma_{ji/i}^t \quad (13)$$

$$\Sigma_{ij}^{t+1} = (\mathbb{I} - K_i F_i) \Sigma_{ij}^t - K_i F_j \Sigma_{jj}^t \quad (14)$$

$$\Sigma_{ik}^{t+1} = (\mathbb{I} - K_i F_i) \Sigma_{ik}^t - K_i F_j \Sigma_{jk}^t, \quad (15)$$

with linearization $F_{i/j} = \frac{\partial f(X_i, X_j)}{\partial X_{i/j}}(\hat{X}_{i/j}^t)$, and Kalman gain

$$K = \begin{bmatrix} K_i \\ K_j \end{bmatrix} = \begin{bmatrix} \Sigma_{ii}^t F_i^T + \Sigma_{ij}^t F_j^T \\ \Sigma_{jj}^t F_j^T + \Sigma_{ji}^t F_i^T \end{bmatrix} S^{-1},$$

²Please note that Julier [9] has used the same structure to reduce the computational cost of the related problem of *simultaneous localization and mapping*.

with residual covariance

$$S = [F_i, F_j] \begin{bmatrix} \Sigma_{ii}^t & \Sigma_{ij}^t \\ \Sigma_{ji}^t & \Sigma_{jj}^t \end{bmatrix} \begin{bmatrix} F_i^T \\ F_j^T \end{bmatrix} + Q_{ij}^t.$$

Our update equations (12)-(15) generalize the update equations by Roumeliotis and Bekey [22] to arbitrary measurement models. As argued in Section III-B, we approximate

$$\begin{aligned} \Sigma_{kk}^{t+1} &\approx \Sigma_{kk}^t \\ \hat{X}_k^{t+1} &\approx \hat{X}_k^t, \end{aligned}$$

for $k \in 2^{\{1, \dots, N\} \setminus \{i, j\}}$ and arrive at update equations equivalent to the Schmidt-Kalman filter [23]. Communication between i and j is allowed at the time of the relative measurement and they share the measurement r_{ij}^t and their beliefs $bel_{i/j}^t = \{\hat{X}_{i/j}^t, \Sigma_{ii/jj}^t\}$. Additionally, they reproduce their correlation by sharing $\sigma_{ij/ji}$. We can correctly perform the updates (12)-(14). The exactly calculated cross-correlation (14) is decomposed into $\Sigma_{ij}^{t+1} = \sigma_{ij}^{t+1} (\sigma_{ji}^{t+1})^T$ and distributed among the two participating robots.

The only terms we have not taken care of yet are the correlations between participating and non-participating robots (15). In contrast to [22], we do not rely on communication with other teammates. The problem in this case is that – unlike (9) for private measurements – the update cannot be written as a simple matrix multiplication of the form $\Sigma_{ik}^{t+1} = M^{rel} \Sigma_{ik}^t$. In particular, we cannot reproduce the terms $\Sigma_{ik}^t = \sigma_{ik}^t (\sigma_{ki}^t)^T$ and $\Sigma_{jk}^t = \sigma_{jk}^t (\sigma_{kj}^t)^T$ that are necessary for the correct update. This is because we do not rely on communication with robot k , that carries the terms σ_{ki}^t and σ_{kj}^t .

To arrive at the general form of approximations that maintain the decentralized structure, we rewrite the Schmidt-Kalman filter update step (13)-(15). For better readability and w.l.o.g. we define $a := \{i, j\}$ and $b := \{1, \dots, N\} \setminus \{i, j\}$. With help of this notation we get

$$\Sigma_{aa}^{t+1} = (\mathbb{I} - K [F_i, F_j]) \Sigma_{aa}^t$$

and thus we have

$$\begin{aligned} \Sigma_{ab}^{t+1} &= (\mathbb{I} - K [F_i, F_j]) \Sigma_{ab}^t = \Sigma_{aa}^{t+1} (\Sigma_{aa}^t)^{-1} \Sigma_{ab}^t \\ &= \underbrace{\begin{bmatrix} \Sigma_{ii}^{t+1} & \Sigma_{ij}^{t+1} \\ \Sigma_{ji}^{t+1} & \Sigma_{jj}^{t+1} \end{bmatrix} \begin{bmatrix} \Sigma_{ii}^t & \Sigma_{ij}^t \\ \Sigma_{ji}^t & \Sigma_{jj}^t \end{bmatrix}^{-1}}_{:=A} \begin{bmatrix} \Sigma_{ib}^t \\ \Sigma_{jb}^t \end{bmatrix}. \quad (16) \end{aligned}$$

To avoid the need for communication, we have to use a block-diagonal approximation of A .

