Distributed Cooperative Localization with Efficient Pairwise Range Measurements

Anwar Quraishi and Alcherio Martinoli

Distributed Intelligent Systems and Algorithms Laboratory,
School of Architecture, Civil and Environmental Engineering,
École Polytechnique Fédérale de Lausanne, Switzerland
anwar.quraishi@epfl.ch, https://disal.epfl.ch/auvdistributedsensing

Abstract. We present a method based on covariance intersection for cooperative localization with pairwise range-only relative measurements. Our method was designed for underwater robots equipped with an acoustic communication and ranging system. Range measurements are not sufficient to compute a complete relative 3D position. Therefore, covariance intersection is performed in a transformed space along their relative estimated positions, while preserving cross-correlations between other state variables. Given the characteristics of the acoustic channel, only one robot can transmit data or a ranging request at a time, hence the pairwise limitation. We also present a heuristic for choosing a peer robot for a range measurement by maximizing mutual information. Our method places no further restrictions on the order, timing or scheduling of relative measurements. We evaluated our method for accuracy and consistency, and present results from simulations as well as outdoor experiments.

1 Introduction

A number of robot actions, such as path planning and spatial information gathering depend on accurate localization. Robots operating on land or in the air can often exploit external positioning references such as Global Navigation Satellite Systems (GNSSs), cameras and range finders. However, access to such external references is limited in many scenarios, such as indoor places, caves and underwater environments. In Cooperative Localization (CL), a team of robots operating together shares and fuses information and relative measurements to improve their individual localization accuracy. This allows new, more accurate information about position acquired by one robot to be propagated to other robots. We have previously demonstrated acoustic navigation for underwater robots with static and moving surface beacons [1]. However, by using CL and taking turns to resurface periodically for GNSS reception, underwater robots can function without relying on external beacons, as illustrated in Fig. 1.

An important characteristic of CL is that information sharing makes position estimates of all robots correlated [2]. Position estimates of other robots coupled with relative measurements are key ingredients for CL. However, measurement models in localization frameworks used typically, such as Kalman filtering, often

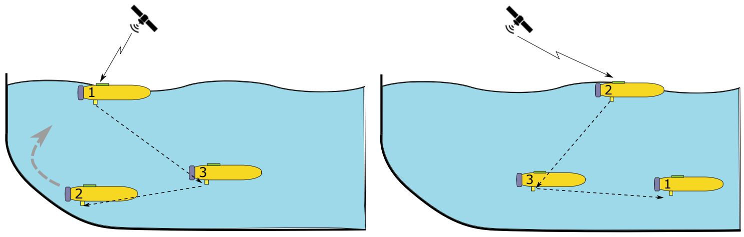

Fig. 1: Underwater robots periodically resurface for GNSS position reception, and then use CL to share the improved position estimate with other robots.

assume that the measurement information is uncorrelated with robot state estimates [3]. Therefore, relative measurement updates in CL need to be handled carefully to avoid inconsistent and overconfident position estimates. A further problem of peer-selection arises in CL with pairwise relative measurements. The ideal choice of a peer is that which provides a relative measurement resulting in the maximum information gain.

CL has received considerable attention in the realm of multi-robot systems. The simplest approach for CL is to gather all robot observations and relative measurements and process them at a central location [4]. Roumeliotis et al. in [5] showed that a centralized Kalman filter for CL can be decomposed into smaller, communicating filters which are distributed among the robots. Later, Luft et al. in [6] limited communication exclusively between the pair of robots that obtained a relative measurement. This reduced the communication cost for N robots from $\mathcal{O}(N)$ to $\mathcal{O}(1)$. However, both approaches require inter-robot cross-correlation terms to be communicated along with robot position estimates. This is necessary but adds communication cost.

Treating correlated information as independent, as in [7], can make new position estimates overconfident [8]. In the work of Bahr et al. [2], each robot maintains a set of several state estimates, and keeps track of their dependencies with other peers through careful book-keeping. A robot can use only those estimates of another peer that are not correlated (directly or via other peers) with its own estimate. However, the memory, computation and communication requirements of this approach grow exponentially with the number of robots.

In [9] and [10], the authors use an approach based on distributed Maximum Likelihood Estimation, where each robot optimizes its own state given the relative measurements. The problem of inter-robot correlations does not arise in this approach. While there is no formal proof of convergence, it performs well in practice.

Another technique for addressing the problem of inter-robot correlation is Covariance Intersection (CI) [8] [11] [12] [13], which treats estimates as if they were maximally correlated. They do not require robots to communicate cross-correlation estimates, saving communication bandwidth. The price to pay is that they are pessimistic in that they overestimate uncertainty. This overestimation is addressed in a hybrid approach, called Split-CI [14, 15]. It splits the covariances into dependent and independent components. However, it requires communication of both, the independent and the dependent covariance matrices.

