The 2010 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 18-22, 2010, Taipei, Taiwan

Decentralized Cooperative Simultaneous Localization and Mapping for Dynamic and Sparse Robot Networks

Keith Y. K. Leung*, Timothy D. Barfoot†, Hugh H. T. Liu‡

University of Toronto Institute for Aerospace Studies
Toronto, Ontario, Canada, M3H 5T6
keith.leung@robotics.utias.utoronto.ca*, tim.barfoot@utoronto.ca†, liu@utias.utoronto.ca‡

***Abstract*— Communication among robots is key to performance in cooperative multi-robot systems. In practice, communication connections for information exchange between all robots are not always guaranteed, which adds difficulty to state estimation. This paper examines the decentralized cooperative *simultaneous localization and mapping (SLAM)* problem under a sparsely-communicating and dynamic network. We mathematically prove how the centralized-equivalent estimate can be obtained by all robots in the network in a decentralized manner. Furthermore, a robot only needs to consider its own knowledge of the network topology to detect when the centralized-equivalent estimate is obtainable. Our approach is validated through more than 250 minutes of experiments using a team of real robots, with accurate groundtruth data of all robots and landmark features.**

I. Introduction

A cooperative multi-robot system is beneficial in many applications. It allows for the implementation complex strategies that require more than a single robot. Multiple robots can also provide a certain degree of redundancy to ensure the completion of tasks should a portion of the multi-robot team become disabled. Communication and the mutual exchange of information are key performance factors for many cooperative multi-robot systems. However, only a few number of researchers have examined the limitation of communication range and its impact on a multi-robot system (and on cooperative SLAM in particular). Leung et al. [1] presented an algorithm for performing decentralized cooperative localization in sparsely-communicating robot networks. The algorithm is versatile in that it can use different recursive filtering methods within it. They showed with a proof in [2] that the algorithm allows all robots to obtain decentralized estimates that are equivalent to the centralized estimate whenever possible, even when the communication network between robots is never fully connected. Furthermore, simulation results show how localization performance of the decentralized algorithm changes with respect to network connectivity. In this paper, we examine the more difficult cooperative decentralized *simultaneous localization and mapping (SLAM)* problem, in which each robot is required to estimate the map and the state of all robots in a sparsely-communicating and dynamic network. More specifically, we examine how robots can obtain the centralized-equivalent state estimate for the system under the assumption of a network that is never guaranteed to be fully connected. Our decentralized

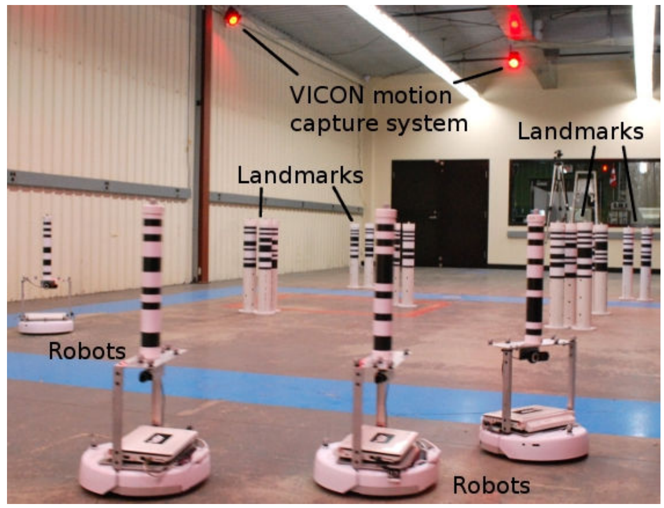

Fig. 1. Our decentralized cooperative SLAM experiment includes a fleet of five robots built using the iRobot Create platform. A 10-camera Vicon motion capture system provides accurate groundtruth information for both robots and landmarks during our experiments.

cooperative SLAM algorithm is validated through more than 250 minutes of experiments with real hardware (as shown in Fig. 1) and the results are compared with groundtruth data for both robot poses and landmark positions. To the knowledge of the authors, this is one of the few multi-robot SLAM experiments where millimetre-accurate groundtruth data is available.

We begin with a review of some of the past work in multi-robot state estimation in section II. Section III contains the problem formulation for cooperative SLAM. In section IV, we present and prove several key theorems that form the basis of our decentralized algorithm, and also discuss the necessary initial conditions for the decentralized cooperative SLAM problem. In section V, we present our decentralized algorithm. Section VI describes our experimental setup, and the experimental results are presented in section VII.

II. Related Work

The study of distributed and decentralized state estimation first started with sensor networks, and later progressed to the study of the cooperative localization problem. This problem involves a team of robots communicating relative measurements and odometry information with one another

978-1-4244-6676-4/10/$25.00 ©2010 IEEE

to determine the state of the team. Refer to [2] for a more thorough review of past work on sensor networks and cooperative localization.

Cooperative SLAM extends the idea of cooperative localization to estimate a map of the environment. Fenwick et al. [3] generalized the single robot Extended Kalman Filter (EKF) SLAM approach for the multi-robot scenario. Using the information form of the EKF and taking advantage of the additive property of updating an estimate with observation information, Nettleton et al. [4, 5] devised a decentralized SLAM algorithm that involves low-bandwidth transmission of sub-map information between robots. Reece and Roberts [6] pointed out that Nettleton's approach could yield highly conservative estimates and proposed a method to improve on this. Later, Thrun et al. [7] introduced the *Sparse Extended Information Filter (SEIF)* for SLAM, which could be extended for the multi-robot case. Howard et al. [8] presented a novel approach to multi-robot SLAM that used manifold maps. In their experiments, each robot maintained its own manifold map until it encountered another robot, at which point their maps were merged. Howard [9] also looked at performing multi-robot SLAM, wherein each robot was unaware of each other's initial pose and began state estimation in a decentralized manner. When robots encountered each other for the first time, their individual maps were combined into a common map using relative poses. The mapping process then continued as robots broadcasted new observations to one another. The notion of a virtual robot traveling backward in time was introduced to allow the incorporation of information gathered by a robot before the common map was merged. Ko et al. [10], Zhou and Roumeliotis [11] and Wang et al. [12] also examined the unknown initial correspondence problem. Zhou and Roumeliotis [11] performed experiments where robots used relative range and bearing measurements as well as landmarks in the robots' maps to create a merged map. Also related is the work by Stachniss [13] on multi-robot mapping with known poses, where robots keep track of information communicated to other robots to maintain the centralized estimate. In contrast to this, an important aspect of our current work is that robots do not need to keep track of what other robots know, and we can still guarantee that the centralized-equivalent estimate is obtainable by all robots.