A naive approximation of A and thus of (15) can be achieved by taking the block-diagonal matrix of $(\mathbb{I} - K [F_i, F_j])$, which leads to

$$\Sigma_{ik}^{t+1} = (\mathbb{I} - K_i F_i) \Sigma_{ik}^t - K_i F_j \Sigma_{jk}^t \approx (\mathbb{I} - K_i F_i) \Sigma_{ik}^t. \quad (17)$$

In our experiments, approximation (17) yields unsatisfying results. A key contribution of this paper is the more reasoned approximation

$$\Sigma_{ik}^{t+1} \approx M^{rel} \Sigma_{ik}^t = \Sigma_{ii}^{t+1} (\Sigma_{ii}^t)^{-1} \Sigma_{ik}^t, \quad (18)$$

where we assume that the joint covariance Σ^t is positive definite at any time, which can be assumed due to imperfect sensors and due to motion noise. Approximation (18) corresponds to setting

$$A \approx \tilde{A} := \begin{bmatrix} \Sigma_{ii}^{t+1} (\Sigma_{ii}^t)^{-1} & 0 \\ 0 & \Sigma_{jj}^{t+1} (\Sigma_{jj}^t)^{-1} \end{bmatrix}$$

in (16). It has the following mathematical justification:

Proposition 1: Under the assumption that the position estimates of two robots are either strongly correlated or totally uncorrelated, \tilde{A} is the unique block diagonal approximation, which minimizes the approximation error with respect to the operator norm corresponding to the 1-norm for vectors, i.e.,

$$\tilde{A} = \underset{B = \begin{bmatrix} B_1 & 0 \\ 0 & B_2 \end{bmatrix}, B_1, B_2 \in \mathbb{R}^{d \times d}}{\operatorname{argmin}} \|\Sigma_{ab}^{t+1} - B \Sigma_{ab}^t\|_1,$$

where d is the dimension of a single robot's state space. The statement also holds for every entry-wise or column-wise matrix norm.

Especially, if the estimates of robots i and j are strongly correlated before the measurement then the approximation (18) is exact.

Proof: See Appendix B. ■

Note that there are instances where both (17) and (18) underestimate cross-correlations. Assume, for example, $\Sigma_{ik}^t = 0$ and $\Sigma_{jk}^t \neq 0$ before the relative measurement. Then (15) implies in general $\Sigma_{ik}^{t+1} \neq 0$ after the relative measurement. But with (17) or (18), it remains $\Sigma_{ik}^{t+1} = 0$. However, our experiments indicate, that (18) does not only outperform (17), but also leads to good overall localization performance in practice.

All terms necessary to calculate M^{rel} are known to robot i , which can now realize the update by setting $\sigma_{ik}^{t+1} = M^{rel} \sigma_{ik}^t$. In Algorithm 1, functions $\operatorname{relCorrectBelief}^{a/b}$ are given by (12)-(13) and function $\operatorname{relCorrectCorrelation}$ is given by (14).

IV. EXPERIMENTS WITH REAL WORLD DATA

We test our approach on the publicly available *UTIAS Multi-Robot Cooperative Localization and Mapping Dataset* [15]. A fleet of five (two-wheel differential drive) robots get range and bearing measurements with known correspondences to each other and up to 15 distinguishable landmarks. The dataset includes odometry and measurements together with pose ground truth. The duration is over 4.7 hours, spread over nine different runs ranging from 15 to 70 minutes. In the ninth run, the environment contains some barriers as depicted in Fig. 1. We let the robots perform localization based on six different strategies (listed below), using dead reckoning and the relative measurements. Additionally, we allow one of the robots to process its landmark measurements. To increase the amount of data fivefold, we consecutively allow each of the robots (one at a time) to use the landmark measurements. This provides us with over 23.6 hours (not entirely independent) data. Then we drop all relative bearing

measurements, which leaves the robots with relative range only measurements. This serves to demonstrate that our approach can deal with generic measurement models. Following, we replace all relative measurements by simulated relative pose measurements, consisting of the pose of the observed robot in the observing robot's reference frame. This experiment allows us to compare against Covariance Intersection [4], which relies on relative measurements providing the observing robot with an estimate of the observed robot's pose.

We compare the following EKF-based approaches:

- Collaborative Localization using the centralized joint extended Kalman Filter (EKF), e.g., [22]. (ground truth)
- Our approach (DCL).
- Naive version of our approach (NDCL): we use approximation (17) instead of (18).
- Covariance Intersection (CI) [4]: information to be fused is assumed to be maximally correlated.
- Total Naive Collaborative Localization (NCL): cross-correlations are neglected.
- Single Robot Localization (SL): robots can detect landmarks but no other robots.

Note that to realize the centralized equivalent approach (EKF), each measurement demands for $N - 1 = 4$ edges in the communication graph. In contrast, our approach (DCL) requires only one edge for each relative measurement and no communication at all for private measurements. The datasets include a total number of $n_{rel} = 173289$ relative measurements and $n_{priv} = 42605$ private measurements on average for one of the robots performing landmark measurements. Thus, EKF needs $(N - 1)(n_{rel} + n_{priv}) / (n_{rel}) = 4.98$ times more edges than our approach.