The second problem addressed in this paper is optimal peer-selection, which is similar to optimal sensor or beacon selection. In [16], three best (fixed) ultrasonic ranging beacons are selected based on Geometric Dilution of Precision (GDOP) for indoor localization. An entropy-minimization-based sensor selection approach for fixed target tracking is presented in [17]. [18] presents a selection criteria based on mutual information for general Bayesian filtering problems.

In this paper, we present a CI-based fully distributed cooperative localization algorithm for range-only relative measurements. We have developed our method keeping in mind limitations of underwater acoustic ranging and communication. To that end, CI offers important advantages. It does not need inter-robot correlations to be computed and communicated. While this is inefficient in that the uncertainty in robot positions is *overestimated*, it allows for a completely distributed implementation of CL. Additionally, CI is provably consistent.

When a relative measurement comprises of a full relative pose along with the position estimate of another robot, it is straightforward to perform CI. This is not the case with range measurements, which are one dimensional, whereas robot positions can be two- or three-dimensional. Therefore, a range measurement update can have a direct influence on robot position only along the relative position vector between two robots. We perform CI in a transformed space aligned with an estimate of this vector to update the robot position and the corresponding position covariance. It is important to note that range measurements can indirectly influence all state variables via the cross-correlations between them. Our method accounts for and preserves the cross-correlation between the state variables. While our method adds computational cost, the cost of internal computation is much lower than that of acoustic communication.

We also derive a peer-selection heuristic for choosing the best peer for performing a pairwise range measurement. In the trivial case when uncertainty in positions of other robots is not known, the best choice is a peer robot that is along (or closest to) the direction of highest uncertainty. However, the knowledge of the said uncertainty exists because robots broadcast their position estimates during a range measurement. We use the mutual information between current position estimate of the robot and a potential range measurement to derive a mathematical expression for scoring the 'usefulness' of peer robots.

In summary, our work consists of two main contributions. (1) We derive a linear transform of the robot state in which to perform CI with range-only relative measurements, while preserving the cross-covariances in the robot state, and (2) we derive a mathematical expression to rank peer robots in the order of the amount of information a relative update would provide.

2 Methodology

We consider a team of N underwater robots navigating in a three dimensional space. Robots choose a peer with which to perform a range measurement and

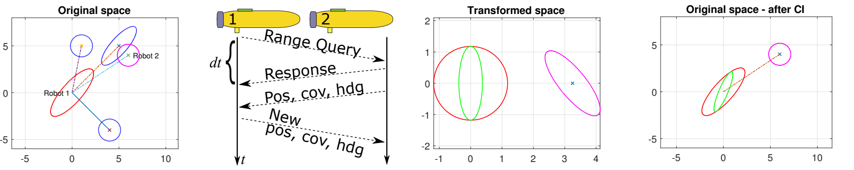

(a) Peer selection (b) Query (c) CI (d) New estimate

Fig. 2: (a) Robot 1 chooses Robot 2 for a range measurement. (b) R2 responds to the range query of R1 and also sends its position estimate. (c) CI is performed in a transformed space. (d) Updated estimate of R1 (in green).

transmit a ranging query. On receiving a response, CI is performed using our our *projected covariance intersection* method, while preserving all the cross-correlation terms in the state. The sequence of steps are illustrated in Fig. 2, and explained in the rest of this section.

2.1 State Description

The states of the robots are assumed to be Gaussian random variables and are expressed using the mean-covariance parametrization,

$$s_i^t \sim \mathcal{N}(X_i^t, \Sigma_i^t), \tag{1}$$

where $i \in \{1, \dots, N\}$. Our algorithm is agnostic to the formulation of the state variable, except that it requires the three-dimensional position represented in a fixed frame to be a part of the state.

$$X_i^t = [x_i, y_i, z_i, \dots] \tag{2}$$

2.2 Motion and Individual Measurements

Motion and private measurement updates are purely internal to a robot do not require any communication between them. They are processed individually by robots using any kind of sensor fusion framework such as a Kalman filter.

2.3 Range-Based Covariance Intersection

Communication between two robots is required only when they perform a range measurement. No other robots are required to be involved in the communication. However, we assume all communication is broadcast, so other robots can listen to this communication.

