Multiple Relative Pose Graphs for Robust Cooperative Mapping

Been Kim, Michael Kaess, Luke Fletcher, John Leonard, Abraham Bachrach, Nicholas Roy and Seth Teller

Abstract— This paper describes a new algorithm for cooperative and persistent simultaneous localization and mapping (SLAM) using multiple robots. Recent pose graph representations have proven very successful for single robot mapping and localization. Among these methods, incremental smoothing and mapping (iSAM) gives an exact incremental solution to the SLAM problem by solving a full nonlinear optimization problem in real-time. In this paper, we present a novel extension to iSAM to facilitate online multi-robot mapping based on multiple pose graphs. Our main contribution is a relative formulation of the relationship between multiple pose graphs that avoids the initialization problem and leads to an efficient solution when compared to a completely global formulation. The relative pose graphs are optimized together to provide a globally consistent multi-robot solution. Efficient access to covariances at any time for relative parameters is provided through iSAM, facilitating data association and loop closing. The performance of the technique is illustrated on various data sets including a publicly available multi-robot data set. Further evaluation is performed in a collaborative helicopter and ground robot experiment.

I. INTRODUCTION

This paper describes a novel extension to incremental smoothing and mapping (iSAM) [1] that facilitates multi-robot mapping based on multiple pose graphs. The goal is a scalable, general purpose technique that can apply to multiple mobile robots and/or multiple missions executed in the same environment simultaneously or at different times. The approach combines inter-robot constraints, which arise when one robot encounters another robot, or two robots observe the same part of the environment, with conventional single robot pose graphs in a manner that maintains computational efficiency and preserves consistency, using the concept of "base nodes" [2]. This provides advances over previous work, such as Howard [3] and Andersson [4] in several ways. Our approach deals with uncertainty in inter-robot constraints, and handles multiple encounters between robots. In addition, our approach is fully recursive and provides a globally consistent solution for all robots at any time. With consistent solution we mean that the correct least squares estimate based on all measurements is obtained.

SLAM techniques that use pose graphs (also known as view-based representations) [1], [5]–[10] are attractive for their robustness and computational efficiency. These methods have demonstrated successful performance on a variety of challenging large-scale vision [5], [6], [9] and lidar [1], [7], [8], [10] data sets. These methods efficiently solve the

The authors are with the Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology (MIT), Cambridge, MA 02139, USA {beenie22, kaess, lukesf, jleonard, abachrac, nickroy, teller}@mit.edu

Fig. 1: Cooperative mapping using a quadrotor and a ground robot. The quadrotor can access spaces which the ground robot cannot enter. Both robots are equipped with Hokuyo laser range finders with 30 meter range and 270 degrees viewing angle. Encounters are generated using a camera on the quadrotor to detect the checkerboard pattern on the ground robot and are refined by scan matching. The experiment was performed in the loading dock area of the Stata Center. We show an enlarged part of the map at the time of an encounter, overlaid on a floorplan. Data from the ground robot is shown in blue, for the quadrotor in red, and the encounters are shown in green. See Fig. 8 for the full map and more details.

core batch/recursive SLAM state estimation problem on a graph of constraints derived from sensor observations and proprioceptive measurements. The motivation for our work is to extend these pose graph SLAM techniques so that they may be applied in a distributed, multi-robot setting.

Multi-robot mapping introduces several additional challenges beyond those encountered in single robot SLAM. Measurements between robots introduce additional constraints, which we refer to as encounters [3], see Fig. 1 for an example. Direct encounters occur when one robot observes another robot, providing a constraint on the relative pose of the robots at a single time step. An indirect encounter occurs when two robots observe the same part of an environment (not necessarily at the same time), allowing a constraint to be estimated between the positions of two robots at the respective time steps. Initialization is a key issue – when another robot is first encountered, how is the initial

estimate for the robots' relative pose estimated? Multi-robot mapping also raises the issues of consistency, computational efficiency, and communications requirements. It is important to ensure that measurements are not utilized more than once, which would lead to overconfidence, as information flows through the distributed pose graph network. State-of-the-art algorithms exploit sparsity to efficiently handle large numbers of constraints. This paper does not consider the potential impact of constrained communications between robots, and assumes full, high-bandwidth connectivity between robots when they encounter one another.

