

An Experimental Study of Robust Distributed Multi-robot Data Association from Arbitrary Poses

Erik Nelson, Vadim Indelman, Nathan Michael and Frank Dellaert

Abstract In this work, we experimentally investigate the problem of computing the relative transformation between multiple vehicles from corresponding inter-robot observations during autonomous operation in a common unknown environment. Building on our prior work, we consider an EM-based methodology which evaluates sensory observations gathered over vehicle trajectories to establish robust relative pose transformations between robots. We focus on experimentally evaluating the performance of the approach as well as its computational complexity and shared data requirements using multiple autonomous vehicles (aerial robots). We describe an observation subsampling technique which utilizes laser scan autocovariance to reduce the total number of observations shared between robots. Employing this technique reduces run time of the algorithm significantly, while only slightly diminishing the accuracies of computed inter-robot transformations. Finally, we provide discussion on data transfer and the feasibility of implementing the approach on a mesh network.

1 Introduction and Related Work

In this work, we investigate the problem of computing the relative transformation between multiple vehicles based on corresponding inter-robot observations developed during autonomous operation in a common unknown environment. Applications that rely on distributed mapping and coordinated control must, in general, assume

E. Nelson (✉) · N. Michael
Carnegie Mellon University, Pittsburgh, PA 15213, USA
e-mail: enelson@cmu.edu

N. Michael
e-mail: nmichael@cmu.edu

V. Indelman · F. Dellaert
Georgia Institute of Technology, Atlanta, GA 30332, USA
e-mail: indelman@cc.gatech.edu

F. Dellaert
e-mail: dellaert@cc.gatech.edu

the existence of a shared environment representation in order to establish a common reference frame for integration of distributed observations and joint cooperative control decisions. Therefore, a fundamental capability required by these applications is a robust strategy to establish the relative pose between individual vehicles.

Several approaches exist to address the problem of establishing a consistent reference frame in the multi-robot SLAM literature based on landmarks [1], direct relative inter-robot observations [2], and multi-robot data association [3]. Landmark-based strategies require additional prior knowledge of the environment or a means to instrument the environment (e.g., beacons [4]). Approaches that leverage direct inter-robot observations, such as when one robot detects the relative presence of another robot using onboard vision [5, 6] or RF ranging [7], assume that vehicles will proximally operate at the same time. In this work, we focus on environments and systems that do not admit prior or external instrumentation as well as temporal assumptions on spatial operation. Such scenarios can occur in complex or expansive environments where vehicles operate independently with the expectation of frequent and infrequent interactions (e.g., tunnel networks [8], large buildings [9]) or at disparate time schedules. Consequently, we assume that vehicles autonomously navigate the unknown environment, concurrently estimating their location and the map of the environment, and opportunistically coordinate with other vehicles toward furthering the application objective. For this reason, techniques most related to the emphasis of this work build on data association methods which seek to establish a consistent relative transform based on the existence of mutual environment observations made by each vehicle. Cunningham et al. [10] and Montijano et al. [11] propose robust methods for establishing a relative inter-robot transformation without requiring prior knowledge of the initial relative inter-robot pose based on variations of the RANSAC algorithm. Indelman et al. [12] suggest an alternative method based on the observation that vehicles will share common incremental observations in areas historically traversed by multiple vehicles, and pursue an EM-based methodology which evaluates present and historic observations developed along the trajectory transited by the vehicles.

In this work, we pursue an experimental sequel to the method presented by Indelman et al. [12]. The study focuses on analyzing the correctness of the resulting relative transformation as well as the relationship between the algorithm's computational complexity, shared data requirements, and team size, using trials of multiple autonomous vehicles (aerial robots). We briefly summarize the technical approach in Sect. 2. Section 3 details experiments designed to evaluate the technique using a multi-robot system. The experiments assess the correctness and robustness of the approach as well as key considerations toward reducing its computational complexity for real-time performance. Specifically, we describe an observation subsampling approach which utilizes laser scan autocovariance to select salient scans for sharing, and discuss considerations for implementing the approach over a capacity constrained network. Section 4 reports on the accuracy of inter-robot transforms resulting from three indoor and outdoor multi-robot trials, and analyzes the impact of laser scan saliency on both accuracy, and time consumed by individual algorithmic steps. Section 5 closes with a discussion and summarization of the experimental design and results.

2 Technical Approach

We now briefly review the formulation proposed in our prior work and defer to this work for a detailed discussion on the approach [12]. We consider a group of R robots deployed to collaboratively operate in some unknown environment and assume the robots start from different locations, without knowledge of the existence of other vehicles. Each robot r is assumed to be capable of estimating its trajectory X^r based on observations Z^r from its onboard sensors. We represent this estimation problem in a pose graph probabilistic formulation

$$p(X^r | Z^r) \propto p(x_0^r) \prod_i p(u_{i-1,i}^r | x_{i-1}^r, x_i^r) \tag{1}$$

where $x_i^r \in X^r$ is the robot's pose at time t_i , expressed relative to some reference frame, and $p(x_0^r)$ is a prior term. Since we assume no a priori knowledge about the environment and the initial pose of the robots, the reference frame of each robot is arbitrarily set to coincide with the initial pose.