Consider a robot, i with position and position covariance $\vec{p}_i, R_i$. It chooses to query another robot j and obtains a range measurement r_{ij} with standard deviation σ_r. Robot j also transmits its position estimate $\vec{p}_j, R_j$. We define the range vector as

$$\vec{r}_{ij} = r_{ij} \frac{\vec{p}_j - \vec{p}_i}{\|\vec{p}_j - \vec{p}_i\|}, \tag{3}$$

Note that only magnitude of $\vec{r}_{ij}$, i.e., r_{ij}, is measured, the actual vector is obtained from estimated quantities.

A transform $\mathcal{F}$ is applied to the state of robot i so that the dimension along $\vec{r}_{ij}$ is decoupled from the rest of the elements of the state.

$$q = \mathcal{F}(s_i - X_i), \tag{4}$$
$$\Sigma'_q = \mathcal{F}\Sigma_i\mathcal{F}^{\mathrm{T}}, \tag{5}$$

where q is the transformed state vector. Such a transform can be obtained by setting

$$\mathcal{F} = VW^{-\frac{1}{2}}T^{\mathrm{T}}, \tag{6}$$

where T and W are obtained from the eigenvalue decomposition of the state covariance matrix Σ_i. V is obtained from the Gram-Schmidt orthogonalization [19] starting with the vector $\vec{r}_{ij}$. The key feature of this transform is that q is zero-mean and the covariance matrix is identity, i.e., all the elements are uncorrelated.

$$q \sim \mathcal{N}(0_{M\times 1}, \mathrm{I}_{M\times M}). \tag{7}$$

This method has similarly been used in [20] for truncating Gaussian PDFs given a hard constraint. The correlations between elements of X_i are not lost but encoded in T and W. The matrix V in the transform rotates the space in a way that the first dimension in the transformed space is along the range vector. CI is then easily performed in one dimension, and only applies to the first element of q, q_x.

$$q_x \sim \mathcal{N}(0, 1). \tag{8}$$

Next, the transformation is applied to the other relevant quantities, namely $\vec{r}_{ij}$, σ_r, $\vec{p}_j$ and R_j. Note that $\mathcal{F}$ applies rotation as well as scaling.

$$\vec{r}' = \mathcal{F}\vec{r}_{ij}, \tag{9}$$
$$\sigma'_r = \|\mathcal{F}\hat{r}\|\sigma_r, \tag{10}$$
$$\vec{p}'_j = \mathcal{F}\vec{p}_j, \tag{11}$$
$$R'_j = \mathcal{F}R_j, \tag{12}$$

where $\hat{r}$ is the unit range vector. To perform CI, we are interested in the conditional probability distribution of $\vec{p}'_j$ along the transformed range vector $\vec{r}'$ as below. This is trivial to obtain for a Gaussian distribution.

$$P_j(p'_j | y' = 0, z' = 0) \sim \mathcal{N}(x'_j, \sigma'^2_j). \tag{13}$$

In the transformed space, using the range measurement and position of robot j, an estimate of the position of robot i is calculated as

$$\hat{x}'_i = x'_j - \vec{r}', \tag{14}$$
$$\hat{\sigma}'^2_i = \sigma'^2_j + \sigma'^2_r. \tag{15}$$

Thus, we have

$$\hat{q}_x \sim \mathcal{N}(\hat{x}'_i, \hat{\sigma}'^2_i). \tag{16}$$

CI is now performed between q_x and $\hat{q}_x$ (see Eqs. 8, 16) to obtain an updated estimate of the position of robot i in the transformed space.

$$\left[\sigma_{x'}^{\prime 2}\right]^{-1} = \omega \left[1\right]^{-1} + (1-\omega)\left[\hat{\sigma}_i^{\prime 2}\right]^{-1}, \tag{17}$$

$$\mu_{x'} = \left[\sigma_{x'}^{\prime 2}\right] \cdot \left[0 \cdot \omega \left[1\right]^{-1} + \hat{x}_1' \cdot (1-\omega)\left[\hat{\sigma}_i^{\prime 2}\right]^{-1}\right] \tag{18}$$

where ω is chosen to minimize $\sigma_{x'}^{\prime 2}$. After this update, the resulting state and covariance of robot 1 in the transformed space are

$$q = \left[\mu_{x'}, 0, 0, \ldots\right], \tag{19}$$

$$\Sigma_q' = \mathrm{diag}\left(\sigma_{x'}^2, 1, 1, \ldots\right). \tag{20}$$

This is illustrated in Fig. 2c. Finally, the new estimate of the state of robot 1 in the original space is computed by performing an inverse transform.

$$\hat{X}_{i|j} = TW^{\frac{1}{2}}V^{\mathrm{T}}q + X_1, \tag{21}$$

$$\hat{\Sigma}_{i|j} = TW^{\frac{1}{2}}V^{\mathrm{T}}\Sigma_q' VW^{\frac{1}{2}}T^{\mathrm{T}}, \tag{22}$$

which follows from the inverse of the transform, $\mathcal{F}^{-1} = TW^{\frac{1}{2}}V^{\mathrm{T}}$, given that T and V are orthonormal matrices (see Fig. 2d).