In the works mentioned above, methods that can obtain the centralized estimate requires a static and fully connected network, and methods that do not rely on a static and fully connected network are sub-optimal in the sense that they can not show that the centralized-equivalent estimate can be obtained. Furthermore, in some the works mentioned, robots do not estimate the state of one another other. In this paper, we present an algorithm that will address all the above issues. That is, we will present an algorithm that is able to provide the centralized-equivalent estimate of all robots and observed landmarks whenever possible in a sparsely-communicating and dynamic network that never has to be fully connected.

III. Problem Formulation

In a multi-robot system, let N represent the set that contains the unique identification indices of all robots, and let M represent the set that contains the unique identification indices of all landmarks (which make up the map). We assume a general system model:

$$\begin{aligned}
\mathbf{x}_{i,k} &= \mathbf{g}\left(\mathbf{x}_{i,k-1}, \mathbf{u}_{i,k}, \epsilon_k\right) \ (\forall i \in N) \\
\mathbf{x}_{j,k} &= \mathbf{x}_{j,k} \ (\forall j \in M) \\
\mathbf{y}_{i,k}^{j,i} &= \mathbf{h}\left(\mathbf{x}_{i,k}, \mathbf{x}_{j,k}, \delta_k\right), (\forall j \in N \cup M)(d_k^{j,i} \leq r_{\mathsf{obs}})
\end{aligned}$$

where for timestep k, $\mathbf{x}_{i,k}$ $(i \in N)$ represents the state (pose) of robot i, $\mathbf{x}_{j,k}$ $(j \in M)$ represents the state (position) of the landmark j, $\mathbf{u}_{i,k}$ represents the odometry information of robot i, $\mathbf{g}(\cdot)$ is the state transition function (with process noise, ϵ_k), $\mathbf{y}_{i,k}^{j,i}$ represents the measurement (e.g., range/bearing) of robot j with respect to robot i, $\mathbf{h}(\cdot)$ is the measurement function (with measurement noise, δ_k), $d_k^{j,i}$ is the distance between robot i and object (robot or landmark) j, and r_{obs} is the measurement range limit. Let

$$X_k = \{\mathbf{x}_{i,k}\}, \ (\forall i \in N \cup M_k)$$

represent the set of all states at timestep k, where M_k is the set of landmarks that has been observed by at least one robot up to time k. Let

$$X_{i,k} = \{\mathbf{x}_{j,k}\}, \ (\forall j \in N_{i,k} \cup M_{i,k})$$

represent the states of all robots and landmarks known to robot i up to time k, where $N_{i,k}$ is the set of robots known to robot i up to time k, and $M_{i,k}$ is the set of landmarks known to robot i up to time k. Let

$$Y_{i,k} = \{\mathbf{y}_{i,k}^{j,i}\}, \ (\forall j)(d_k^{j,i} \leq r_{\mathsf{obs}})$$

represent the set of all measurements made by robot i at timestep k. Due to uncertainty in both state transition and measurements, the true state of the system cannot be found deterministically, but can only be estimated using odometry and measurement data. In general, the centralized *belief*, $\mathsf{bel}(X_k)$, is represented by a probability density function, $p(\cdot)$, over all states, X_k:

$$\mathsf{bel}(X_k) := p\left(X_k | \mathsf{bel}(X_0), \{\mathbf{u}_{i,1:k}\}, \{Y_{i,1:k}\}, (\forall i)\right),$$

which is conditioned on the initial belief, $\mathsf{bel}(X_0)$, past odometry data, and past range and bearing measurements.

The *knowledge set*, $S_{i,k}$, consists of all odometry and measurement data, as well as the previous state estimates known to robot i at time k. We assume initial that

$$S_{i,0} = \{\mathsf{bel}(\mathbf{x}_{i,0})\}, \ (\forall i \in N).$$

Robots within communication range r_{comm} of each other are able to exchange and relay state estimates, odometry data, and measurement data. Let $S_{i,k}^-$ represent the knowledge set after state transition and observations, but before communication is established with any other robot:

$$S_{i,k}^- = S_{i,k-1} \cup \{\mathbf{u}_{i,k}, Y_{i,k}\} \tag{1}$$

When communication occurs between robots i and j, they will make their knowledge sets available to each other, and the knowledge set of both robots will become identical:

$$S_{i,k} = S_{j,k} = S^-_{i,k} \cup S^-_{j,k}, (\forall j)(d^{j,i}_k \leq r_{\mathsf{comm}}) \tag{2}$$

From a practical and computation point of view, it is helpful to apply the *Markov property*,

$$p\left(X_k | \mathsf{bel}(X_0), U_{1:k}, Y_{1:k}\right) = p\left(X_k | \mathsf{bel}(X_{k-1}), U_k, Y_k\right),$$

when performing state estimation, as it limits memory and processing requirements and allows for recursive state estimation. The difficulty is that in a decentralized framework, the Markov property can only be applied once a robot obtains sufficient information regarding other robots through communication. Furthermore, each robot must ensure that other robots will no longer require any of the past information that will be discarded when applying the Markov property. Our objective is for each robot to perform SLAM to estimate the state of all robots and known landmarks (i.e., find $\mathsf{bel}(X_{i,k})$) in a decentralized manner, given that the robots are only occasionally exchanging information with one another.

IV. Decentralized SLAM in a Dynamic Network

In this section, we will examine the decentralized cooperative SLAM problem theoretically. We will show the conditions under which a robot can obtain the centralized-equivalent estimate and show how it relates to the decentralized cooperative localization problem. We will also discuss the initial conditions necessary to guarantee that the centralized-equivalent estimate is obtainable by all robots.