The measurement likelihood term $p(u_{i-1,i}^r | x_{i-1}^r, x_i^r)$ in (1) involves the relative pose measurement $u_{i-1,i}^r$ that can be either directly obtained from odometry measurements or calculated from vision or laser sensor observations at the two time instances t_{i-1} and t_i . We follow the standard assumption in the SLAM community and model the measurement likelihood as a Gaussian:

$$p(u_{i-1,i}^r | x_{i-1}^r, x_i^r) \propto \exp\left(-\frac{1}{2} \|u_{i-1,i}^r \ominus h(x_{i-1}^r, x_i^r)\|_{\Sigma}^2\right)$$

with Σ being the measurement noise covariance and h the measurement model that, in the case of relative pose observations and robot poses expressed in the same reference frame is $h(x_{i-1}^r, x_i^r) \doteq x_{i-1}^r \ominus x_i^r$. We follow Lu and Milios [13] and use the notation \ominus in $a \ominus b$ to express b locally in the frame of a for any two poses a, b .

The maximum a posteriori (MAP) estimate of the r th robot pose X^r using only local information is then given by

$$\hat{X}^r = \arg \max_{X^r} p(X^r | Z^r)$$

We denote by \mathcal{F} the set of multi-robot data association, with each individual data association $(r_1, r_2, k, l) \in \mathcal{F}$ representing a relative pose constraint $u_{k,l}^{r_1,r_2}$ relating between the pose of robot r_1 at time t_k and the pose of robot r_2 at time t_l . This constraint can represent both direct observation of one robot pose relative to another robot, and also the estimated relative pose based on observation of a common scene by two robots. In the latter case, it is computed from the measurements of the two robots $z_k^{r_1} \in Z^{r_1}$ and $z_l^{r_2} \in Z^{r_2}$, that can represent, for example, laser scans or image observations.

Assuming multi-robot data association \mathcal{F} has been established and appropriate constraints $u_{k,l}^{r_1,r_2}$ have been calculated, we can write a probabilistic formulation for the multi-robot joint pdf for the robots as follows:

$$p(X|Z) \propto \prod_r p(X^r|Z^r) \prod_{(r_1,r_2,k,l) \in \mathcal{F}} p(u_{k,l}^{r_1,r_2} | x_k^{r_1}, x_l^{r_2}) \quad (2)$$

where X and Z represent, respectively, the robot trajectories and the measurements.

As the robots express their local trajectories with respect to different reference systems, the measurement likelihood term in (2) is

$$p(u_{k,l}^{r_1,r_2} | x_k^{r_1}, x_l^{r_2}) \propto \exp\left(-\frac{1}{2} \|\text{err}(u_{k,l}^{r_1,r_2}, x_k^{r_1}, x_l^{r_2})\|_{\Sigma}^2\right)$$

with

$$\text{err}(u_{k,l}^{r_1,r_2}, x_k^{r_1}, x_l^{r_2}) \doteq u_{k,l}^{r_1,r_2} \ominus h(x_k^{r_1}, x_l^{r_2})$$

and

$$h(x_k^{r_1}, x_l^{r_2}) \doteq x_k^{r_1} \ominus (T_{r_2}^{r_1} \oplus x_l^{r_2})$$

The notation \oplus represents the compose operator [13], and $T_{r_2}^{r_1}$ is a transformation between the reference frames of robots r_1 and r_2 . Since the robots start operating from different unknown locations, this transformation is initially unknown and arbitrary.

While the formulation (2) assumes multi-robot data association \mathcal{F} is given, in practice it is unknown ahead of time and should therefore be established. In [12], we propose an expectation-maximization (EM) based framework to reliably infer the multi-robot data association \mathcal{F} in a multi-robot pose SLAM framework, without assuming prior knowledge on initial relative poses between the robots, i.e., unknown $T_{r_j}^{r_i}$ for all pairs $r_i, r_j \in [1, \dots, R]$. The remainder of this work experimentally evaluates the efficacy of this EM-based methodology to accurately estimate the unknown transformations, $T_{r_j}^{r_i}$, for all pairs of robots.

3 Experimental Study

3.1 Experimental Design and Approach

We consider the problem of developing accurate relative transforms between multiple aerial robots while operating in a common environment through shared laser sensor observations. As the system relies on a capacity constrained network, we focus on experimental questions relating computational complexity, shared data requirements,