A. Accuracy Analysis

We calculate the root-mean-square error averaged over all five robots (RMSE_t^E) of estimator E at time step t . As measure for the performance of an estimator in one run, we subtract the corresponding value of the centralized filter

$$P^E = \frac{1}{T} \sum_{t=1}^T (\operatorname{RMSE}_t^E - \operatorname{RMSE}_t^{EKF}),$$

where T is the absolute number of time steps in that run. This choice is justified by the following observation. Due to linearization errors, approximate measurement and motion models, and randomness in the system, even EKF sometimes produces wrong estimates. In these cases it is unlikely (though possible) that any approximation of EKF yields better results. This is illustrated in Fig. 2. It shows the root-mean-square error (RMSE) averaged over all five robots in the second run, where the relative measurements consist of range information only. In the time between second 400 to 500, all algorithms (including the correct EKF) encounter problems. For the rest, our approach (DCL) yields localization with RMSE below 50 cm. After around 1200 seconds, EKF yields bad localization results, while our algorithm remains stable. This is due to a robot processing landmark measurements with wrong associations. In EKF the robot instantaneously communicates these

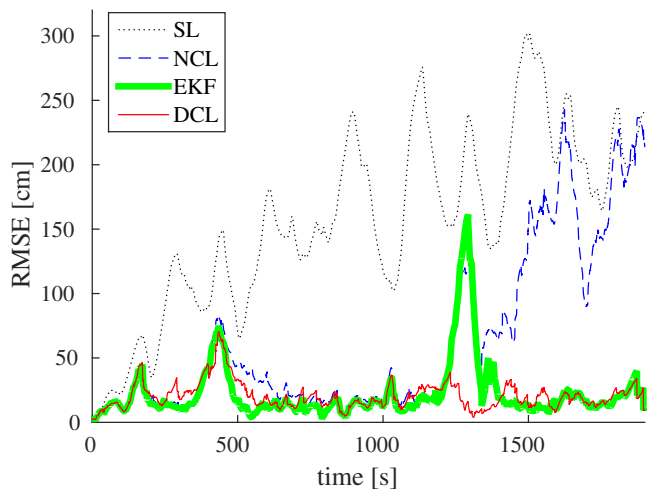


Fig. 2. Root-mean-square error averaged over all five robots in the second run. With our approach (DCL), the error remains below 50 cm throughout the whole experiment. Using single robot localization (SL), the error diverges due to the fact that only a subset of robots detect landmarks. At first, the naive approach (NCL) keeps up with DCL. But due to overconfidence its localization performance starts to decrease after around 1200 seconds.

measurements which corrupts all other robots’ estimates. We emphasize that in this case our algorithm (DCL) outperforms EKF just accidentally by not communicating measurement information. For the rest of the time, it performs almost identically well as EKF. Using single robot localization (SL), the error diverges due to the fact that only a subset of robots detect landmarks. At first, the naive approach (NCL) keeps up with DCL. But due to overconfidence its localization performance starts to decrease after around 1200 seconds.

Table I shows the values P^E in centimeters for each of the four approximate algorithms in the nine runs.

TABLE I
POSITIONING ACCURACY (P^E IN CM) ON REAL WORLD DATA WITH RANGE AND BEARING MEASUREMENTS

Run	DCL	NDCL	NCL	SL
1	1.19	1.94	6.16	171.19
2	1.52	1.75	4.27	148.15
3	-0.22	2.89	7.62	161.63
4	1.38	1.29	10.14	131.59
5	1.06	0.92	3.91	117.79
6	0.79	1.26	3.53	91.40
7	1.32	1.74	3.79	80.40
8	2.27	4.08	10.63	207.19
9	2.45	37.15	12.10	339.05
∅	1.48	6.65	7.70	177.71

In every single run, our approach (DCL) yields the best absolute localization results – except for the fifth run where the naive version of our algorithm (NDCL) is slightly better. The average is taken over all runs weighted with the respective duration. Here, DCL outperforms all other approaches. Additionally, the correct joint EKF outperforms our approximate algorithm only by less than 1.5 centimeters although making use of nearly five times as many communication links.

We repeat this experiment four times while consecutively

allowing each of the robots (one at a time) to receive the landmark measurements. That way we get 45 runs with relative range and bearing measurements. We erase all relative bearing measurements to get another 45 runs with range-only relative measurements. We generate another 45 runs by replacing all relative measurements by simulated relative pose measurements. Tab. II shows mean and variance of P^E over all 45 runs (weighted by the duration) for the different measurement types and algorithms. For all algorithms, we additionally count the number of runs in which it outperforms all other approximate algorithms. No matter which kind of relative measurements we use, our approach outperforms all of the compared methods in absolute accuracy.