A. Obtaining the Centralized-Equivalent Estimate

The decentralized cooperative SLAM problem is very similar to the decentralized cooperative localization problem, with the difference being that in decentralized cooperative SLAM, there are landmarks positions that need to be estimated in addition to the robot states. The concepts of checkpoint and partial checkpoint were previously defined in [2] for decentralized cooperative localization. The existence of these events indicate that it is possible (in the decentralized cooperative localization problem) for the entire team of robots (and thus each single robot) to obtain the centralized-equivalent estimate at the event occurrence time. We will now examine how these events can be applied to the cooperative SLAM problem.

Definition 1: *A* checkpoint, $C(k_c, k_e)$, *is an event that occurs at the checkpoint time,* k_c, *that first comes into existence at* k_e, *in which* $S_{i,k_e} \supseteq S_{j,k_c}, \forall i \in N, \forall j \in N$.

Definition 2: *A* partial checkpoint, $C_p(k_{c,i}, k_{e,i})$, *is an event that occurs for robot* i *at time* $k_{c,i}$, *that first comes into existence at* $k_{e,i}$, *in which* $S_{i,k_{e,i}} \supseteq S_{j,k_{c,i}}, \forall j \in N$.

Using the above definitions, we want to define when it is possible to obtain the centralized-equivalent estimate in cooperative SLAM (under a sparsely-communicating network).

Lemma 1.1: *In decentralized cooperative SLAM, all robots can obtain the centralized-equivalent estimate,* $\mathsf{bel}(X_{k_c})$, *if and only if each robot's knowledge set,* $S_{i,k}$, *contains* $\{\mathsf{bel}(X_{k_c})\}$ *or* $\{\mathsf{bel}(X_{j,k_{s,j}}), \mathbf{u}_{j,k_{s,j}+1:k_c}, Y_{j,k_{s,j}+1:k_c}\}$, *where* $k_{s,j} < k_c, \forall j \in N$.

Proof. Assume that the centralized-equivalent estimate is obtainable by all robots. This implies that all robots can calculate

$$\begin{aligned}
\mathsf{bel}(X_{k_c}) &= p\left(X_{k_c} | \mathsf{bel}(\mathbf{x}_{j,0}), \{\mathbf{u}_{j,1:k_c}, Y_{j,1:k_c}\}, \forall j \in N\right) \\
&= p\Big(X_{k_c} | \mathsf{bel}(X_{j,k_{s,j}}), \left\{\mathbf{u}_{j,k_{s,j}+1:k_c}\right\}, \\
&\qquad \left\{Y_{j,k_{s,j}+1:k_c}\right\}, \forall j \in N\Big), \text{ if } k_{s,j} < k_c \\
&= p\left(X_{k_c} | \mathsf{bel}(X_{j,k_{s,j}}), \forall j \in N\right), \text{ if } k_{s,j} = k_c.
\end{aligned}$$

In order for all robots to generate the above belief, each robot must have the following information in its knowledge set:

$$\begin{aligned}
S_{i,k} &\supseteq \begin{cases} \{\mathsf{bel}(X_{j,k_{s,j}}), \mathbf{u}_{j,k_{s,j}+1:k_c}, Y_{j,k_{s,j}+1:k_c}\}, \\ \qquad \text{if } k_{s,j} < k_c \\ \{\mathsf{bel}(X_{j,k_{s,j}})\}, \text{ if } k_{s,j} = k_c \end{cases}, \forall j \in N \\
&\supseteq \begin{cases} \{\mathsf{bel}(X_{j,k_{s,j}}), \mathbf{u}_{j,k_{s,j}+1:k_c}, Y_{j,k_{s,j}+1:k_c}\}, \\ \qquad \text{if } k_{s,j} < k_c \\ \{\mathsf{bel}(X_{k_c}))\}, \text{ if } k_{s,j} = k_c \end{cases}, \forall j \in N.
\end{aligned}$$

This completes the first half of our proof because $S_{i,k}$ contains the centralized-equivalent estimate at k_c or the centralized-equivalent estimate from robot j at some earlier time $k_{s,j}$ with odometry and measurement data up to time k_c, $\forall j \in N$. Next we assume that each robot's knowledge set contains $\{\mathsf{bel}(X_{k_c})\}$ or $\{\mathsf{bel}(X_{j,k_{s,j}}), \mathbf{u}_{j,k_{s,j}+1:k_c}, Y_{j,k_{s,j}+1:k_c}\}$, where $k_{s,j} < k_c, \forall j \in N$. This allows each robot to generate a probability density function with the aforementioned information as conditional dependencies:

$$\begin{aligned}
&\begin{cases} p\left(X_{k_c} | \mathsf{bel}(X_{j,k_{s,j}}), \left\{\mathbf{u}_{j,k_{s,j}+1:k_c}\right\}, \left\{Y_{j,k_{s,j}+1:k_c}\right\}, \right. \\ \qquad\qquad \left. \forall j \in N\right), \text{ if } k_{s,j} < k_c \\ p\left(X_{k_c} | \mathsf{bel}(X_{k_c})\right), \text{ if } k_{s,j} = k_c \end{cases} \\
&= \mathsf{bel}(X_{k_c}),
\end{aligned}$$

where both densities are equivalent to the centralized estimate. This completes the second half of the proof. $\square$

Theorem 1.1: *In decentralized cooperative SLAM, a checkpoint,* $C(k_c, k_e)$, *exists if and only if all robots can obtain the centralized equivalent estimate,* $\mathsf{bel}(X_{k_c})$.