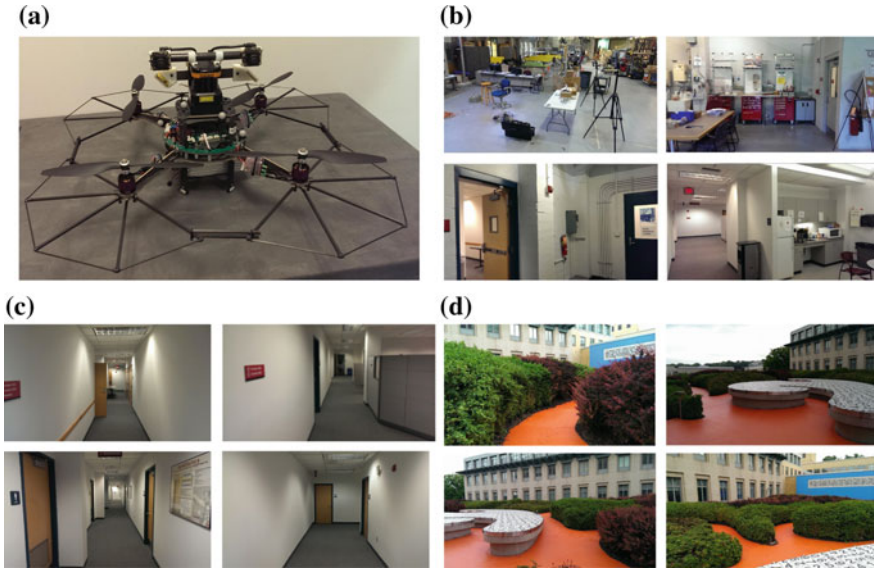


Fig. 1 **a** Quadrotor platform, equipped with an onboard computer (1.86 GHz Intel Core 2 Duo processor), IMU, laser, and beam deflector mirrors. Stereo cameras are not utilized in the described experiments. **b–d** Experiment environments for **T1**, **T2**, and **T3**

and team size. Toward studying this relationship, we propose a strategy that seeks to reduce data transfer and algorithmic complexity at the cost of reducing the accuracy of the computed relative transforms. We show the feasibility of the proposed technique experimentally and extend this discussion to larger teams where mild reductions in accuracy can permit real-time performance.

Trials of sensory information (laser scans) are captured from fleets of quadrotors operating in three different environments (Fig. 1). These environments consist of a set of paths transitioning from a wide open room with clutter to a hallway environment (trial **T1**), a series of connected corridors and hallways (trial **T2**), and an outdoor hedge maze (trial **T3**).

- T1** Three robots navigating along a path from the same initial pose before diverging and traveling long distances in different directions.
- T2** Three robots navigating through hallways from different starting poses. All robots meet, and travel in the same direction around a 10×16m loop before diverging.
- T3** Three robots navigating an outdoor structured environment. Robots navigate low to the ground to capture laser scans of bushes and walls. Robots 2 and 3 have no trajectory overlap, but capture laser scans which share features.

Pose estimates along the robot trajectories are generated through local instances of SLAM running on each quadrotor during operation. We use a laser and inertial

based SLAM implementation similar to that of [14], which leverages ICP for laser odometry [15], a histogram filter for localization, and a UKF to fuse estimates [16]. Octomap is used for 3D mapping capabilities [17].

3.2 Implementation

We pursue a centralized implementation where laser scan observations are distributed to and evaluated by a single vehicle (i.e., client-server model). Each new observation received by the server is matched against a history of all observations from other robots using a variant of ICP [18]. If there is a significantly low covariance in any one ICP match, a data association $(r_i, r_j, k, l) \in \mathcal{F}$ between the two robots, r_i and r_j , links the poses from which the matched scans were captured.

After a specified number of observations are shared by all robots, transform hypotheses between pairs of robots are established using EM [12]. For each transform hypothesis generated this way, data associations formed between robot pairs which share similarity in translation and rotation are separated into a set of inliers. The transform hypothesis which contains the highest number of inlier data associations is chosen as the most probable transform. If the number of inliers is small for all transform hypotheses, none are selected, signifying that the robots are operating in disjoint environments. If a reliable transform has not yet been discovered, hypotheses are constantly generated and reassessed as new observations are shared between robots.

3.3 Computational Complexity and Saliency of Information

Given this implementation, we choose to evaluate three steps of the algorithm in an analysis of computational complexity (Table 1). SLAM is executed locally on each robot, and is therefore omitted. The server is responsible for managing all computations in these three steps, and executes them in sequence when presented with new pose and sensor information from a robot.

- Step 1 Updates to a robot's pose graph to incorporate a new pose and any new data associations formed with other robots since the previous update.
- Step 2 ICP and transform hypothesis generation between one robot's history of shared scans and any new scans recently shared by other robots.
- Step 3 Sensor observation autocovariance calculation.

Table 1 Complexities of individual steps of the data association strategy

Algorithm step	1	2	3
Complexity	$O(E^2)$	$O(R^2 \bar{Z}' ^2)$	$O(R Z')$

Step 1 Edges are added to the multi-robot pose graph upon individual robot pose updates as well as multi-robot pose correspondences generated through data association. The pose graph is implemented with the GTSAM optimization library [19], which computes individual updates (nearly) linearly in the number of edges in the pose graph. Continuing the notation introduced in Sect. 2, the shared pose graph has $E = R(|X^r| - 1) + |\mathcal{F}|$ edges if the poses of all robot are updated at a shared frequency. Since edges are added sequentially during run time, the total number of computations performed in Step 1 is roughly $\frac{1}{2}(E^2 + E)$, resulting in a complexity of $O(E^2)$.