TABLE II
POSITIONING ACCURACY (P^E IN CM) ON REAL WORLD DATA

RANGE AND BEARING					
	DCL	NDCL	NCL	CI	SL
mean(P^E)	2.68	29.83	6.25	–	183.56
variance(P^E)	3.42	42.94	4.77	–	63.55
# best runs	31	11	3	–	0
RANGE ONLY					
	DCL	NDCL	NCL	CI	SL
mean(P^E)	0.69	18.65	9.93	–	155.03
variance(P^E)	7.89	32.70	11.60	–	57.49
# best runs	23	10	12	–	0
SIMULATED RELATIVE POSE					
	DCL	NDCL	NCL	CI	SL
mean(P^E)	2.59	13.83	8.88	7.84	180.78
variance(P^E)	2.03	20.11	7.39	5.68	57.96
# best runs	37	2	2	4	0

B. Robustness Analysis

As stated before, even the correct EKF might diverge. However – if the covariance is not underestimated – the filter might recover from that after subsequent measurements. To demonstrate robustness of our approximation, we compare mean time to failure and the relative number of recoveries after failures for the different approaches. We define a failure to be a point in time where the root-mean-square error of the joint system exceeds 50 cm. We define a recovery as a point in time where the root-mean-square error of the joint system drops below 10 cm after a failure. The recovery ratio R is the number of recoveries divided by the number of total failures. Mean time to failure T is the mean time difference between a recovery (or the start of a run) and the next failure. Fig. 3 serves as an example. It shows the estimation error of the x-component over time for one of the robots that only rely on dead reckoning and relative measurements. All compared algorithms - including EKF (not shown) - lose positioning. In contrast to the naive approach (NCL), our approach (DCL) is able to recover from the bad localization around $t = 600s$, and yield an accurate estimate later on.

Table III displays the mean time to failure in minutes and the recovery ratio in percent. For each measurement type, these values are taken over the complete 23.6 hours of data. Our algorithm (DCL) outperforms all other algorithms – including

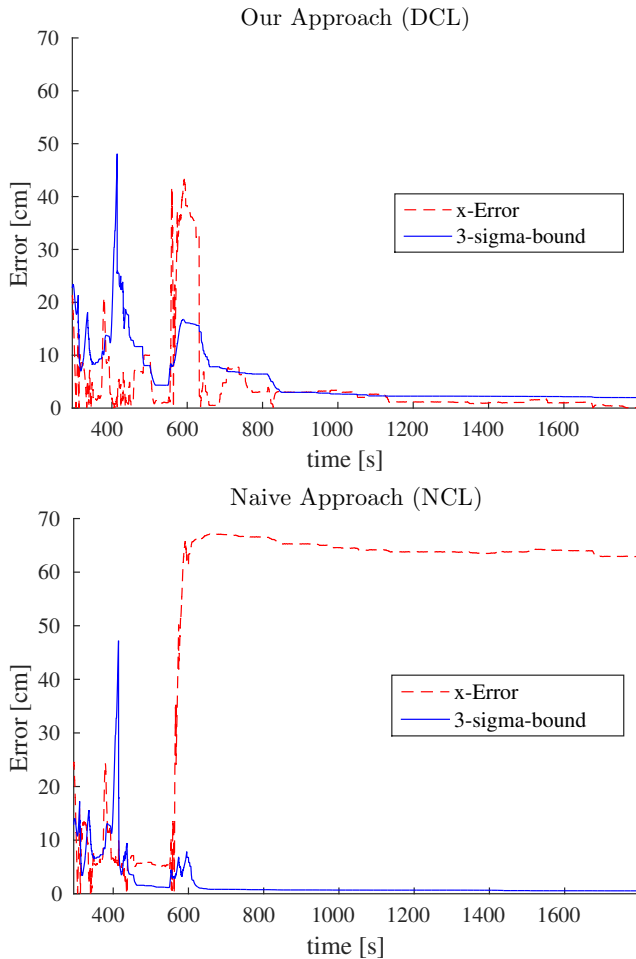


Fig. 3. Absolute difference $|x|$ between the correct x-component (of robot 4 in run 3) and the corresponding value estimated with our algorithm (DCL, plot on top) and the naive approach (NCL, plot at the bottom). This is an example, where all algorithms (including the correct centralized EKF) lose positioning around time 550s-650s. In contrast to the naive one, our approach is able to recover from that and yields accurate localization.

EKF – regarding mean time to failure. With respect to recovery rate, our approach is approximately identical to EKF. Note that mean time to failure is very low for Covariance Intersection (CI). We accredit this to the fact that in CI only the observed robot profits from a relative measurement. However, conservative approximation in CI achieves a very high recovery ratio. The experiments demonstrate our algorithm to be as robust as the correct EKF while using nearly five times less communication links and forgoing storage of measurements and the need for global communication.