Proof. First, assume that $C(k_c, k_e)$ exists, which by definition implies

$$S_{i,k_e} \supseteq S_{j,k_c}, \forall i \in N, \forall j \in N.$$

In other words, each robot's knowledge set at time k_e must contain all robot's knowledge sets at time k_c. At a minimum, the knowledge set of robot j must contain its own state estimate, as well as odometry and measurements from its state estimate time, $k_{s,j}$, to time k_c. We can write this as

$$S_{j,k_c} \supseteq \begin{cases} \{\mathsf{bel}(X_{j,k_{s,j}}), \mathbf{u}_{j,k_s+1:k_c}, Y_{j,k_s+1:k_c}\}, \\ \qquad\qquad\qquad\qquad \text{if } k_{s,j} < k_c \\ \{\mathsf{bel}(X_{j,k_{s,j}})\}, \text{ if } k_{s,j} = k_c \end{cases}.$$

Substituting this into the definition of a checkpoint, when a checkpoint exists, we have

$$S_{i,k_e} \supseteq \begin{cases} \{\mathsf{bel}(X_{j,k_{s,j}}), \mathbf{u}_{j,k_s+1:k_c}, Y_{j,k_s+1:k_c}\}, \\ \qquad \text{if } k_{s,j} < k_c \\ \{\mathsf{bel}(X_{k_c}))\}, \text{ if } k_{s,j} = k_c \end{cases} \begin{array}{r} \forall j \in N, \\ \forall i \in N. \end{array}$$

Using Lemma 1.1, this implies that all robots can obtain the centralized-equivalent estimate. For the second half of the proof, we assume that all robots can obtain the centralized-equivalent estimate. By using Lemma 1.1, we know that this implies

$$S_{i,k_e} \supseteq \begin{cases} \{\mathsf{bel}(X_{j,k_{s,j}}), \mathbf{u}_{j,k_s+1:k_c}, Y_{j,k_s+1:k_c}\}, \\ \qquad \text{if } k_{s,j} < k_c \\ \{\mathsf{bel}(X_{k_c}))\}, \text{ if } k_{s,j} = k_c \end{cases} \begin{array}{r} \forall j \in N, \\ \forall i \in N. \end{array}$$

By our knowledge set update rules, this implies

$$S_{i,k_e} \supseteq S_{j,k_c}, \forall i \in N, \forall j \in N.$$

Or in other words, $C(k_c, k_e)$ exists. □

The following lemma and theorem relate to partial checkpoints. Their proofs are similar to those in Lemma 1.1 and Theorem 1.1 and will not be shown.

Lemma 1.2: *In decentralized cooperative SLAM, robot i can obtain the centralized-equivalent estimate,* $\mathsf{bel}(X_{k_{c,i}})$*, if and only if the robot's knowledge set,* $S_{i,k}$*, contains* $\{\mathsf{bel}(X_{k_{c,i}})\}$ *or* $\{\mathsf{bel}(X_{k_{s,j}}), \mathbf{u}_{j,(k_{s,j}+1):k_{c,i}}, Y_{j,(k_{s,j}+1):k_{c,i}}\}$ *where* $k_{s,j} < k_{c,i} \forall j \in N$.

Theorem 1.2: *In decentralized cooperative SLAM, a partial checkpoint,* $C_p(k_{c,i}, k_{e,i})$*, exists for robot i if and only if robot i can obtain the centralized equivalent estimate,* $\mathsf{bel}(X_{k_{c,i}})$.

We have shown that the existence of a checkpoint (and partial checkpoint) indicates when all robots (and a single robot) can obtain the centralized-equivalent estimate in decentralized cooperative SLAM. From this, theorems that apply to checkpoints and partial checkpoints from [2] can now be applied to the decentralized cooperative SLAM problem. These theorems include a practical method of detecting the existence of a partial checkpoint (Theorem 2.2 in [2]), which says that a partial checkpoint, $C_p(k_{c,i}, k_{e,i})$ exists for robot i if and only if $S_{i,k_{e,i}} \supseteq \{\mathbf{u}_{j,k_{c,i}} \text{ or } \mathsf{bel}(X_{j,k_{c,i}})\}, \forall j \in N$. This will be used directly within our decentralized cooperative SLAM algorithm. It is of interest to note that we are using partial checkpoints, and not checkpoints, because checkpoints are only detectable by an outside observer of our robot system (see [1] for more details).

Two other important theorems that we will use are Theorems 3.1 and 3.2 from [2]. They tell us that a robot's decision to invoke the Markov property when it detects a partial checkpoint has no effect on all other robots' abilities to obtain a partial checkpoint that occurs at the same timestep. Hence, a robot is only required to consider its own knowledge set when applying the Markov property. This is a subtle, but an important aspect of our algorithm, as robots do not need to keep track of what other robots know, or what it has communicated with other robots.

B. Initial Knowledge of Robots

In this section, we will show that it is necessary for all robots to know the total number of robots in the system to obtain centralized-equivalent estimates in decentralized cooperative SLAM. This may seem intuitive, but it is in contrast to the decentralized cooperative localization problem, where it was shown that the number of robots does not need to be known in advance to obtain the centralized-equivalent estimate [2]. For a checkpoint, we will use the term *detectable* to imply that an outside observer of the system can determine the existence (or non-existence) of the checkpoint based on information in the knowledge set of all robots. For a partial checkpoint, we will use the term *detectable* to imply that a robot can determine the existence of the partial checkpoint based solely on information in its own knowledge set. We will begin by showing that checkpoints and partial checkpoints can only be detected if and only if the number of robots is known to all robots. We will then examine whether it is possible for robots to (unknowingly) obtain the centralized-equivalent estimate without detecting it (i.e., the number of robots in the system is unknown), and show that there is no guarantee that this will ever happen in decentralized cooperative SLAM.

Theorem 2.1: *The existence of a checkpoint is detectable if and only if the number of robots in the network is known by all robots.*

Proof. Assume that the checkpoint $C(k_c, k_e)$ is detectable. This implies that we can determine whether the checkpoint exists. Theorem 1.2 from [2] indicates that a checkpoint exists if and only if

$$S_{i,k_e} \supseteq \{\mathbf{u}_{j,k_c} \text{ or } \mathsf{bel}(X_{j,k_c})\}, \forall i \in N, \forall j \in N$$

In order to apply this, each robot needs to know the number of robots in the network, $|N|$. Now let us assume that the number of robots in the network is known by all robots. By the reverse logic, we can check whether the expression shown above (of Theorem 1.2 from [2]) is true or false. This implies that each robot can validate if a checkpoint exists. In other words, if the expression above is true, each robot can detect the existence of a checkpoint. □

Corollary 2.1: *It is possible to detect when the centralized-equivalent estimate is obtainable by all robots if and only if the number of robots in the network is known by all robots.*