Step 2 To identify data associations from arbitrary local robot coordinate frames in a temporally invariant manner, every laser scan shared by one robot must be compared against each other robot's history of laser scans. Assuming robot poses are updated with a shared frequency, every laser scan shared by an individual robot will be accompanied by $R - 1$ sensor observations shared by other robots. Growth in the laser scan histories of each robot results in a growth in the number of ICP comparisons performed between observations over time. Let each robot r gather a set of observations, Z^r , and share a subset, $\bar{Z}^r \subset Z^r$, with other robots. Then the total number of ICP comparisons performed over the duration of the trial run time is

$$\begin{aligned} & R(R - 1) \sum_{i=1}^{|\bar{Z}^r|} i \\ &= \frac{1}{2} (R^2 - R) (|\bar{Z}^r|^2 + |\bar{Z}^r|) \end{aligned}$$

This result implies that the run time complexity of Step 2 is $O(R^2|\bar{Z}^r|^2)$.

Step 3 The dominant source of complexity in Step 2 arises from $R^2|\bar{Z}^r|^2$ laser scan ICP comparisons from potentially large initial offsets over the duration of the run time, as well as EM on a number of transform hypotheses quadratic in $|\bar{Z}^r|$. The algorithm can be made significantly more efficient through a reduction the number of observations in the set \bar{Z}^r as $|\bar{Z}^r|$ is not restricted to be equal to $|X^r|$. We therefore propose a subsampling strategy that seeks to reduce $|\bar{Z}^r|$ based on laser scan autocovariance [20], which can be used as a scalar measure of laser scan saliency. Autocovariance is calculated by randomly perturbing the pose from which a laser scan was captured, performing ICP to match the perturbed scan against the original, and storing the resulting transformation mean and covariance. After N such iterations, autocovariance is computed by

$$\delta = \frac{1}{\text{trace}(\Sigma)},$$

where Σ is the covariance of the N -Gaussian mixture.

To subsample, we skip nine of ten sequential laser scans and maintain at least 0.1 m in normed (x, y) pose estimate between scans. Saliencies are calculated for the remaining set. Only scans with δ greater than a threshold, δ_s , are shared with

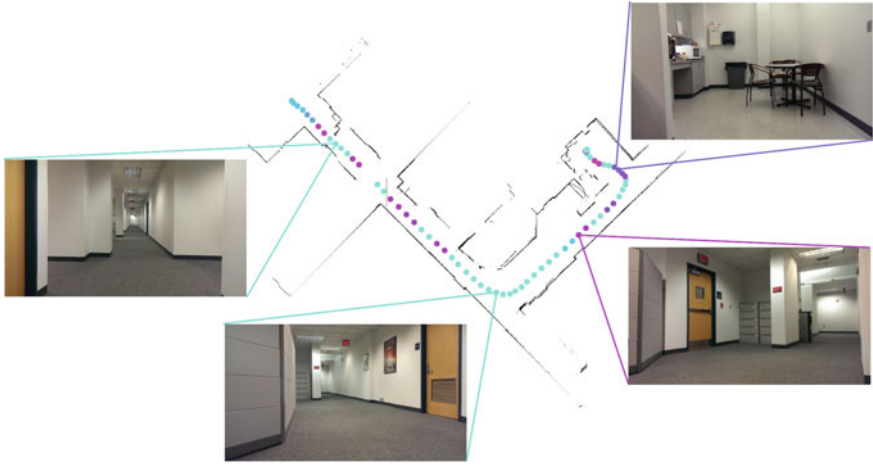


Fig. 2 Laser scan saliency along a robot's trajectory with inlaid images of the environment. Magenta corresponds to high saliency, while teal corresponds to low saliency

other robots. Because laser scan saliency is only computed once per laser scan, the run time complexity of Step 3 is $O(R|Z'|)$. Laser scan saliency is plotted along a robot trajectory from trial **T2** in Fig. 2.

3.4 Network Complexity

Constrained network capacity is a concern when requiring the distribution of large amounts of shared data. Further, mesh networks exhibit a reduced capacity when distributing packets between multiple systems [21]. A conservative model of three robots sharing uncompressed scans (approximately 34 kB per scan) limits the data sharing rate of each robot to approximately 4Hz. The rate reduces further with the number of robots where in practice an update rate of 1 Hz is expected for six vehicles [21].

4 Results and Discussion

The data association algorithm was used to calculate inter-robot relative transforms between robots in **T1**, **T2**, and **T3**. All reported transforms, $T_{r_i}^{r_1}$, are expressed as a rotation followed by a translation from the local coordinate frame of robot r_i to that of robot r_1 .