V. CONCLUSION

This paper introduces a fully decentralized, EKF-based algorithm for Collaborative Localization. To the best of our knowledge, it is the first approach that tracks inter-robot correlations while fulfilling all of the following relevant conditions: the algorithm does not require storage of measurements, it supports generic measurement models, and communication is limited to pairs of robots that obtain a relative measurement.

TABLE III
MEAN TIME TO FAILURE T IN MINUTES AND RECOVERY RATIO R IN PERCENT FOR REAL WORLD DATA.

RANGE AND BEARING						
	EKF	DCL	NDCL	NCL	CI	SL
T	42.37	46.89	24.79	38.98	–	2.65
R	75.00	72.00	74.42	71.43	–	0
RANGE ONLY						
	EKF	DCL	NDCL	NCL	CI	SL
T	11.14	11.99	10.81	7.64	–	2.63
R	57.14	57.63	54.10	41.94	–	0
SIMULATED RELATIVE POSE						
	EKF	DCL	NDCL	NCL	CI	SL
T	41.53	45.13	40.51	44.41	8.41	2.65
R	78.57	76.00	74.07	79.17	93.80	0

Each robot maintains only the latest estimate of its own pose. These particularly hard conditions make the approach applicable to a wide range of applications.

We test our approach on real world data. It outperforms existing decentralized EKF-based approaches with respect to accuracy and robustness. Albeit our estimator is not provably consistent in general, our experiments imply that it is as robust as the correct EKF, which – in contrast to our method – rests on the assumption of persistent availability of global communication or has substantial memory requirements.

VI. APPENDIX

A. Derivation of the Update Equations

We derive the update equations (12)-(15) for relative measurements $r_{ij}^t = f(X_i^t, X_j^t) + \mathcal{N}(0, Q)$. W.l.o.g. we assume a system of three robots where robot 1 detects robot 2. Thus, the Jacobian can be written as $F := \frac{\partial f(X_1, X_2)}{\partial \hat{X}}(\hat{X}) = [F_1, F_2, 0]$ with $F_{1/2} = \frac{\partial f(X_1, X_2)}{\partial X_{1/2}}(\hat{X})$.

We write out the standard Kalman Filter update equations [24, Ch. 1] for the joint system with covariance Σ . The Kalman gain is calculated as

$$\begin{aligned}
 K &= \Sigma F^T (F \Sigma F^T + Q)^{-1} \\
 &= \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \Sigma_{13} \\ \Sigma_{21} & \Sigma_{22} & \Sigma_{23} \\ \Sigma_{31} & \Sigma_{32} & \Sigma_{33} \end{bmatrix} \begin{bmatrix} F_1^T \\ F_2^T \\ 0 \end{bmatrix} \\
 &= \left([F_1, F_2, 0] \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \Sigma_{13} \\ \Sigma_{21} & \Sigma_{22} & \Sigma_{23} \\ \Sigma_{31} & \Sigma_{32} & \Sigma_{33} \end{bmatrix} \begin{bmatrix} F_1^T \\ F_2^T \\ 0 \end{bmatrix} + Q \right)^{-1} \\
 &= \begin{bmatrix} \Sigma_{11} F_1^T + \Sigma_{12} F_2^T \\ \Sigma_{21} F_1^T + \Sigma_{22} F_2^T \\ \Sigma_{31} F_1^T + \Sigma_{32} F_2^T \end{bmatrix} \cdot S^{-1} =: \begin{bmatrix} K_1 \\ K_2 \\ K_3 \end{bmatrix},
 \end{aligned}$$

with residual covariance

$$S = [F_1, F_2] \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix} \begin{bmatrix} F_1^T \\ F_2^T \end{bmatrix} + Q.$$

For the updated covariance Σ^{t+1} , we have

$$\begin{aligned}
 \Sigma^{t+1} &= (\mathbb{I} - KF)\Sigma \\
 &= \begin{bmatrix} \mathbb{I} - K_1 F_1 & -K_1 F_2 & 0 \\ -K_2 F_1 & \mathbb{I} - K_2 F_2 & 0 \\ -K_3 F_1 & -K_3 F_2 & \mathbb{I} \end{bmatrix} \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \Sigma_{13} \\ \Sigma_{21} & \Sigma_{22} & \Sigma_{23} \\ \Sigma_{31} & \Sigma_{32} & \Sigma_{33} \end{bmatrix}.
 \end{aligned}$$

Element-wise evaluation yields

$$\begin{aligned}\Sigma_{ii}^{t+1} &= (\mathbb{I} - K_i F_i) \Sigma_{ii} - K_i F_j \Sigma_{ji} \\ \Sigma_{ij}^{t+1} &= (\mathbb{I} - K_i F_i) \Sigma_{ij} - K_i F_j \Sigma_{jj} \\ \Sigma_{i3}^{t+1} &= (\mathbb{I} - K_i F_i) \Sigma_{i3} - K_i F_j \Sigma_{j3}\end{aligned}$$

for $i, j \in \{1, 2\}, i \neq j$.