Proof. Theorem 1.1 tells us that the existence of a checkpoint is equivalent to when the centralized-equivalent estimate is obtainable. Since Theorem 2.1 indicates that the existence of a checkpoint can only be detected if the number of robots in the network is known, it follows that detecting when the centralized-equivalent estimate is obtainable by all robots is only possible if and only if the number of robots in the network is known by all robots. □

We will omit the proofs of the following theorem and corollary as they are similar to those in Theorem 2.1 and Corollary 2.1

Theorem 2.2: *The existence of a partial checkpoint is detectable by robot i if and only if the number of robots in the network is known to robot i.*

Corollary 2.2: *It is possible to detect when the centralized-equivalent estimate is obtainable by robot i if and only if the number of robots in the network is known to robot i.*

The question that arises now is whether a robot can obtain the centralized-equivalent estimate without detecting it (i.e., without knowing the number of robots in the system). In cooperative localization, this was possible under certain conditions on communication and observation ranges. This caused the initially uncorrelated estimate from a robot (or a subgroup of robots) to be statistically independent (i.e., remain uncorrelated) before encountering another robot (or another subgroup) where an inter-robot measurement occurs. We will show that this can never be guaranteed in decentralized cooperative SLAM due to the presence of landmarks. Consider the Fig. 2 as an example. There are two sub-groups consisting of a single robot each. Acting independently, each robot's self estimate will be correlated with the landmark, but they will not be aware that the landmark has been observed by both robots (i.e., the estimate of both robots and the landmark should be correlated) until they encounter each other. Each robot will also apply the Markov property at every timestep because they only know of their own existence. When the two robots finally encounter each other, they are unable to merge their estimates to form the centralized-equivalent estimate. We generalize this example in the following lemma and theorem.

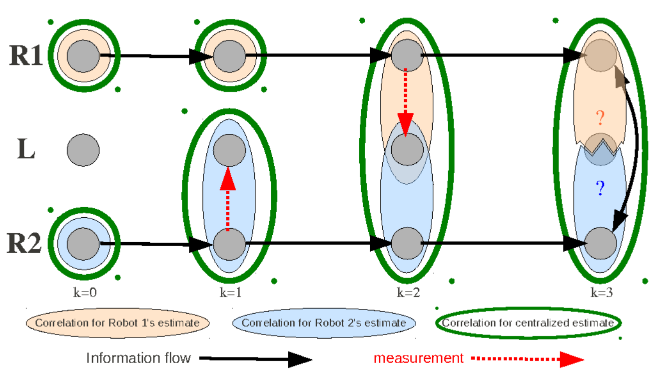

Fig. 2. A scenario illustrating Theorem 2.3.

Lemma 2.1: *Assume that a team of robots is divided into mutually-exclusive sub-groups in a decentralized cooperative SLAM scenario. The subgroups' centralized-equivalent estimates are conditionally independent of one another if and only if the subgroups do not measure any common landmarks or each other.*

Proof. Let Q represent an arbitrary sub-group of robots, and let $\overline{Q}$ represent all other robots not in Q such that $Q \cup \overline{Q} = N$. Also, let $X_{Q,k}$ represent the state of the robots in Q at time k, as well as the state for landmarks observed by the robots in Q at or prior to time k. We will first assume that sub-groups do not measure the same landmark or each other. In this case, we can write the estimate over subgroup Q as

$$\begin{aligned} \mathsf{bel}(X_{Q,k}) &= p(X_{Q,k}|\mathsf{bel}(\mathbf{x}_{i,0}), \mathbf{u}_{i,0}, Y_{i,0}), \forall i \in Q \\ &= p(X_{Q,k}|\mathsf{bel}(\mathbf{x}_{i,0}), \mathbf{u}_{i,0}, Y_{i,0}, \mathsf{bel}(\mathbf{x}_{j,0}), \\ &\qquad \mathbf{u}_{j,0}, Y_{j,0}), \forall i \in Q, \forall j \in \overline{Q} \\ &= p(X_{Q,k}|\mathsf{bel}(\mathbf{x}_{i,0}), \mathbf{u}_{i,0}, Y_{i,0}), \forall i \in N. \end{aligned}$$

From the first to second line above, we have given additional conditional dependencies to the probability density function. This has no effect since we assumed that sub-groups do not measure the same landmark or each other. Now by taking the product of the estimates from all subgroups,

$$\begin{aligned} \prod_{Q} \mathsf{bel}(X_{Q,k}) &= p(X_k|\mathsf{bel}(\mathbf{x}_{i,0}), \mathbf{u}_{i,0}, Y_{i,0}), \forall i \in N \\ &= \mathsf{bel}(X_k), \end{aligned}$$

we obtain the centralized estimate of all robots. This implies that the sub-groups' centralized-equivalent estimates are conditionally independent of one another. Next we will start by assuming that sub-groups' centralized-equivalent estimates are conditionally independent of one another, which we can express as

$$\mathsf{bel}(X_k) = \prod_{Q} \mathsf{bel}(X_{Q,k}).$$

This implies that $\mathsf{bel}(X_{Q,k})$ is statistically independent from the estimates of all other subgroups. This is only possible if they do not measure the same landmark or each other. □

Theorem 2.3: *Assume that a team of robots is divided into mutually-exclusive sub-groups in a decentralized cooperative SLAM scenario. The product of estimates from sub-groups is not guaranteed to be the centralized-equivalent estimate for the entire robot team.*

Proof. Lemma 2.1 showed that estimates from sub-groups are statistically independent if and only if the subgroups do not measure each other or the same landmarks. In decentralized cooperative SLAM, it is always possible for robots to observe the same landmarks. Therefore, the product of estimates from sub-groups is not guaranteed to be the centralized-equivalent estimate for the entire robot team. □

In summary, we cannot guarantee that a robot can obtain the centralized-equivalent estimate without detecting a partial checkpoint, which requires the robot to know the number of robots in the network. The alternate way of ensuring that the centralized-equivalent estimate is obtainable is to keep all robots from applying the Markov property (i.e., have all robots retain all information in their knowledge set through time). This however, is impractical as computation and memory usage will increase indefinitely. In summary, we need to allow robots to know when they can apply the Markov property, which requires knowing the number of robots in the system according to Corollary 2.2. This requirement is necessary due to the presence of landmarks, and although it may seem intuitive, it is of interest to show this because the same requirement is not necessary in a cooperative localization scenario.