While there is a substantial body of work on cooperative SLAM by multiple robots [3], [11]–[17], there has been limited work on cooperative SLAM using a pose graph approach. Thrun and Liu [14] formulated the multi-robot SLAM problem using the sparse extended information filter. This work highlighted the importance of the initialization problem – determining the relative pose of one robot to another at the first encounter. Whereas most early work in multi-robot SLAM assumed that initial relative poses of the robots are known [12], [18], Thrun and Liu handled initialization within the SEIF framework.

Howard [3] developed a particle filter based SLAM algorithm with experimental results for cooperative mapping by four robots using laser sensing. The approach assumes that the first encounter between two robots is perfectly accurate, making it unnecessary to consider subsequent encounters. A key motivation for our work is to relax this assumption, allowing data to be shared for multiple, uncertain encounters, converging towards the optimal solution over time, sacrificing neither consistency nor computational efficiency.

Our work uses incremental smoothing and mapping (iSAM) [1] as the core state estimation engine. iSAM extends the original batch square root SAM work by Dellaert [19] to provide an efficient incremental solution to the full SLAM problem (solving for the full vehicle trajectory, without filtering) based on updating a factorization of the information matrix. Based on this square root information matrix, iSAM also provides efficient access to the exact marginal covariances needed for data association and loop closing [20].

In the literature on smoothing and mapping (SAM) approaches, two works stand out as most relevant to this paper, Tectonic SAM [2] and C-SAM [4]. Firstly, Ni *et al.* [2] developed Tectonic SAM, a large-scale mapping approach based on submap decomposition. Tectonic SAM uses "base nodes" to represent the relative pose of each submap, providing a computational speedup as submaps can be efficiently separated from one another. The base nodes in Tectonic SAM are mathematically equivalent to the "anchor nodes" we use in this paper (described below), however the context in which they are applied is very different. Tectonic SAM addresses batch, single robot, large-scale mapping using submaps, whereas our focus is on online multi-robot SLAM in a common global frame without separately optimizing submaps and a separator. Also, in contrast to Tectonic SAM we do not make any approximations related to the separator, but rather perform one large incremental optimization over

all pose graphs. Secondly, Andersson *et al.* [4] developed a multi-robot version of SAM called C-SAM. The C-SAM algorithm has some similarities with our approach, but several crucial differences: (a) the focus of C-SAM is on the alignment and merging of multiple feature-based landmark maps [21], whereas our approach is based on pose graph representations and (b) C-SAM is a batch algorithm, whereas our algorithm is recursive. Further, our algorithm has been extensively tested with real data whereas published results for C-SAM only cover simulations with two robots.

Unlike Tectonic SAM and C-SAM, our method provides a recursive solution, while still providing a globally consistent solution for all robot trajectories and the relative parameters in real time. The relative parameters are optimized together with the pose graphs by solving a nonlinear least squares problem based on a QR factorization [1], providing the exact solution for all robot trajectories in the global frame. The uncertainty of the relative parameters can also be calculated efficiently. The performance of the technique is illustrated with a variety of data sets, including a publicly available multi-robot laser data set provided by Howard [3] and a laser data set acquired jointly by a helicopter and a ground robot, using vision for inter-robot observations.

II. APPROACH

We present a probabilistic formulation of the multi-robot SLAM problem based on pose graphs. Pose graphs are a common solution for single robot localization and mapping, in which all current and past robot poses form a Markov chain connected by odometry measurements, with additional constraints between arbitrary poses for loop closing events. Solving a pose graph formulation is also known as smoothing, and in contrast to filtering allows one to correctly deal with nonlinear measurements. Several algorithms for solving the pose graph in a recursive formulation have been presented. We choose to use incremental smoothing and mapping (iSAM) [1], because it provides an efficient solution without need for approximations, and allows efficient access to the estimation uncertainties. We first extend iSAM to deal with multiple robot trajectories simultaneously in a global reference frame. We then present a relative formulation that is advantageous in terms of faster convergence speed, while still providing a globally consistent solution.