We derived the covariance update equations for relative measurements. The corresponding update for \hat{X} can be derived in the same, straightforward manner. Notice, that – in our notation – a private measurement is a special case of a relative measurement. Thus, to derive the update equations (7)-(11), we just replace $F = [F_1, F_2, 0]$ by $H = [H_1, 0, 0]$.

B. Proof of Proposition 1

Let X, Y, Z be d -dimensional random variables and $\langle X, Y \rangle := \text{cov}(X, Y) = E((X - E(X))(Y - E(Y))^T)$. Under the assumptions of imperfect sensors and motion noise, one can assume that the joint covariance Σ^t is positive definite at any time. Thus we have

$$\begin{pmatrix} \langle X, X \rangle & \langle X, Y \rangle & \langle X, Z \rangle \\ \langle Y, X \rangle & \langle Y, Y \rangle & \langle Y, Z \rangle \\ \langle Z, X \rangle & \langle Z, Y \rangle & \langle Z, Z \rangle \end{pmatrix} = \begin{pmatrix} \Sigma_{ii} & \Sigma_{ij} & \Sigma_{ik} \\ \Sigma_{ji} & \Sigma_{jj} & \Sigma_{jk} \\ \Sigma_{ki} & \Sigma_{kj} & \Sigma_{kk} \end{pmatrix}$$

and all of its principal sub-matrices are positive definite and invertible. Then we get with [13, Corollary 1] that

$$v^T \langle X, Y \rangle \langle Y, Y \rangle^{-1} \langle Y, X \rangle v \leq v^T \langle X, X \rangle v,$$

for all $v \in \mathbb{R}^d$. With help of the symmetric and positive definite square roots, this is equivalent to

$$\left\| \langle Y, Y \rangle^{-\frac{1}{2}} \langle Y, X \rangle v \right\| \leq \left\| \langle X, X \rangle^{\frac{1}{2}} v \right\|,$$

for all $v \in \mathbb{R}^d$. Using the induced matrix norm we obtain

$$\begin{aligned}\|\langle Y, X \rangle\| &:= \max_{\|v\|=1} \|\langle Y, X \rangle v\| \\ &\leq \|\langle Y, Y \rangle^{\frac{1}{2}}\| \max_{\|v\|=1} \|\langle Y, Y \rangle^{-\frac{1}{2}} \langle Y, X \rangle v\| \\ &\leq \|\langle Y, Y \rangle^{\frac{1}{2}}\| \|\langle X, X \rangle^{\frac{1}{2}}\|.\end{aligned}\quad (19)$$

This inequality is related to the usual Cauchy-Schwarz-Inequality valid for scalar products. Decompose now Z and X in an Y -dependent part and a part orthogonal to Y with respect to $\langle \cdot, \cdot \rangle$, namely

$$\begin{aligned}Z &= \langle Z, Y \rangle (\langle Y, Y \rangle)^{-1} Y + \hat{Z}, \\ X &= \langle X, Y \rangle (\langle Y, Y \rangle)^{-1} Y + \hat{X}.\end{aligned}$$

The goal is to apply (19) to $\langle \hat{X}, \hat{Z} \rangle$. Therefore we need to prove that $\langle \hat{Z}, \hat{Z} \rangle$ is positive definite. We obtain, that

$$\langle \hat{Z}, \hat{Z} \rangle = \langle Z, Z \rangle - \langle Z, Y \rangle \langle Y, Y \rangle^{-1} \langle Y, Z \rangle$$

is the Schur complement of a positive definite matrix and thus it is positive definite. Thus (19) implies

$$\begin{aligned}\|\Sigma_{ik} - \Sigma_{ij} \Sigma_{jj}^{-1} \Sigma_{jk}\| &= \|\langle \hat{X}, \hat{Z} \rangle\| \\ &\leq \|\langle \hat{X}, \hat{X} \rangle^{\frac{1}{2}}\| \|\langle \hat{Z}, \hat{Z} \rangle^{\frac{1}{2}}\| \\ &= \left\| (\Sigma_{ii} - \Sigma_{ij} \Sigma_{jj}^{-1} \Sigma_{ji})^{\frac{1}{2}} \right\| \cdot \left\| (\Sigma_{kk} - \Sigma_{kj} \Sigma_{jj}^{-1} \Sigma_{jk})^{\frac{1}{2}} \right\|.\end{aligned}\quad (20)$$