V. Decentralized SLAM Algorithm

Algorithm (1) is our proposed decentralized cooperative SLAM algorithm, developed based on the theory provided in section IV. Fig. 3 is a graphical representation.

Algorithm 1: DecentralizedSLAM(k, $\mathbf{u}_{i,k}$, $Y_{i,k}$, $S_{i,k-1}$, $S_{j,k}$ $(i \in N)(\forall j)(d_k^{j,i} \leq r_{\mathsf{comm}})$)

1 $S_{i,k} \leftarrow S_{i,k-1} \cup \{\mathbf{u}_{i,k}\} \cup \{Y_{i,k}\} \cup \{S_{j,k}\ (\forall j, d_k^{j,i} \leq r_{\mathsf{comm}})$
2 $K \leftarrow \{k_r\}$such that$\{\mathbf{u}_{i,k_r}\} \in S_{i,k} (\forall i \in N)$
3 $k_{c,i} \leftarrow \max(K)$
4 $\tilde{S}_{i,k_{c,i}} \leftarrow S_{i,k} - \{\mathbf{u}_{j,k_r}, Y_{j,k_r}\}\ \ (\forall j, \forall k_r > k_{c,i})$
5 $\mathsf{bel}^* \left(X_{k_{c,i}}\right) \leftarrow p\left(X_{k_{c,i}} | \tilde{S}_{i,k_{c,i}}\right)$
6 $S_{i,k} \leftarrow S_{i,k} \cup \mathsf{bel}^* \left(X_{k_{c,i}}\right)$
7 $S_{i,k} \leftarrow S_{i,k} - \{\mathbf{u}_{j,k_r}, Y_{j,k_r}, \mathsf{bel}^* \left(X_{k_r}\right)\}\ \ (\forall k_r \leq k_c, \forall j)$
8 $\mathsf{bel}\left(X_k\right) \leftarrow p\left(X_k | S_{i,k}\right)$
9 **return** $\{\mathsf{bel}\left(X_k\right), S_{i,k}\}$

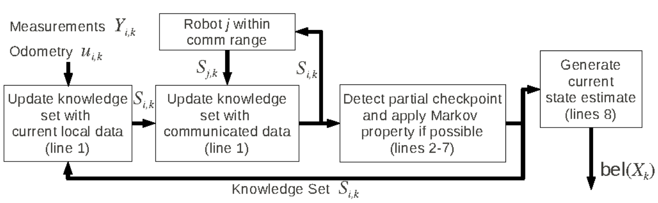

Fig. 3. The decentralized cooperative SLAM algorithm, which provides centralized-equivalent state estimates of all robots and map features whenever possible.

Algorithm 1 executes at every timestep on all robots. On line 1, we update the knowledge set of robot i by carrying out (1) and (2). In lines 2 and 3, we search for the latest partial checkpoint. Line 4 defines the subset of knowledge that includes all information up to the partial checkpoint time. On line 5, we calculate an estimate for the partial checkpoint time. Note that bel^* indicates the centralized-equivalent estimate. On line 6, we proceed to discard information replaceable by bel^* (invoking the Markov property), and enter the newly-calculated belief into the knowledge set on line 7. Finally, on line 8, we use all available information in the knowledge set to produce the current state estimate. Note that this is only a temporary estimate until we can obtain the next centralized-equivalent estimate. Robot i assumes robot j maintains its last known velocity until its next odometry measurement is received. It is important to point out again that robots are only using their local knowledge in this algorithm (i.e., they do not track what other robots know).

A number of recursive filtering (and data association) methods can be applied on lines 5 and 8. Complexity analysis of the algorithm is difficult because it is not possible to predict the topology of the dynamic network. The complexity of the algorithm is at least on the order of the filtering and data association methods used. The computational requirement for a robot will increase in proportion to the number of timesteps since its last partial checkpoint. Using the *Extended Kalman filter (EKF)* as an example, suppose that system state has n dimensions. The complexity of our algorithm will be on the order of $O((k - k_{c,i})n^4)$ for the worst case scenario (i.e., when all but one robot makes observations and communicates with each other for an extended duration). This is the cost of obtaining the centralized-equivalent estimate in a dynamic and sparse network.

VI. Experimental Setup

Our decentralized approach to SLAM was validated through experiments. Fig. 4 shows our experimental equipment, which includes a fleet of robots and a collection of landmarks. The robots use a vision based approach to obtain range and bearing measurements to other robots and landmarks. A 10-camera Vicon motion capture system is used to provide groundtruth data for both robots and landmarks at 100Hz with millimetre accuracy.

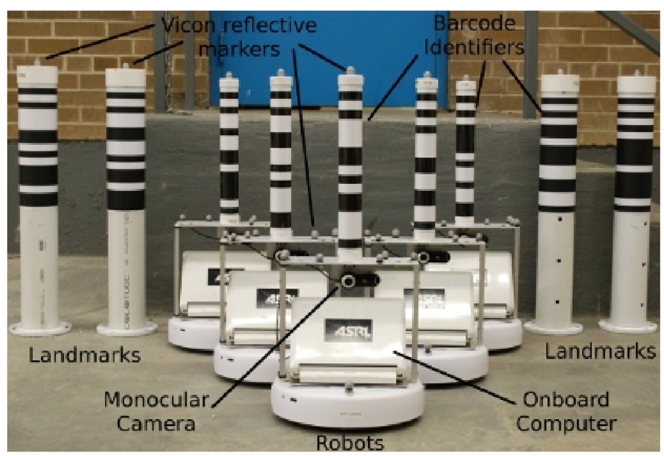

Fig. 4. Robots and landmarks for our decentralized cooperative SLAM experiments

The experiments were carried out in an area of approximately 120m^2, with 5 robots and 15 landmarks placed inside the workspace in a variety of configurations. During the experiments, the robots drove randomly in the lab space while avoiding obstacles. At the same time, the robots recorded all their odometry data, as well as range and bearing measurements with timestamps. The clocks on the robots were synchronized by running the *Network Time Protocol daemon (NTPD)* on each robot. The dataset collected from each experimental trial was then processed offline. The duration of each test ranged from approximately 20 minutes to 70 minutes, and we have collected and processed 250 minutes of data for this paper. Additional information regarding our dataset is available in [14].