2.4 Peer Selection

Consider a robot i that needs to choose another robot $j \in [1, n] \setminus i$ and perform a pairwise range update (we performed experiments with $n = 3$ and $n = 4$). We seek to choose a robot j that results in the lowest posterior uncertainty, as shown in Fig. 2a. Doing so would require current position estimates of all robots, which is not feasible in view of communication constraints. However, since all communication is broadcast, all robots receive position estimates as well as heading of other robots during a range measurement (see Fig. 2b and Section 2.5). This information, coupled with a constant velocity model is used to compute the current estimates for all other robots.

For the following analysis, we assume that the state consists of only the 3D positions. We have

$$s_k \sim \mathcal{N}(p_k, R_k). \tag{23}$$

Let z_{ij} be the random variable describing the potential range measurement between i and j (including the position of a peer robot j). We would like to choose a j for which the mutual information $I(s_i; z_{ij})$ is maximized. Formally, we seek to solve the problem

$$j^* = \underset{j \in [1,n] \setminus i}{\arg\max}\ I(s_i; z_{ij}). \tag{24}$$

We know that this mutual information between two random variables can be written in terms of their differential entropy as

$$I(s_i; s_j) = h(s_i) - h(s_i | z_{ij}), \tag{25}$$

where $h(s)$ is the differential entropy of random variable s [21]. When s has a Gaussian distribution, the $h(s)$ can be evaluated as

$$h(s) = \frac{1}{2}\log(2\pi e)^n |\Sigma_s| \text{ bits}, \tag{26}$$

where $|\Sigma_s|$ is the determinant of the covariance matrix Σ_s. In Eq. 25, $P(s_i)$ is fixed, and only the posterior density $P(s_i|z_j)$ depends on the choice of j. Therefore, we can combine Eqs. 25, 26 to reformulate the problem in Eq. 24 as

$$j^* = \underset{j\in[1,n]\setminus i}{\arg\min}\ |\hat{\Sigma}_{i|j}|, \tag{27}$$

where $\hat{\Sigma}_{i|j}$ is the covariance of the posterior probability density. This was previously computed in Eq. 22. On the right side of the equation, note that T and V are orthonormal matrices, hence their determinant is 1. The posterior covariance in the transformed space Σ'_q is shown in Eq. 20 to be a diagonal matrix such that its determinant will evaluate to $\sigma^2_{x'}$. Therefore, we have

$$|\hat{\Sigma}_{i|j}| = |W||\Sigma'_q| = |W|\sigma^2_{x'}. \tag{28}$$

Considering that that W is independent of j, and following from Eqs. 15 and 17, the problem can be further simplified to

$$j^* = \underset{j\in[1,n]\setminus i}{\arg\min}\ \sigma'^2_j + \sigma'^2_{r,j}. \tag{29}$$

We recall here that the first term is the conditional distribution of the position of robot j, and the second term is the variance of the range measurement, both in the transformed space. Note that it can be easily deduced that if the first term is ignored, the minimum is obtained for robot j which is along the direction of highest uncertainty of robot i.

However, broadcasts from other robots coupled with a constant velocity model provide a current estimate for the first term, σ'^2_j, as explained earlier. This serves as a heuristic for solving the optimization problem in Eq. 29.

2.5 Range Queries and Communication

The limitation of pairwise measurements is because of constraints of most underwater acoustic transceivers. To avoid interference, only one robot can transmit a signal at a time. However, all robots can listen to any transmitted signal, if they are within communication range. We use a Time Division Multiple Access (TDMA) scheme, where each robot is assigned pre-determined time slots by rotation. Practical approaches for implementing such a scheme in underwater robots have been discussed in [1].

During its assigned slot, a robot initiates a range query and the subsequent exchange shown in Fig. 2b. Range is computed by measuring two-way-travel-time. Finally, the robot performs CI and broadcasts its updated position estimate. The duration of individual time slots depends on the time it takes to perform the exchange in Fig. 2b, which is a characteristic of the acoustic transceiver hardware used.