We first report on the accuracy of transforms calculated without saliency thresholding on laser scans. Trajectories, inlier data associations, and outlier data associations from **T1** are shown in the local frame of each robot in Fig. 3. Robots begin from the same initial position and diverge after sharing a large number of laser scans in

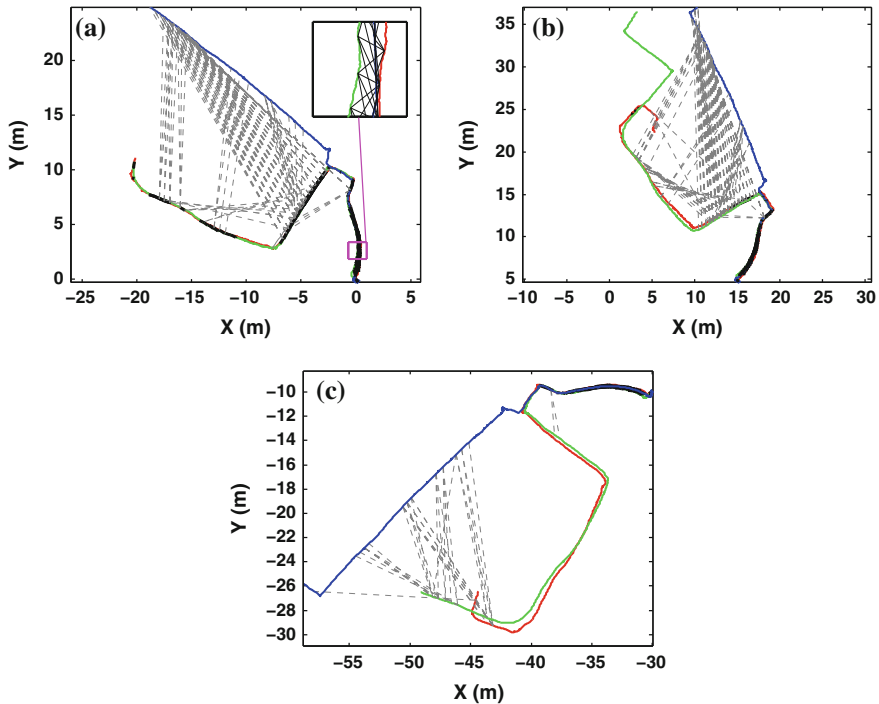


Fig. 3 Robot trajectories from **T1** in a common frame after applying transforms resulting from data association. Estimates are expressed in the local frame of each robot. Inlier (*black*) and outlier (*gray dashed*) correspondences after a common reference frame has been established are shown. **a** r_1 local frame (*red*). **b** r_2 local frame (*green*). **c** r_3 local frame (*blue*)

common. As such, strong hypotheses are formed in the initial stretch of the trajectory by the high number of inlier ICP correspondences. In the latter half of each robot’s trajectory, incorrect data associations are made between robots due to the similarity in laser features throughout the hallway environment. However, given the strength of the transform established in the initial poses, these data associations are considered outliers.

Table 2 displays computed and measured transforms from all three datasets with no laser scan saliency thresholding. The transform with the largest error among all trials, $T_{r_3}^{r_1}$ from **T1**, has a translation error norm of 0.31 m. Rotational errors in all computed transforms are between zero and eight degrees.

Figure 4a displays saliency plotted across trajectories in **T2**. Structural symmetries in the hallway environment cause ICP comparisons to converge to local minima during autocovariance computation, resulting in low saliency in areas such as four-way crossroads and corridors. The set of salient laser scans is roughly a subset of the set of shared scans that form high numbers of ICP correspondences with other robots (Fig. 4b). This inclusion signifies that autocovariance is a suitable metric for laser scan subsampling.

Table 2 Computed and measured transformations $T_{r_2}^{r_1}$ and $T_{r_3}^{r_1}$, with x, y in m and θ in rad

Transform		Trial T1		Trial T2		Trial T3	
		$T_{r_2}^{r_1}$	$T_{r_3}^{r_1}$	$T_{r_2}^{r_1}$	$T_{r_3}^{r_1}$	$T_{r_2}^{r_1}$	$T_{r_3}^{r_1}$
Computed	x:	-0.12	0.15	2.62	-4.53	1.41	-13.59
	y:	-0.03	-0.27	7.45	-4.09	-3.99	-1.24
	θ :	-0.02	0.03	-1.57	0.00	0.97	2.05
Measured	x:	0.00	0.00	2.48	-4.60	1.42	-13.63
	y:	0.00	0.00	7.50	-3.99	-3.90	-1.02
	θ :	0.00	0.00	-1.57	0.00	1.08	2.01