### A. Multiple Pose Graphs in a Global Frame

We first formulate the multi-robot mapping problem in a single global coordinate frame using one pose graph for each robot trajectory. Fig. 2a shows an example for two robots, with the Markov chains for each robot clearly visible. The pose variables are shown as blue shaded circles, and measurements as small black discs. For R robots, the trajectory of robot $r \in \{0 \ldots R-1\}$ is given by $M_r + 1$ pose variables $\{\mathbf{x}_i^r\}_{i=0}^{M_r}$. As each trajectory by itself is under-constrained, we fix that gauge freedom by introducing a prior $\mathbf{p}^r$ for each trajectory r. The prior can be chosen arbitrarily, but for simplicity is chosen to be the origin. Measurements between poses of a single trajectory are of one of two types: The most

(a) Two pose graphs for two robots without encounters. Each trajectory is anchored by a prior p^r on its first pose that can be chosen arbitrarily (typically chosen to be the origin). For simplicity, we omit the concept of loop closing constraints here.

(b) Two encounters expressed as additional constraints connecting the pose graphs of the two robots. Note that encounters can be based on one robot observing the other (yielding a synchronized constraint such as c_1), or more generally by a common observation of the environment detected for example by scan matching (yielding a constraint that connects poses created at arbitrary times).

(c) The same encounters as in (b), but using our relative formulation. We introduce anchors Δ^r for each trajectory that specify the offset of the trajectory with respect to a common global frame. Each encounter measurement now additionally connects to the anchors of both trajectories.

(d) Our relative pose graph formulation generalizes to more than two robots, and does not require the pose graphs to be synchronized or even to start at the same time or location. This example shows three encounters between three robots.

Fig. 2: From single robot pose graphs to multiple relative pose graphs. In contrast to typical SLAM, pose constraints in our formulation can influence more than two variables. For visualization, we therefore use the factor graph representation, which directly corresponds to the entries in the measurement Jacobian as described in [22]. Small black discs represent measurements (factors), and larger blue shaded circles represent variables. The lines indicate dependencies, where black lines are used for normal pose constraints and red lines for encounters.

frequent type of measurement connects successive poses, and is based on odometry measurements or scan matching. Another type of measurement connects two arbitrary poses, representing loop closing constraints. Although supported in our implementation, for simplicity we omit loop closing constraints from the presentation in this paper and show only constraints $\mathbf{u}_i^r$ between successive poses $\mathbf{x}_{i-1}^r$ and $\mathbf{x}_i^r$ in the example in Fig. 2a.

In our discussion so far the robot trajectories are completely independent, which we now change by introducing the concept of encounters. An encounter $\mathbf{c}$ between two robots r and r' is a measurement that connects two robot poses $\mathbf{x}_i^r$ and $\mathbf{x}_{i'}^{r'}$. An example is shown in Fig. 2b, with the dependencies between measurement and poses shown as red lines. Note that we do not make any assumptions about synchronization here. In particular, measurement $\mathbf{c}_1$ in Fig. 2b connects poses taken at the same time, which can for example arise from an observation of one robot by the other. But more generally, an observation can connect poses taken at different times, and can for example be derived from scan matching of observations taken at arbitrary times. An example is the measurement $\mathbf{c}_2$ in Fig. 2b.

We take a probabilistic approach for estimating the actual robot trajectories based on all measurements. The joint probability of all pose variables $X = \{\mathbf{x}_i^r\}_{r=0,i=0}^{R-1,M_r}$, measurements and priors $Z = \{\mathbf{u}_i^r\}_{r=0,i=1}^{R-1,M_r} \cup \{\mathbf{p}^r\}_{r=0}^{R-1}$ and N encounters $C = \{\mathbf{c}_j\}_{j=1}^N$ is given by

$$P(X,Z,C) \propto$$
$$\prod_{r=0}^{R-1} \left(P(\mathbf{x}_0^r|\mathbf{p}^r) \prod_{i=1}^{M_r} P(\mathbf{x}_i^r|\mathbf{x}_{i-1}^r,\mathbf{u}_i) \right) \prod_{j=1}^N P(\mathbf{x}_{i_j}^{r_j}|\mathbf{x}_{i'_j}^{r'_j},\mathbf{c}_j) \quad (1)$$

where the data association (i_j, i'_j, r_j, r'_j) for the encounters is known. We assume Gaussian measurement models, as is standard in the SLAM literature. The process model

$$\mathbf{x}_i^r = f_i(\mathbf{x}_{i-1}^r, \mathbf{u}_i^r) + \mathbf{w}_i^r \quad (2)$$

describes the odometry sensor or scan-matching process, where $\mathbf{w}_i^r$ is a normally distributed zero-mean process noise with covariance matrix Λ_i^r. The Gaussian measurement equation

$$\mathbf{x}_{i_j}^{r_j} = h_j(\mathbf{x}_{i'_j}^{r'_j},\mathbf{c}_j) + \mathbf{v}_j \quad (3)$$

models the encounters between robots, where $\mathbf{v}_j$ is a normally distributed zero-mean measurement noise with covariance Γ_j.