3 Experimental Setup

The proposed method was implemented and tested in a group of three to four robots. We evaluated our algorithms in terms of localization accuracy and estimation consistency, both in simulation as well as with outdoor experiments. The kinematics in simulation were roughly calibrated to that of real robots. In both, simulation and real experiments, the robots were programmed to follow a pre-planned trajectory. The state estimator uses proprioceptive sensors and a motion model of the robot for inertial navigation, combined with range measurements. To emulate periodic surfacing events in a team of underwater robots, two of the three or four robots in the group were allowed periodic access to GNSS positioning information, which would be passed on to other robots through the cooperative localization framework. Otherwise, GNSS positions were used only for following the trajectory and as ground truth.

Experiments are performed with two strategies of choosing a peer for range measurement. In the first one, called *cyclic*, each robot queries other robots for a range measurement turn-by-turn, in a cycle. The second one, called best-selection or *bsel*, uses the proposed peer-selection approach.

Range measurement updates are performed using the proposed CI approach, as well as an Extended Kalman Filter (EKF) approach. The EKF approach uses the standard update equations, ignoring the correlation between inter-robot position estimates. We do not implement any centralized EKF for joint estimation of positions of all robots. Each robot runs an independent state estimator.

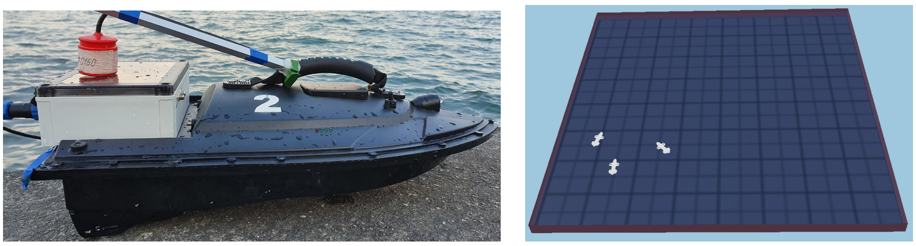

(a) ASV with acoustic modem (b) Webots simulation

Fig. 3: (a) The ASV with the acoustic modem. When in operation, the modem, attached to a rod, is suspended in water from the rear of the robot. (b) A screenshot of the Webots simulation environment with three AUVs.

3.1 Simulation

Experiments were performed in simulation with three to four robots using Webots [22], a high-fidelity robotics simulation software. A picture of the simulation environment is shown in Fig. 3b. An acoustic modem is also simulated with the same propagation delay, ranging accuracy and bandwidth as the real acoustic modem. Simulations are performed with two different sets of trajectories, one with four robots and another with three robots. We restrict access to GNSS to two robots at 0.1 Hz or 0.05 Hz (once in 10 s or 20 s). The time-slot length for acoustic transmission is set to 5 s, which means there is one range measurement in the whole system every 5 s. We perform five experimental runs for each combination of trajectory, GNSS access, peer-selection method (cyclic or bsel) and range update method (EKF or CI).

3.2 Outdoor Experiments

Outdoor experiments were performed using three Autonomous Surface Vehicles (ASVs). The ASVs, pictured in Fig. 3a, are equipped with the same sensing and computing hardware as the Vertex Autonomous Underwater Vehicles [23] from Hydromea SA. Additionally, they have continuous GNSS reception, which serves as ground truth (but the state estimator has limited access to it). All robots are equipped with a Beringia Microlink acoustic transceiver for communication as well as range measurements. The maximum data transfer speed is 10 Bytes/s, and the ranging accuracy is about 2.5 m. A range measurement, including related data exchange (Fig. 2b), takes about 3.5 s. Experiments were performed in Lake Geneva.

The robots follow a pre-planned trajectory, which is occasionally disturbed by strong waves. Robot 1 has no access to GNSS for position estimation. Robots 2 and 3 are allowed GNSS updates once in 20 s (0.05 Hz). Acoustic range is recorded every 1.5 s between each of the three pairs of robots. For the purpose of this experiment, shorter acoustic time-slots were used without any guard times for avoiding echoes and interference, and only range measurements were recorded. Position information was recorded via radio.

The recorded data is re-processed offline with different range update methods (CI, EKF) and peer-selection strategies (cyclic, bsel). Depending on the peer-selection strategy, and subject to a more realistic time-slot length of 5 s, only a subset of the recorded range measurements are used. This gave us accuracy and consistency metrics for both range measurement update methods as well as peer-selection strategies.