First we show that Approximation (18) is exact if the estimates of robots i and j are strongly correlated before the relative measurement, namely if $\Sigma_{ii}^t = \Sigma_{ij}^t (\Sigma_{jj}^t)^{-1} \Sigma_{ji}^t$. With (20) and assuming the second term of its right hand side to be bounded, we get $\Sigma_{jk}^t = \Sigma_{ji}^t (\Sigma_{ii}^t)^{-1} \Sigma_{ik}^t$, after exchanging i and j . Thus, we arrive at (18):

$$\begin{aligned}\Sigma_{ik}^{t+1} &= (\mathbb{I} - K_i F_i) \Sigma_{ik}^t - K_i F_j \Sigma_{ji}^t (\Sigma_{ii}^t)^{-1} \Sigma_{ik}^t \\ &= \Sigma_{ii}^{t+1} (\Sigma_{ii}^t)^{-1} \Sigma_{ik}^t.\end{aligned}$$

Let B be any block-diagonal matrix as in Proposition 1, i.e.,

$$B = \begin{bmatrix} B_1 & 0 \\ 0 & B_2 \end{bmatrix}, \quad B_1, B_2 \in \mathbb{R}^{d \times d},$$

where d is the dimension of a single robot's state space. For the case $\Sigma_{ij}^t = 0$ we have

$$\begin{aligned}M &:= \Sigma_{ab}^{t+1} - B \Sigma_{ab}^t \\ &= \begin{bmatrix} \Sigma_{ii}^{t+1} (\Sigma_{ii}^t)^{-1} \Sigma_{ib}^t + \Sigma_{ij}^{t+1} (\Sigma_{jj}^t)^{-1} \Sigma_{jb}^t - B_1 \Sigma_{ib}^t \\ \Sigma_{ji}^{t+1} (\Sigma_{ii}^t)^{-1} \Sigma_{ib}^t + \Sigma_{jj}^{t+1} (\Sigma_{jj}^t)^{-1} \Sigma_{jb}^t - B_2 \Sigma_{jb}^t \end{bmatrix},\end{aligned}$$

with $M = (m_{ij})_{ij} \in \mathbb{R}^{m \times n}$, $m = 2d$, and $n = (N - 2)d$. As we are interested in minimizing the 1-norm of M , i.e.,

$$\|M\|_1 = \max_{1 \leq j \leq n} \sum_{i=1}^m |m_{ij}|,$$

it suffices to minimize subsets of columns independently (the same holds for any entry-wise or column-wise norm). In particular, we can set $b = k$.

For $\Sigma_{ik}^t = 0$ it is easy to see that

$$\begin{bmatrix} B_1 & 0 \\ 0 & \Sigma_{jj}^{t+1} (\Sigma_{jj}^t)^{-1} \end{bmatrix} \in \underset{A}{\text{argmin}} \|M\|$$

for all $B_1 \in \mathbb{R}^{d \times d}$ and thus $\tilde{A} \in \text{argmin}_A \|M\|$. Analogously we get $\tilde{A} \in \text{argmin}_A \|M\|$ for $\Sigma_{jk}^t = 0$, which proves the proposition.

Note that, if the estimate of robot k is strongly correlated to both i and j before the relative measurement, then also estimates i and j are strongly correlated – the case which we have already covered. This can be seen as follows. We have $\Sigma_{ii}^t = \Sigma_{ik}^t (\Sigma_{kk}^t)^{-1} \Sigma_{ki}^t$ and $\Sigma_{kk}^t = \Sigma_{kj}^t (\Sigma_{jj}^t)^{-1} \Sigma_{jk}^t$, or equivalently $\mathbb{I} = (\Sigma_{jj}^t)^{-1} \Sigma_{jk}^t (\Sigma_{kk}^t)^{-1} \Sigma_{ki}^t$. We combine this with $\Sigma_{ik}^t = \Sigma_{ij}^t (\Sigma_{jj}^t)^{-1} \Sigma_{jk}^t$ from (20) and conclude

$$\begin{aligned}\Sigma_{ii}^t &= \Sigma_{ik}^t (\Sigma_{kk}^t)^{-1} \Sigma_{ki}^t \\ &= \Sigma_{ij}^t (\Sigma_{jj}^t)^{-1} \Sigma_{jk}^t (\Sigma_{kk}^t)^{-1} \Sigma_{kj}^t (\Sigma_{jj}^t)^{-1} \Sigma_{ji}^t \\ &= \Sigma_{ij}^t (\Sigma_{jj}^t)^{-1} \Sigma_{ji}^t,\end{aligned}$$

which shows that the estimates of robots i and j are strongly correlated.

ACKNOWLEDGMENTS

This work has been partially supported by the European Commission under ERC-AG-PE7-267686-LIFENAV, the Graduate School of Robotics in Freiburg, and the State Graduate Funding Program of Baden-Württemberg.