We solve the least squares formulation that corresponds to the maximum a posteriori (MAP) estimate X^* for the robot trajectories X in (1)

$$X^* = \arg\max_X P(X,Z,C) = \arg\min_X -\log P(X,Z,C) \quad (4)$$

Combining (4) with the measurement models (2) and (3), and choosing covariance Σ for the priors, leads to the following nonlinear least squares problem:

$$X^* = \arg\min_X \left\{ \sum_{r=0}^{R-1} \left(\|\mathbf{p}^r - \mathbf{x}_0^r\|_\Sigma^2 \right. \right.$$
$$\left. + \sum_{i=1}^{M_r} \left\| f_i(\mathbf{x}_{i-1}^r, \mathbf{u}_i^r) - \mathbf{x}_i^r \right\|_{\Lambda_i^r}^2 \right)$$
$$+ \sum_{j=1}^N \left\| h_j(\mathbf{x}_{i'_j}^{r'_j}, \mathbf{c}_j) - \mathbf{x}_{i_j}^{r_j} \right\|_{\Gamma_j}^2 \right\} \quad (5)$$

where we use the notation $\|\mathbf{e}\|_\Sigma^2 = \mathbf{e}^T \Sigma^{-1} \mathbf{e}$ for the squared Mahalanobis distance with covariance matrix Σ.

We solve the non-linear least squares problem (5) using the incremental smoothing and mapping (iSAM) algorithm [1]. If the process models and constraint functions are nonlinear and a good linearization point is not available, nonlinear optimization methods are used, such as the Gauss-Newton or the Levenberg-Marquardt algorithm, which solve a succession of linear approximations to (5) to approach the minimum. Linearization of measurement equations and subsequent collection of all components in one large linear system yields a standard least squares problem of the form

$$\theta^* = \arg\min_\theta \|A\theta - \mathbf{b}\|^2 \tag{6}$$

where the vector $\theta \in \mathbb{R}^n$ contains the poses of all robot trajectories, where n is the number of variables. The matrix $A \in \mathbb{R}^{m \times n}$ is a large, but sparse measurement Jacobian, with m the number of measurements, and $\mathbf{b} \in \mathbb{R}^m$ is the right-hand side vector. iSAM solves this equation by QR factorization of the measurement Jacobian

$$A = Q \begin{bmatrix} R \\ 0 \end{bmatrix} \tag{7}$$

where $R \in \mathbb{R}^{n \times n}$ is the upper triangular square root information matrix (note that the information matrix is given by $R^T R = A^T A$) and $Q \in \mathbb{R}^{m \times m}$ is an orthogonal matrix, that is not explicitly stored in practice. Instead, the right-hand side vector $\mathbf{b}$ is modified accordingly during the QR factorization to obtain $\mathbf{d} \in \mathbb{R}^n$. The solution is then obtained by backsubstitution

$$R\theta = \mathbf{d} \tag{8}$$

To avoid refactoring an increasingly large measurement Jacobian each time a new measurement arrives, iSAM modifies the existing factorization to include the new measurement rows. The key to efficiency is to keep the square root information matrix sparse, which requires choosing a suitable variable ordering. iSAM periodically reorders the variables according to some heuristic and performs a batch factorization that also includes a relinearization of the measurement equations. More details can be found in [1].

B. Multiple Relative Pose Graphs

Initialization is a major problem of the approach to multi-robot mapping described so far. If we do not have an encounter at the beginning of the sequence, such as c_1 in the example in Fig. 2b, then we have a gauge freedom on the second trajectory. The gauge freedom requires us to also anchor the second robot trajectory using a prior. However, we do not have a good initial estimate for that prior before we obtain a first encounter. Therefore, any choice is likely to conflict with the first encounter once that is added. Removing the prior at that point is also problematic, as the variable estimates can be far from a good linearization point, with the likely consequence of getting stuck in a local minimum.