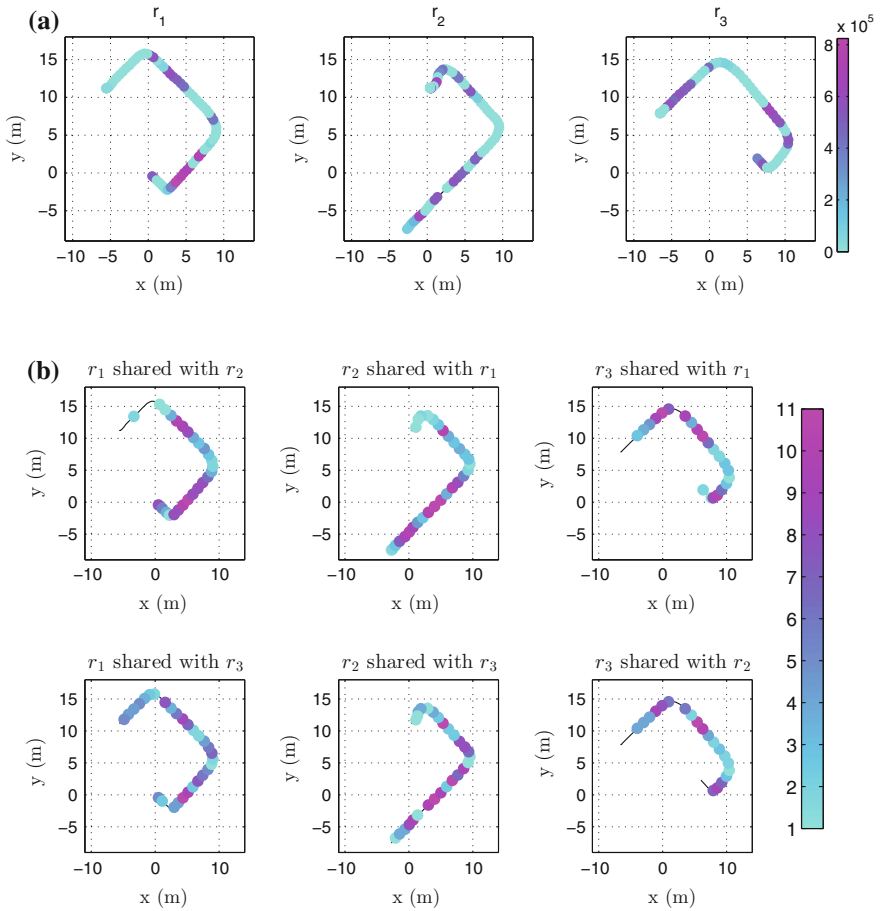


Fig. 4 **a** Saliency, δ , of shared scans in T2. **b** Number of ICP scan correspondences found between robot pairs in T2, plotted along robot trajectories. Counts mark the number of ICP correspondences made between the scan shared from the marked pose with all scans shared by the other robot

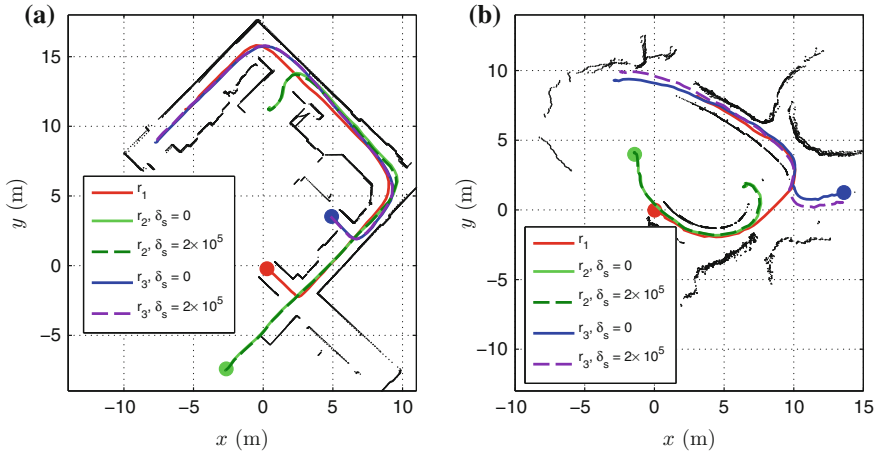


Fig. 5 **a, b** Robot trajectories from **T2** and **T3** in the frame of r_1 after applying transforms resulting from the data association strategy with and without saliency thresholding. Saliency thresholding makes little difference to transform estimates. Laser scans from r_1 are shown

Transform accuracies for **T2** and **T3** are shown in Table 3 with varying saliency thresholds, δ_s . Thresholding with $\delta_s = 2 \times 10^5$ reduces the average number of scans used to 36 percent of the total. By increasing the saliency threshold to $\delta_s = 4 \times 10^5$, there are no longer enough inlier ICP correspondences between robots r_1 and r_2 from **T2** to compute a transform. While thresholding scans by their saliency increases error in both translation and rotation, this increase, at worst, raises normed (x, y) error by 0.37 m, and rotation error by 0.08 radians over a trajectory roughly 20 m in length. At $\delta_s = 8 \times 10^5$, no robot pairs have enough inlier correspondences to form a transform.

Trajectories from **T2** and **T3** are shown in Fig. 5 with and without salient scans thresholded at $\delta_s = 2 \times 10^5$. Thresholding laser scans from robots in **T2** by saliency leads to a 0.04 m increase in translation error for both robots, with no difference in rotation error. In trial **T3**, thresholding increases translation error by 0.13 and 0.37 m, and rotation error by -0.02 and 0.08 radians for $T_{r_2}^{r_1}$ and $T_{r_3}^{r_1}$, respectively.