REFERENCES

- [1] A. Bahr, M. R. Walter, and J. J. Leonard. Consistent cooperative localization. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, pages 3415–3422, 2009.
- [2] T. Bailey, M. Bryson, H. Mu, J. Vial, L. McCalman, and H. F. Durrant-Whyte. Decentralised cooperative localisation for heterogeneous teams of mobile robots. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, pages 2859–2865, 2011.
- [3] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE/ACM Transactions on Networking (TON)*, 14(SI):2508–2530, 2006.
- [4] L. C. Carrillo-Arce, E. D. Nerurkar, J. L. Gordillo, and S. I. Roumeliotis. Decentralized multi-robot cooperative localization using covariance intersection. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1412–1417, 2013.
- [5] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. A Probabilistic Approach to Collaborative Multi-Robot Localization. *Autonomous Robots*, 8(3):325–344, 2000.
- [6] A. Howard, M. J. Mataric, and G. Sukhatme. Localization for mobile robot teams using maximum likelihood estimation. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 434–439, 2002.
- [7] A. Howard, M. J. Mataric, and G. Sukhatme. Putting the 'I' in 'team': an ego-centric approach to cooperative localization. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, pages 868–874, 2003.
- [8] G. P. Huang, N. Trawny, A. I. Mourikis, and S. I. Roumeliotis. Observability-based consistent EKF estimators for multi-robot cooperative localization. *Autonomous Robots*, 30(1):99–122, 2011.
- [9] S. J. Julier. A sparse weight Kalman filter approach to simultaneous localisation and map building. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1251–1256, 2001.
- [10] N. Karam, F. Chausse, R. Aufrere, and R. Chapuis. Cooperative Multi-Vehicle Localization. In *IEEE Intelligent Vehicles Symposium*, pages 564–570, 2006.
- [11] S. S. Kia, S. F. Rounds, and S. Martinez. A centralized-equivalent decentralized implementation of Extended Kalman Filters for cooperative localization. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3761–3766, 2014.
- [12] R. Kurazume, S. Nagata, and S. Hirose. Cooperative positioning with multiple robots. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, pages 1250–1257, 1994.
- [13] P. Lavergne. A Cauchy-Schwarz inequality for expectation of matrices. Discussion papers, Department of Economics, Simon Fraser University, 2008.
- [14] K. Y. K. Leung, T. D. Barfoot, and H. H. T. Liu. Decentralized Localization of Sparsely-Communicating Robot Networks: A Centralized-Equivalent Approach. *IEEE Transactions on Robotics*, 26(1):62–77, 2010.
- [15] K. Y. K. Leung, Y. Halpern, T. D. Barfoot, and H. H. T. Liu. The UTIAS Multi-robot Cooperative Localization and Mapping Dataset. *The International Journal of Robotics Research*, 30(8):969–974, 2011.
- [16] H. Li and F. Nashashibi. Cooperative Multi-Vehicle Localization Using Split Covariance Intersection Filter. *IEEE Intelligent Transportation Systems Magazine*, 5(2): 33–44, 2013.
- [17] A. Martinelli. Improving the precision on multi robot localization by using a series of filters hierarchically distributed. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1053–1058, 2007.
- [18] A. Martinelli, F. Pont, and R. Siegwart. Multi-Robot Localization Using Relative Observations. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, pages 2797–2802, 2005.
- [19] E. D. Nerurkar, S. I. Roumeliotis, and A. Martinelli. Distributed maximum a posteriori estimation for multi-robot cooperative localization. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, pages 1402–1409, 2009.
- [20] S. Panzneri, F. Pascucci, and R. Setola. Multirobot Localisation Using Interlaced Extended Kalman Filter. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2816–2821, 2006.
- [21] A. Prorok, A. Bahr, and A. Martinoli. Low-cost collaborative localization for large-scale multi-robot systems. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, pages 4236–4241, 2012.
- [22] S. I. Roumeliotis and G. A. Bekey. Distributed multi-robot localization. *IEEE Transactions on Robotics and Automation*, 18(5):781–795, 2002.
- [23] S. F. Schmidt. Applications of state space methods to navigation problems. *Advances in Control Systems*, 3: 293–340, 1966.
- [24] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. The MIT Press, 2005.
- [25] T. R. Wanasinghe, G. K. I. Mann, and G. G. Raymond. Decentralized Cooperative Localization for Heterogeneous Multi-robot System Using Split Covariance Intersection Filter. In *Proceedings of IEEE International Conference on Computer and Robot Vision (CRV)*, pages 167–174, 2014.
- [26] T. R. Wanasinghe, G. K. I. Mann, and R. G. Gosine. Distributed Leader-Assistive Localization Method for a Heterogeneous Multirobotic System. *IEEE Transactions on Automation Science and Engineering*, 12(3):795–809, 2015.