To solve the initialization problem, we introduce the concept of an anchor. The anchor Δ^r for robot trajectory r specifies the offset of the complete trajectory with respect to a global coordinate frame. That is, we keep the individual pose graphs in their own local frame, which is typically anchored at the origin with a prior, as shown in Fig. 2c. Poses are transformed to the global frame by $\Delta^r \oplus \mathbf{x}_i^r$, where we use the notation $\oplus$ from Lu and Milios [23] for pose composition. When operating in the 2D plane, composing a pose $\mathbf{b}$ with a difference (odometry) $\mathbf{d}$ is defined as

$$\mathbf{p} = \mathbf{b} \oplus \mathbf{d} = \begin{pmatrix} x_b + x_d \cos\theta_b - y_d \sin\theta_b \\ y_b + x_d \sin\theta_b + y_d \cos\theta_b \\ t_b + t_d \end{pmatrix} \tag{9}$$

and the reverse transformation $\ominus$ expresses $\mathbf{p}$ locally in the frame of $\mathbf{b}$ (odometry from $\mathbf{b}$ to $\mathbf{p}$)

$$\mathbf{d} = \mathbf{p} \ominus \mathbf{b} = \begin{pmatrix} (x_p - x_b)\cos\theta_b + (y_p - y_b)\sin\theta_b \\ -(x_p - x_b)\sin\theta_b + (y_p - y_b)\cos\theta_b \\ \theta_p - \theta_b \end{pmatrix} \tag{10}$$

The formulation of encounter measurements in the relative formulation differ from the global frame, as they now involve the anchors. As shown in Fig. 2c, the encounter now references the anchor nodes of both pose graphs. That makes sense, because the encounter is a global measure between the two trajectories, but the pose variables of each trajectory are specified in the robot's own local coordinate frame. The anchor nodes are used to transform the respective poses of each pose graph into the global frame, where a comparison with the measurement becomes possible. The measurement model h is accordingly modified to also include the two anchor variables: $h_j'(\mathbf{x}_{i_j'}^{r_j'}, \mathbf{c}_j, \Delta^{r_j}, \Delta^{r_j'})$. The difference $\mathbf{c}$ between a pose $\mathbf{x}_i^r$ from trajectory r and a pose $\mathbf{x}_{i'}^{r'}$ from another trajectory r' is given by $\mathbf{c} = (\Delta^r \oplus \mathbf{x}_i^r) \ominus (\Delta^{r'} \oplus \mathbf{x}_{i'}^{r'})$. The respective term of the global formulation (5) is therefore replaced by

$$\sum_{j=1}^N \left\| h_j'(\mathbf{x}_{i_j'}^{r_j'}, \mathbf{c}_j, \Delta^{r_j}, \Delta^{r_j'}) - \mathbf{x}_{i_j}^{r_j} \right\|_{\Gamma_j}^2 \tag{11}$$

Our concept of relative pose graphs generalizes well to a larger number of robot trajectories. An example with three pose graphs is shown in Fig. 2d. Note that the number of anchor nodes depends only on the number of robot trajectories. Note however that the overall system again faces a gauge freedom that is resolved by adding a prior to one of the anchor nodes. In practice, we therefore only add anchor nodes once they are needed, and add a prior to the first one that is added.

III. EXPERIMENTS AND RESULTS

We illustrate the performance of our technique with a variety of data sets, including a publicly available multi-robot laser data set and a laser data set acquired jointly by a helicopter and a ground robot. An overview of the properties of the data sets is given in Table I. For our experiments, we use a scan matcher that remembers a short history of scans corresponding to key frames, fits contours to the data, and for a new scan, finds a globally optimal alignment within a

(a)   (b)

Fig. 3: Stata Center data set, which contains over 11000 laser scans recorded in about 40 minutes, covering an area of approximately 100m × 55m. We split the data into two parts to simulate two robots traversing the same environment. (a) The merged map based on our algorithm. The blue part of the map is created by one robot, the green part by the other. (b) Encounters visualized by plotting time as elevation of the robot trajectory. The vertical red lines show loop closing constraints within a single robot trajectory, while the vertical black lines show the encounters between the trajectories.