Without saliency thresholding, the total number of shared scans in all trials remains below the data sharing limit of 4 Hz described in Sect. 3.4 for three robots (Table 4). During trials, laser scans were captured on each robot at 20 Hz and uniformly subsampled to one tenth of the original amount. Mean sharing frequency across all robots with no laser scan subsampling remains within a small margin of 2 Hz, implying no robots were limited by the 0.1 m distance constraint between sequential laser scans. After thresholding shared scans by saliency with $\delta_s = 2 \times 10^5$, the mean sharing frequency decreases to below 1 Hz for all robots. The maximum sharing frequencies for both values of δ_s permit unconstrained data transfer on a mesh network with three robots. With a mesh network of six robots, the maximum sharing frequencies indicate that durations of increased sharing would exceed the network transfer limit of 1 Hz per robot by up to 208 percent for $\delta_s = 0$ and up to 131 percent for $\delta_s = 2 \times 10^5$.

Table 4 Trial durations, number of shared scans, and mean and maximum sharing frequencies for $\delta_s = 0$ and $\delta_s = 2 \times 10^5$

Robot	Duration (s)	$\delta_s = 0$			$\delta_s = 2 \times 10^5$		
		$ Z^r $	Max (Hz)	Mean (Hz)	$ Z^r $	Max (Hz)	Mean (Hz)
T2: r_1	37.4	75	2.08	2.00	22	1.01	0.59
T2: r_2	39.0	77	2.02	1.97	26	1.28	0.67
T2: r_3	32.5	65	2.02	2.00	23	0.95	0.71
T3: r_1	35.5	71	2.06	2.00	26	0.98	0.73
T3: r_2	27.6	55	1.99	1.99	18	1.31	0.65
T3: r_3	37.4	74	1.98	1.98	36	1.20	0.96

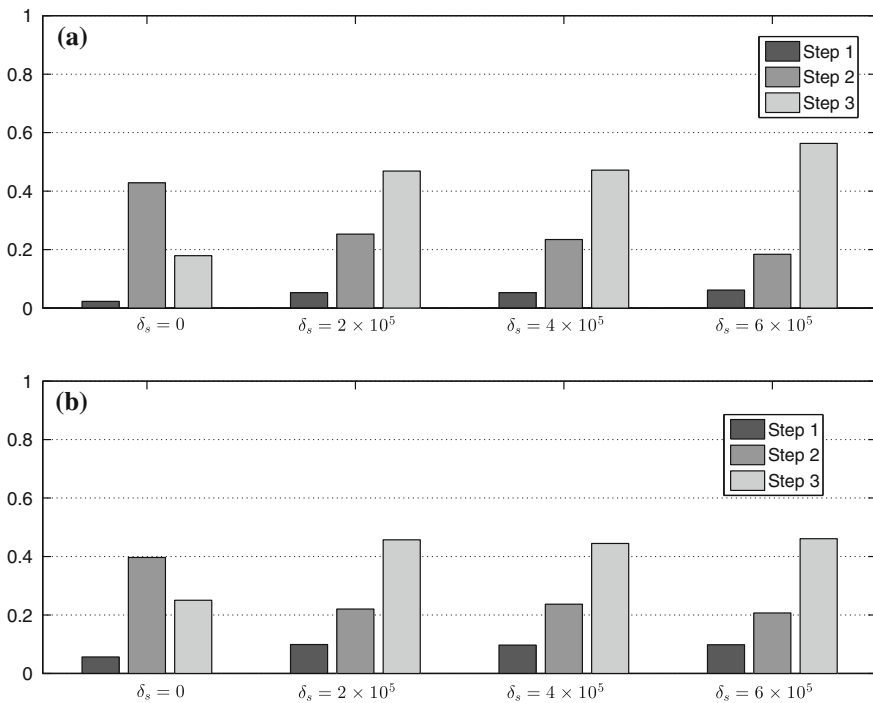


Fig. 6 Proportion of total run time consumed by each step of the algorithm for **a** trial **T2**, and **b** trial **T3** with varying values of δ_s . Time consumed by Step 3 is constant across values of δ_s , signifying that as δ_s increases the total run time of the algorithm decreases

Figure 6 shows the proportion of the process run time consumed by each step of the data association algorithm with $\delta_s = \{0, 2, 4, 6\} \times 10^5$ for trials **T2** and **T3**. After subsampling scans with $\delta_s = 2 \times 10^5$, Step 3 (autocovariance calculations for all laser scans) increases from 21.5 percent of the total computation time (averaged across

both trials) to 46.3 percent. Because total time devoted to Step 3 is not a function of δ_s , thresholding on saliency by $\delta_s = 2 \times 10^5$ decreases the total run time of the algorithm by 46.4 percent. Note that due to the quadratic complexity of Step 2, this estimate is only reflective of trajectories of the same length as those in **T2** and **T3**.

5 Conclusion

In this work we investigated the problem of computing relative transformations between multiple vehicles from shared sensor observation correspondences. Experiments were developed to evaluate the accuracy, computational complexity, and network complexity of a data association strategy introduced by Indelman, et al. [12]. A sensory observation subsampling strategy based on laser scan autocovariance was introduced to reduce the number of laser scans shared between robots, therefore reducing the both the computational complexity of the algorithm as well as the rate of data sharing between robots.