VII. Experimental Results

For the results shown below, the communication range limit is 2m, and communication interval is 0.5s. The maximum observation range is 3m. The EKF is used as the filtering method in our algorithm, which runs on all robots for estimating the pose (x, y, θ) and position (x, y) of all robots and landmarks respectively. The performance of our algorithm is compared with that of the centralized EKF-SLAM algorithm. In order for the centralized estimator to work, we allow it to cheat by ignoring the communication constraints (i.e., a fully connected network is assumed for the centralized estimator). Data association is accomplished using barcode identification as this is not the focus of our work (performing real data association will only influence how we obtain the beliefs on lines 5 and 8 of Algorithm 1). The video attachment accompanying this paper contains footage of our experimental setup, and an animation of our decentralized cooperative SLAM algorithm in use. In presenting our results, we can only show some specific results from one trial (test 1) due to space constraints. Averaged results from other test trials are summarized in Table I. Fig. 5 shows the result of our decentralized cooperative SLAM algorithm at 30 minutes into test 1, in which the network was fully connected only 21% of the time.

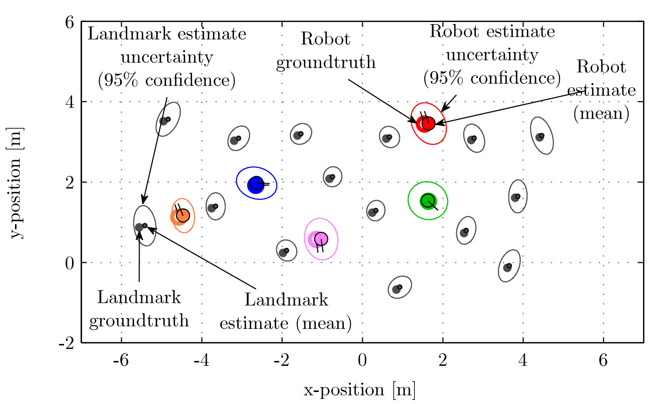

Fig. 5. A graphical representation of the results from test 1 after 1800[s].

Fig. 6 displays the memory usage of robot 1, and shows how memory use is limited since our decentralized cooperative SLAM algorithm makes use of the Markov property. Temporary increases in memory use are caused by the loss of connectivity in the dynamic robot network. These increases are followed by sharp decreases which signal partial checkpoint detections (i.e., invoking the Markov property). Similar memory usage results are observed for all other robots. The average and maximum memory usage (for all robots) in all eight trials are shown under 'Mem' in Table I.

The decentralized estimate errors made by robot 1 on several state components are shown in Fig. 7(a)–7(c) for

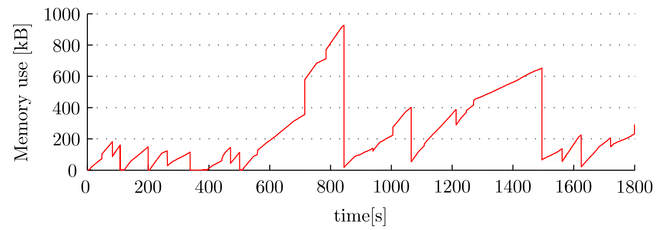

Fig. 6. Memory usage of Robot 1 in test 1.

test 1. Although we are only showing the errors of robot 1's x-position estimate of itself, robot 3, and a landmark, these plots are representative of a robot's estimate of its own pose (for which odometry information is always available), a robot's estimate of another robot's pose (for which odometry data is not always available), and a robot's estimate of a landmark position. In each of the above plots, we also compare with the centralized estimator error. It is important to remember that we are allowing the centralized estimator to cheat by ignoring the communication restrictions. Large differences observed between the centralized and decentralized errors in Fig. 7(b) occur due to the loss of communication between robot 1 and 3 for a long duration. During this time robot 1 assumes the last known velocity of robot 3 to calculate an estimate. However, this estimate is only temporary as discussed in Section V. The uncertainty regions shown on each plot are projections of the 95% confidence ellipsoid for the full-state decentralized estimate (i.e., including all robots and landmarks). Errors that are similar to the above figures are observed in the estimates of the other robots.

The difference between decentralized and centralized estimates from all 8 tests are shown in Table I for $r_{\mathsf{comm}} = 2$, where $e_{x,\mathsf{rms}}$, $e_{y,\mathsf{rms}}$, and $e_{\theta,\mathsf{rms}}$ are the root-mean-squared differences for x, y position, and θ orientation respectively. In general, for a robot's decentralized estimate of its own pose, there is very little difference compared to the centralized estimate because a robot's own odometry is always known. For a robot's estimates of other robots' poses, however, there is on average a greater difference between the centralized and decentralized estimate because the latest odometry from other robots is not always available for the current state estimate (i.e., the network is not always fully connected). Note, however, that the current state estimate is temporary and can be updated when more information is available. Our algorithm ensures that the centralized-equivalent estimate can be obtained at a later time (when a partial checkpoint is detected). As for landmarks, the average differences between the decentralized and centralized estimates are relatively small because landmarks are static. This small difference is mainly contributed by a robot not having the latest measurements from other robots. As a final note, one can expect the differences between the decentralized and centralized estimates to decrease as r_{comm} is increased.