	Stata Center	Fort AP Hill	Loading Dock
Number of robots	2	3	2
Number of scans	11473	77478	42320
Number of poses	881	840	386
Number of encount.	43	12	7
Entries/col. in R	7.4	5.6	5.1
Mapped area	100m × 55m	40m × 28m	45m × 60m
Trajectory length	278m	398m	106m
Optimization time	2.3s	1.4s	0.3s

TABLE I: Overview of the data sets and their properties. The original number of scans is reduced by the key frame process of the scan matcher to a smaller number of poses that eventually is used in the pose graph optimization. The optimization time includes obtaining a full solution after every constraint is added.


Fig. 4: The square root information matrix at the end of the Stata Center sequence. The matrix remains sparse due to variable reordering, which is essential for an efficient solution.


Fig. 5: Positional uncertainty of the origin of the second pose graph with respect to the origin of the first one for the Stata Center sequence.

given search region. For our results, we show plots of the laser measurements instead of evidence grids, as evidence grids can sometimes hide inaccuracies of the map, such as a duplication of a wall.

A. Stata Center Data Set

We present results based on recorded data from the third floor of our building, the Stata Center. The building shape differs from traditional office buildings, yielding its own challenges for robot mapping. The data was recorded with a single robot, and is here split into two segments that are assumed to be created by independent robots. We omitted several scans around the cutting point to ensure the trajectories are independent. Encounters between robots (black lines) as well as loop closing constraints (red lines) within a trajectory are generated by scan matching. For the first encounter, we search the best matching scans. After that, an encounter is added if the robots are within 4m from each other, and more than 85% of the scans overlap. Fig. 3 shows the map created by our algorithm as well as the encounters.

Fig. 4 shows that the final square root information matrix is indeed sparse, with 19633 entries for a matrix of side length 2646, yielding 7.4 entries per column. Fig. 5 shows how the uncertainty of the origin of the second pose graph in the global reference frame (also the origin of the first

one) changes over time. Note that the anchor node is only created at the time of the first encounter (time step 160). The uncertainty remains constant until further encounters and loop closures within a single pose graph occur, for example at time step 231.

B. Fort AP Hill Data Set

We have applied our algorithm to a publicly available multi-robot data set recorded at Fort AP Hill, with the results shown in Fig. 6. This data set was obtained from the Robotics Data Set Repository (Radish) [24]. Three robots travel through the same building on different routes, covering different parts of the building, but with significant overlap,

Fig. 6: Fort AP Hill data set with three robots traveling through the same building. (a) The individual maps from each of the three robots. (b) Map based on the output of our algorithm, using a different color for each of the three robots. (c) The encounters between the three robots visualized by plotting time as elevation of the robot trajectory. The vertical black lines show encounters.


Fig. 7: Comparison of convergence speed between relative and global parameterization, showing faster convergence for our relative formulation.

as shown by the maps of the individual robots in Fig. 6a.

The data set includes encounter information based on fiducials, however, since all the robots start at the same position of the map (bottom right) there are only few encounters later, and we used scan matching instead to define encounters. Encounters are defined based on the same criteria as for the previous data set.

Our algorithm produced the combined map shown in Fig. 6b. We show a 3D view of the planar map in Fig. 6c with the encounters as vertical black lines between robot trajectories plotted by using time as elevation.

C. Loading Dock Data Set

Results for cooperative mapping between a quadrotor helicopter and a ground robot are shown in Fig. 8. This heterogeneous setup allows mapping of areas that are not fully accessible to either type of robot. While the ground robot can easily get into smaller spaces that are problematic for the helicopter, the helicopter can fly over obstacles and cross discontinuities in the floor level. The helicopter setup, control and 2D scan matching is described in our previous work [25]. We performed our experiment in the underground loading dock of the Stata Center that provides continuous hallways and spaces for the ground robot, and a large lower level that is accessible to the helicopter.

Encounters are found using vision to determine inter-robot observations. A checker board on the ground robot was detected from a camera on the helicopter. As the derived relative pose information is noisy, we used scan matching, both to refine the estimate, and to eliminate false positives.

The individual maps generated from the ground robot and the helicopter are shown in Fig. 8a. The complete map shown in Fig. 8b overlaid on a floor plan obtained from MIT facilities shows strong correlation between the learned map and ground truth. This paper is published together with a video of the actual experiment and the map building process.