Multi-robot trials were collected onboard quadrotors operating throughout three dissimilar environments. The trials were used to evaluate the accuracy of the data association algorithm in different domains of operation. We showed that by intelligently selecting which laser scans to share over the network, the total run time of the algorithm could be reduced by 46.4% for three-robot trajectories on the order of 20 m in length. In addition, we demonstrated that the accuracy of transforms resulting from the algorithm does not suffer significantly from subsampling laser scans based on their autocovariance. Finally, we experimentally evaluated data sharing rates to show that this approach can be implemented using a mesh network for at least three robots.

References

1. Fenwick, J.W., Newman, P.M., Leonard, J.J.: Cooperative concurrent mapping and localization. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp. 1810–1817, Washington (2002)
2. Bailey, T., Bryson, M., Hua, M., Vial, J., McCalman, L., Durrant-Whyte, H.: Decentralised cooperative localisation for heterogeneous teams of mobile robots. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp. 2859–2865, Shanghai (2011)
3. Montijano, E., Aragues, R., Sagues, C.: Distributed data association in robotic networks with cameras and limited communications. *IEEE Trans. Robot.* **29**, 1408–1423 (2013)
4. Olson, E., Leonard, J., Teller, S.: Robust range-only beacon localization. In: Auton, I.E.E.E. (ed.) *Underwater Vehicles*, pp. 66–75. Maine, Sebasco (2004)
5. Howard, A., Parker, L.E., Sukhatme, G.S.: Experiments with a large heterogeneous mobile robot team: exploration, mapping, deployment and detection. *Int. J. Robot. Res.* **25**, 431–447 (2006)
6. Zhou, X.S., Roumeliotis, S.I.: Multi-robot SLAM with unknown initial correspondence: the robot rendezvous case. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 1785–1792, Beijing (2006)

7. Charrow, B., Michael, N., Kumar, V.: Cooperative multi-robot estimation and control for radio source localization. In: Proceedings of the International Symposium on Experimental Robotics, Quebec (2012)
8. Thrun, S., Thayer, S., Whittaker, W., Baker, C., Burgard, W., Ferguson, D., Hahnel, D., Montemerlo, M., Morris, A., Omohundro, Z., Reverte, C., Whittaker, W.: Autonomous exploration and mapping of abandoned mines. *IEEE Robot. Autom. Mag.* **11**, 79–91 (2004)
9. Michael, N., Shen, S., Mohta, K., Mulgaonkar, Y., Kumar, V., Nagatani, K., Okada, Y., Kiribayashi, S., Otake, K., Yoshida, K., Ohno, K., Takeuchi, E., Tadokoro, S.: Collaborative mapping of an earthquake-damaged building via ground and aerial robots. *J. Field Robot.* **29**, 832–841 (2012)
10. Cunningham, A., Wurm, K.M., Burgard, W., Dellaert, F.: Fully distributed scalable smoothing and mapping with robust multi-robot data association. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp. 1093–1100, Saint Paul (2012)
11. Montijano, E., Martinez, S., Sagues, C.: Distributed robust data fusion based on dynamic voting. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp. 5893–5898, Shanghai (2011)
12. Indelman, V., Keyes, D., Nelson, E., Michael, N., Dellaert, F.: Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization. In: Proceedings of the IEEE International Conference on Robotics and Automation (To Appear), Hong Kong (2014)
13. Lu, F., Milius, E.: Robot pose estimation in unknown environments by matching 2D range scans. *J. Intell. Robot. Syst.* **18**, 249–275 (1997)
14. Shen, S., Michael, N., Kumar, V.: Autonomous multi-floor indoor navigation with a computationally constrained MAV. In: Proceedings of the IEEE International Conference on Robotics and Automation, pp. 20–25, Shanghai (2011)
15. Pomerleau, F., Colas, F., Siegwart, R., Magnenat, S.: Comparing ICP variants on real-world data sets. *Auton. Robots* **34**, 133–148 (2013)
16. Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics. The MIT Press, Cambridge (2005)
17. Hornung, A., Wurm, K.M., Bennewitz, M., Stachniss, C., Burgard, W.: OctoMap: an efficient probabilistic 3D mapping framework based on octrees. *Auton. Robots* **34**, 189–206 (2013)
18. Censi, A.: An accurate closed-form estimate of ICP’s covariance. In: Proceedings of the IEEE International Conference on Robotics and Automation, IEEE, pp. 3167–3172 (2007)
19. Dellaert, F.: Factor graphs and GTSAM: a hands-on introduction. Technical Report GT-RIM-CP&R-2012-002 (2012)
20. Nieto, J., Bailey, T., Nebot, E.: Recursive scan-matching slam. *Robot. Auton. Syst.* **55**, 39–49 (2007)
21. Jun, J., Sichitiu, M.L.: The nominal capacity of wireless mesh networks. *IEEE Wirel. Commun.* **10**, 8–14 (2003)