4 Results

4.1 Evaluation Metrics

We use the Root Mean Squared Error (RMSE) as a measure of accuracy. We compute it as the Euclidean distance between the true and estimated positions averaged over all time steps since the beginning of the experiment.

$$\mathrm{RMSE(T)} = \frac{1}{T}\sum_{t=1}^{T} \|x^t - \hat{x}^t\|. \tag{30}$$

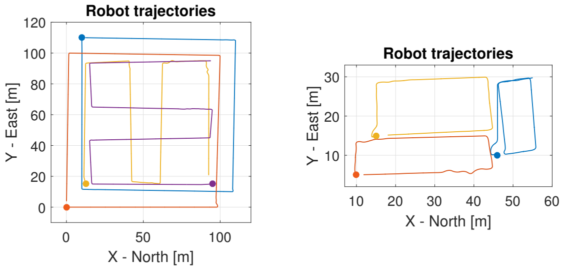

(a) T1, 4 robots (b) T2, 3 robots

Fig. 4: Two sets of trajectories, T1 and T2, in simulation. The starting point of the trajectory for each robot is shown with a dot.

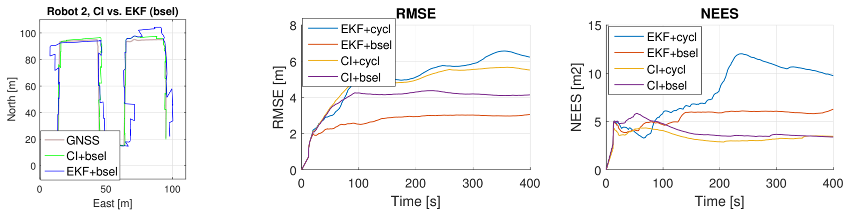

(a) EKF vs CI (b) Evolution of RMSE (c) Evolution of NEES

Fig. 5: For scenario 1: (a) Example estimated trajectory of one of the robots using CI and EKF. (b) Comparison of RMSE over time for the two peer-selection strategies. Our peer-selection method results in higher accuracy. (c) Comparison of NEES for CI and EKF. EKF has higher estimation error owing to ignoring inter-robot correlations.

RMSE is higher when the estimated trajectory is farther off from the true trajectory.

For measuring consistency, we use the Normalized Estimation Error Squared (NEES), averaged over all time steps since the beginning of the experiment.

$$\mathrm{NEES(T)} = \frac{1}{T}\sum_{t=1}^{T}(x^t - \hat{x}^t)^{\mathrm{T}}(\Sigma^t)^{-1}(x^t - \hat{x}^t). \tag{31}$$

Higher values of NEES indicate higher inconsistency between the estimated position covariance and the actual error in estimated position. This metric was introduced by Shalom et. al. in [24].

4.2 Simulations

An example of a trajectory of one of the robots estimated with CI and EKF is shown in Fig. 5a. A comparison of evolution of RMSE is shown in Fig. 5b. Our peer-selection strategy, bsel, results in an improved accuracy of estimated position compared to a cyclic strategy. This is because robots are able to predict which peer can provide most useful information based on the heuristic in Eq. 29.

	CI		EKF	
Scenarios	RMSE-Cyclic NEES-Cyclic	RMSE-bsel NEES-bsel	RMSE-Cyclic NEES-Cyclic	RMSE-bsel NEES-bsel
1. T1,GNSS 0.10Hz,2 Rob/4	5.51 3.45	4.14 3.38	6.22 9.76	3.06 6.27
2. T1,GNSS 0.05Hz,2 Rob/4	6.19 5.43	4.57 4.02	6.77 11.89	3.85 6.10
3. T2,GNSS 0.10Hz,2 Rob/3	2.23 2.71	2.14 3.69	2.34 4.58	2.05 4.31
4. T2,GNSS 0.05Hz,2 Rob/3	2.60 2.74	2.33 3.41	2.88 5.93	1.94 4.17
5. Real,GNSS 0.05Hz,2 Rob/3	3.28 5.03	2.17 4.56	3.23 6.33	2.44 7.22

Table 1: RMSE and NEES for various scenarios and estimation methods.

Evolution of NEES for scenario 1 is compared in Fig. 5c. EKF initially performs similar to CI. However, as uncertainties increase over time, they must be accounted for. Therefore, after a few range updates, NEES for EKF increases in comparison with that of CI. The results of various experimental scenarios, averaged over all runs and across all robots are tabulated in Table 1.

4.3 Outdoor Experiments

The estimated trajectory for one of the three robots along with ground-truth trajectories of the other two peers is shown in Fig. 6. The purely inertial estimate (without any GNSS or acoustic updates) is also shown for comparison. The evolution of RMSE and NEES is shown in Fig. 7a, 7b. The accuracy metrics are shown in the last row of Table 1.