We compare convergence speed between the relative and the global parameterization in Fig. 7. The results show that the relative formulation requires less iterations until the residual falls below a threshold (here 10^{-4}). For iSAM we performed batch optimization every 100 steps, and in most cases a single iteration (reordering and relinearization) was sufficient for the relative formulation, while the global formulation required about twice as many iterations.

IV. CONCLUSION

This paper has presented a new extension to iSAM for multi-robot mapping, with the following properties:

- The algorithm solves the full multi-robot nonlinear SLAM optimization problem in real-time.
- The algorithm provides fast convergence to the global solution, while preserving efficient access to covariances.

(a)

(b)

(c)

Fig. 8: Cooperative mapping using a quadrotor and a ground robot. The robots are shown in Fig. 1. Encounters (green) are generated using a camera on the quadrotor to detect the checkerboard pattern on the ground robot and are refined by scan matching. (a) Map obtained by ground robot (left) and quadrotor (right). The picture in the center shows the robots in operation. (b) Merged map obtained from our algorithm, overlaid on a floor plan that was provided by MIT facilities. (c) Panoramic view of the loading dock environment.

- The algorithm accommodates multiple uncertain encounters between robots.
- The algorithm preserves consistency, even in situations with cycles in the graph of encounters.

The algorithm has been validated experimentally using several real data sets, demonstrating improved convergence speed in comparison to using a global parameterization. Our experimental results include indoor cooperative mapping by a ground robot and a helicopter using visual encounters.

Our relative formulation avoids the need to adjust individual pose graphs when they are combined. This is achieved using anchor nodes, which are mathematically equivalent to the "base nodes" first introduced by Ni *et al.* in the different context of large-scale batch submapping with the Tectonic SAM algorithm [2]. Using these anchor nodes, as soon as robots encounter one another, all the past data from each robot is available to other robots immediately in a common reference frame; there is no need to adjust each pose graph or to reprocess previous data that another robot has acquired. When the first encounter occurs, the pose graphs for the respective robots are not modified at all and only the anchor nodes change; when subsequent encounters occur, information can propagate between the two pose graphs similar to the single robot loop closing case.

Our approach does not at present take into account any issues of communication bandwidth constraints between robots. Nerurkar *et al.* [26] investigated methods for managing communication cost and limiting complexity for large multi-robot teams performing cooperative localization using a distributed conjugate gradient optimization. In future work, it would be interesting to explore the impact of communication constraints on the algorithm presented in this paper.

The technique described here is anticipated to be beneficial for long-term persistent mapping, enabling the accumulation of large databases of relative pose graphs acquired over a robot's lifetime. This will require the development of appropriate techniques for dealing with long-term changes in the environment. Another challenge for future research is to incorporate partial constraints, such as range only or bearing only measurements, for the robot-to-robot encounters.

ACKNOWLEDGMENTS

REFERENCES

[1] M. Kaess, A. Ranganathan, and F. Dellaert, "iSAM: Incremental smoothing and mapping," *IEEE Trans. Robotics*, vol. 24, pp. 1365–1378, Dec 2008.

[2] K. Ni, D. Steedly, and F. Dellaert, "Tectonic SAM: Exact, out-of-core, submap-based SLAM," in *Proc. IEEE Intl. Conf. Robotics and Automation*, pp. 1678–1685, Apr 2007.

[3] A. Howard, "Multi-robot simultaneous localization and mapping using particle filters," *Intl. J. of Robotics Research*, vol. 25, no. 12, pp. 1243–1256, 2006.

[4] L. A. A. Andersson and J. Nygards, "C-SAM : Multi-robot SLAM using square root information smoothing," in *Proc. IEEE Intl. Conf. Robotics and Automation*, pp. 2798–2805, 2008.

[5] K. Konolige and M. Agrawal, "FrameSLAM: From bundle adjustment to real-time visual mapping," *IEEE Trans. Robotics*, vol. 24, no. 5, pp. 1066–1077, 2008.

[6] I. Mahon, S. Williams, O. Pizarro, and M. Johnson-Roberson, "Efficient view-based SLAM using visual loop closures," *IEEE Trans. Robotics*, vol. 24, pp. 1002–1014, Oct 2008.

[7] J. Folkesson and H. Christensen, "Graphical SLAM - a self-correcting map," in *Proc. IEEE Intl. Conf. Robotics and Automation*, vol. 1, pp. 383–390, 2004.

[8] U. Frese, "Treemap: An $O(\log n)$ algorithm for indoor simultaneous localization and mapping," *Autonomous Robots*, vol. 21, no. 2, pp. 103–122, 2006.

[9] R. Eustice, H. Singh, and J. Leonard, "Exactly sparse delayed-state filters for view-based SLAM," *IEEE Trans. Robotics*, vol. 22, pp. 1100–1114, Dec 2006.

[10] E. Olson, J. Leonard, and S. Teller, "Spatially-adaptive learning rates for online incremental SLAM," in *Robotics: Science and Systems (RSS)*, Jun 2007.

[11] S. Williams, G. Dissanayake, and H. Durrant-Whyte, "Towards multi-vehicle simultaneous localisation and mapping," in *Proc. IEEE Intl. Conf. Robotics and Automation*, pp. 2743–2748, May 2002.

[12] J. Fenwick, P. Newman, and J. Leonard, "Cooperative concurrent mapping and localization," in *Proc. IEEE Intl. Conf. Robotics and Automation*, vol. 2, pp. 1810–1817, 2002.

[13] A. Birk and S. Carpin, "Merging occupancy grid maps from multiple robots," *Proc. of the IEEE*, vol. 94, no. 7, pp. 1384–1397, 2006.

[14] S. Thrun and Y. Liu, "Multi-robot SLAM with sparse extended information filters," *Proc. of the Intl. Symp. of Robotics Research (ISRR)*, vol. 15, pp. 254–266, 2005.

[15] J. Ko, B. Stewart, D. Fox, K. Konolige, and B. Limketkai, "A practical, decision-theoretic approach to multi-robot mapping and exploration," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pp. 3232–3238, 2003.

[16] X. Zhou and S. Roumeliotis, "Multi-robot SLAM with unknown initial correspondence: The robot rendezvous case," in *2006 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pp. 1785–1792, 2006.

[17] A. Howard, G. S. Sukhatme, and M. J. Matarić, "Multi-robot mapping using manifold representations," *Proc. of the IEEE*, vol. 94, pp. 1360–1369, Jul 2006.

[18] E. W. Nettleton, H. F. Durrant-Whyte, P. W. Gibbens, and A. H. Goektogan, "Multiple-platform localization and map building," in *Proc. SPIE, Sensor Fusion and Decentralized Control in Robotic Systems III*, vol. 4196, pp. 337–347, 2000.

[19] F. Dellaert, "Square Root SAM: Simultaneous location and mapping via square root information smoothing," in *Robotics: Science and Systems (RSS)*, 2005.

[20] M. Kaess and F. Dellaert, "Covariance recovery from a square root information matrix for data association," *Journal of Robotics and Autonomous Systems, RAS*, vol. 57, pp. 1198–1210, Dec 2009.

[21] J. Tardós, J. Neira, P. Newman, and J. Leonard, "Robust mapping and localization in indoor environments using sonar data," *Int. J. Robotics Research*, vol. 21, pp. 311–330, Apr 2002.

[22] F. Dellaert and M. Kaess, "Square Root SAM: Simultaneous localization and mapping via square root information smoothing," *Intl. J. of Robotics Research*, vol. 25, pp. 1181–1203, Dec 2006.

[23] F. Lu and E. Milios, "Globally consistent range scan alignment for environment mapping," *Autonomous Robots*, pp. 333–349, Apr 1997.

[24] A. Howard and N. Roy, "The robotics data set repository (radish)," 2003.

[25] A. Bachrach, R. He, and N. Roy, "Autonomous flight in unstructured and unknown indoor environments," in *European Micro Aerial Vehicle Conference*, Sep 2009.

[26] E. Nerurkar, S. Roumeliotis, and A. Martinelli, "Distributed maximum a posteriori estimation for multi-robot cooperative localization," in *Proc. IEEE Intl. Conf. Robotics and Automation*, pp. 1402–1409, Jun 2009.

[27] E. Olson, "Real-time correlative scan matching," in *Proc. IEEE Intl. Conf. Robotics and Automation*, (Kobe, Japan), pp. 4387–4393, Jun 2009.