The results demonstrate that the trajectory estimated with CI is more consistent compared to EKF. This is shown by a lower value of NEES, implying that the estimated uncertainty is in better agreement with the actual estimation error. Correct estimation of uncertainty is important because it is used to adaptively weight the influence of incoming measurements (e.g, via the Kalman gain in EKF). Therefore, erroneous uncertainty estimates are likely to cause higher trajectory RMSE eventually. The results also show that regardless of the range update method (CI or EKF), the proposed peer-selection strategy resulted in a lower RMSE. This shows the improvement in accuracy brought by an educated choice of peer for a relative measurement. The experimental results with real robots are in agreement with those obtained from simulations.

4.4 Computation and Data Overhead

The proposed CI approach is more expensive in computation compared to EKF. A range update with EKF would require a 2×2 matrix inversion (assuming position estimation in 2D, the vertical dimension is provided by a depth sensor),

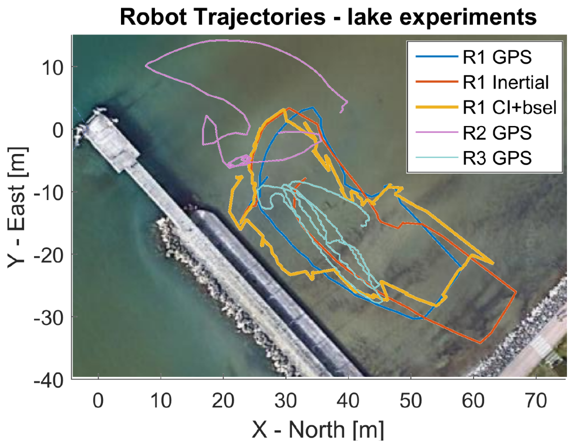

Fig. 6: Robot trajectories from the real experiment. For Robot 1, trajectory estimated with with CI and the proposed peer-selection strategy is shown. The true (GNSS) trajectory and the inertial estimate (without range measurements) are also shown for comparison. For robots 2 and 3, only the true (GNSS) trajectory is shown.

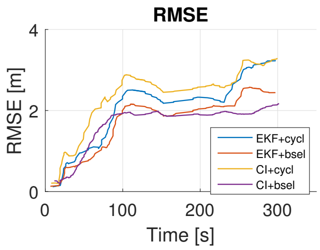

(a) RMSE

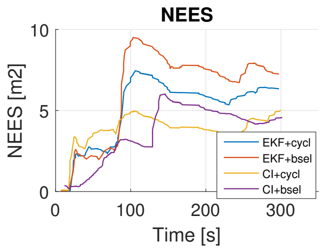

(b) NEES

Fig. 7: RMSE and NEES for the real experiment. The plots show that CI results in lower estimation error (NEES), which demonstrates more accurate estimation of uncertainty. They also show that using the proposed peer-selection strategy results in more accurate trajectory estimation.

regardless of the size of the state variable. The proposed transformation requires eigenvalue decomposition, inversion of a diagonal matrix and several matrix multiplications, the sizes of which depend on the size of the state. This adds an overhead in computation time (EKF: 0.1 ms, CI: 1.5 ms, approx.).

For the peer-selection heuristic, the only additional information added to the communication is the robot heading. All the other information is also needed for range update. Computing this metric also requires many of the operations used for CI-based range update (eigenvalue decomposition, etc.). The computation time was found to be about 1.0 ms per peer on our setup. Our setup consists of C++ code compiled with `-O2` optimization flag, running on a Raspberry Pi Zero with single core 1 GHz processor. Since there is one range update in approximately 5 s, this combined computation delay is affordable.

5 Conclusion

We presented a range-based covariance intersection method for cooperative localization. CI provides a number of advantages for cooperative localization with underwater robots, which have severe communication constraints. CI does not require communication of inter-robot correlation terms. It can be run in a fully distributed fashion, and uses information only from the two robots involved in a relative measurement. It allows robots to exploit information from their peers without any adverse effects such as overconfident estimates. We showed that CI-based cooperative localization results in a better estimation of uncertainty.

We also derived a peer-selection heuristic for performing range measurements based on an information theoretic approach. We showed that our peer-selection strategy improves localization accuracy in comparison to sequentially querying peers for a range measurement.

A number of improvements to the system are possible. During real-world operation, range queries or responses made by a robot may get lost or corrupted. At the moment, we do not try to resend a query or query another peer, resulting in no updates being performed during some time slots. A recovery strategy may reduce this 'dead time' with no updates. The peer-selection heuristic is based on estimating position of peers using a constant velocity model. When a peer changes direction, this estimate becomes invalid. A better estimate can be made if planned trajectories of peer robots are communicated in advance.

Acknowledgement: This work was partially funded by the Swiss National Science Foundation under grant CRSII2_160726/1.

References

1. A. Quraishi, A. Bahr, F. Schill, and A. Martinoli, "A Flexible Navigation Support System for a Team of Underwater Robots," in *International Symposium on Multi-Robot and Multi-Agent Systems*, 2019, pp. 70–75.
2. A. Bahr, J. J. Leonard, and M. F. Fallon, "Cooperative localization for autonomous underwater vehicles," *The International Journal of Robotics Research*, vol. 28, no. 6, pp. 714–728, 2009.
3. S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics.* Cambridge, Massachusetts: MIT Press, 2005.
4. J. Borenstein, "Experimental results from internal odometry error correction with the omnimate mobile robot," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 6, pp. 963–969, 1998.
5. S. I. Roumeliotis and G. A. Bekey, "Distributed multirobot localization," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 781–795, 2002.
6. L. Luft, T. Schubert, S. I. Roumeliotis, and W. Burgard, "Recursive decentralized localization for multi-robot systems with asynchronous pairwise communication," *The International Journal of Robotics Research*, vol. 37, no. 10, pp. 1152–1167, 2018.
7. S. Panzieri, F. Pascucci, and R. Setola, "Multirobot Localisation Using Interlaced Extended Kalman Filter," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2006, pp. 2816–2821.

8. S. J. Julier and J. K. Uhlmann, "A non-divergent estimation algorithm in the presence of unknown correlations," in *Proceedings of the 1997 American Control Conference (Cat. No.97CH36041)*, vol. 4, 1997, pp. 2369–2373 vol.4.
9. A. Howard, M. J. Matarić, and G. S. Sukhatme, "Localization for mobile robot teams using maximum likelihood estimation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1, 2002, pp. 434–439.
10. A. Howard, M. J. Matarić, and G. S. Sukhatme, "Localization for Mobile Robot Teams: A Distributed MLE Approach," in *Experimental Robotics VIII*, B. Siciliano and P. Dario, Eds. Springer Berlin Heidelberg, 2003, pp. 146–155.
11. J. Klingner, N. Ahmed, and N. Correll, "Fault-tolerant covariance intersection for localizing robot swarms," *Robotics and Autonomous Systems*, vol. 122, pp. 103–306, 2019.
12. L. C. Carrillo-Arce, E. D. Nerurkar, J. L. Gordillo, and S. I. Roumeliotis, "Decentralized multi-robot cooperative localization using covariance intersection," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Nov 2013, pp. 1412–1417.
13. M. Vasic, D. Mansolino, and A. Martinoli, "A system implementation and evaluation of a cooperative fusion and tracking algorithm based on a Gaussian Mixture PHD filter," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 4172–4179.
14. H. Li, F. Nashashibi, and M. Yang, "Split covariance intersection filter: Theory and its application to vehicle localization," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1860–1871, 2013.
15. H. Li and F. Nashashibi, "Cooperative multi-vehicle localization using split covariance intersection filter," *IEEE Intelligent Transportation Systems Magazine*, vol. 5, no. 2, pp. 33–44, 2013.
16. J. Park and J. Lee, "Beacon selection and calibration for the efficient localization of a mobile robot," vol. 32, no. 1. Cambridge University Press, 2014, pp. 115–131.
17. H. Wang, K. Yao, G. Pottie, and D. Estrin, "Entropy-based sensor selection heuristic for target localization," in *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks*, 2004, pp. 36–45.
18. E. Ertin, J. W. Fisher, and L. C. Potter, "Maximum mutual information principle for dynamic sensor query problems," in *Information Processing in Sensor Networks*. Springer Berlin Heidelberg, 2003, pp. 405–416.
19. L. Pursell and S. Y. Trimble, "Gram-schmidt orthogonalization by gauss elimination," *The American Mathematical Monthly*, vol. 98, no. 6, pp. 544–549, 1991.
20. D. Simon and D. L. Simon, "Constrained Kalman filtering via density function truncation for turbofan engine health estimation," *International Journal of Systems Science*, vol. 41, no. 2, pp. 159–171, 2010.
21. T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, Inc, 1991.
22. O. Michel, "Cyberbotics ltd. webots tm: Professional mobile robot simulation," *International Journal of Advanced Robotic Systems*, vol. 1, no. 1, pp. 39–42, 2004.
23. F. Schill, A. Bahr, and A. Martinoli, "Vertex: A New Distributed Underwater Robotic Platform for Environmental Monitoring," *International Symposium on Distributed Autonomous Robotic Systems, 2016, Springer Proceedings in Advanced Robotics*, vol. 6, pp. 679–693, 2018.
24. Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2004.