VIII. Conclusions

In this paper, we posed the decentralized cooperative SLAM problem, where it is necessary for all the robots

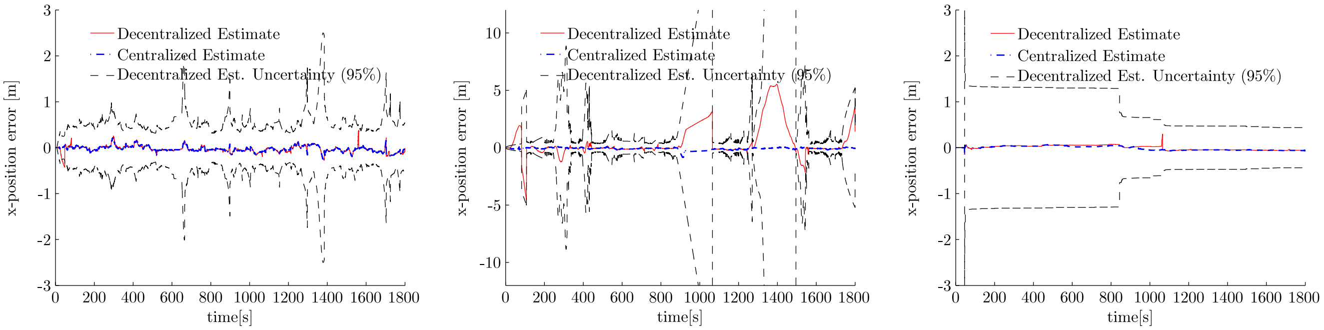

(a) Robot 1's estimate of its own x-position. (b) Robot 1's estimate of robot 3's x-position. (c) Robot 1's estimate of a landmark's x-position

Fig. 7. Components of robot 1's decentralized estimates, which matches closely with the centralized estimates

Test	Self estimate			Estimate of other robots			Estimate of landmarks		Mem [kB]		PC[%]
	$e_{x,\text{rms}}$ [m]	$e_{y,\text{rms}}$ [m]	$e_{\theta,\text{rms}}$ [rad]	$e_{x,\text{rms}}$ [m]	$e_{y,\text{rms}}$ [m]	$e_{\theta,\text{rms}}$ [rad]	$e_{x,\text{rms}}$ [m]	$e_{y,\text{rms}}$ [m]	avg	max	
1	0.0698	0.0871	0.1143	4.7854	4.9774	1.5576	0.0620	0.0457	411	1606	20.70
2	0.1107	0.1604	0.1467	2.7331	2.5393	1.4111	0.1895	0.3032	162	747	40.14
3	0.1550	0.1595	0.1632	1.2900	2.0314	1.3689	0.0853	0.1027	134	491	45.10
4	0.0693	0.0492	0.0532	0.5838	0.9072	0.9871	0.0669	0.0520	53	269	57.76
5	0.1087	0.0760	0.0991	1.2826	1.4030	1.1762	0.0104	0.0080	75	335	44.61
6	0.1073	0.0675	0.0927	1.2104	1.0937	1.0198	0.0967	0.0284	89	442	51.82
7	0.0674	0.0320	0.0594	0.6641	0.7278	0.8566	0.0507	0.0219	44	285	63.79
8	0.1050	0.1233	0.1197	1.0813	1.2597	0.9881	0.0400	0.0247	75	531	49.91

TABLE I

RMS errors between the decentralized and centralized estimates, memory usage (Mem), and percentage of time that the network is fully connected (PC).

in a team to estimate the position of the landmarks in a workspace, as well as the pose of all robots, in a network where full connectivity is never guaranteed. We proved that the conditions for which the centralized-equivalent estimate can be obtained by a robot is analogous to when a partial checkpoint exists, a concept previously applicable only to the decentralized cooperative localization problem [2]. We showed that robots can apply the Markov property based on local knowledge (i.e., they do not need to track what other robots know). Furthermore, we proved that it is necessary in decentralized cooperative SLAM for each robot to initially know the total number of robots in the network, which was not necessary for the cooperative localization case. A decentralized algorithm was presented that guarantees that the centralized-equivalent estimate can be obtained by a robot once it detects a partial checkpoint. To validate the algorithm, we carried out over 250 minutes of experiments and compared the result from our algorithm against that from a centralized estimator (which we allowed to cheat by ignoring communication constraints) and groundtruth data. Our results show how memory usage is limited in our algorithm due to use of the Markov property. The results also show that the centralized-equivalent estimate can always be recovered after a period of poor network connectivity.

Acknowledgment

This research is supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada.

References

[1] K.Y.K. Leung, T.D. Barfoot, and H.H.T. Liu. Decentralized localization for dynamic and sparse robot networks. In *Proc. of IEEE Int'l Conference on Robotics and Automation*, 2009.

[2] K. Y. K. Leung, T. D. Barfoot, and H. H. T. Liu. Decentralized localization of sparsely-communicating robot networks: A centralized-equivalent approach. *IEEE Trans. on Robotics*, 26(1):62–77, 2010.

[3] J.W. Fenwick, P.M. Newman, and J.J. Leonard. Cooperative concurrent mapping and localization. In *IEEE ICRA*, 2002.

[4] E.W. Nettleton, H.F. Durrant-Whyte, P.W. Gibbens, and A.H. Goktogan. Multiple platform localisation and map building. *Proc. SPIE*, 4196:337–347, 2000.

[5] E. Nettleton, S. Thrun, H. Durrant-Whyte, and S. Sukkarieh. Decentralised slam with low-bandwidth communication for teams of vehicles. *Field & Service Robotic*, 24:179–188, 2006.

[6] S. Reece and S. Roberts. Robust, low-bandwidth, multi-vehicle mapping. In *IEEE FUSION*, 2005.

[7] S. Thrun, Y. Liu, D. Koller, A. Ng, Ghahramani Z., and H. Durrant-Whyte. Simultaneous localization and mapping with sparse extended information filters. *IJRR*, 23:693, 2004.

[8] A. Howard, G.S. Sukhatme, and M.J. Matarić. Multirobot simultaneous localization and mapping using manifold representations. *IEEE Special Issue on Multi-robot Systems*, 94(9):1360–1369, 2006.

[9] A. Howard. Multi-robot simultaneous localization and mapping using particle filters. *IJRR*, 25(12):1243–1256, 2006.

[10] 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 IROS*, 2003.

[11] X.S. Zhou and S.I. Roumeliotis. Multi-robot slam with unknown initial correspondence: The robot rendezvous case. In *IEEE/RSJ IROS*, 2006.

[12] Z. Wang, S. Huang, and G. Dissanayake. Multi-robot simultaneous localization and mapping using d-slam framework. In *IEEE ISSNIP*, 2007.

[13] C. Stachniss. *Exploration and mapping with mobile robots*. PhD thesis, Universität Freiburg, 2006.

[14] K. Y. K. Leung, Y. Halpern, T. D. Barfoot, and H. H. T. Liu. The utias multi-robot cooperative localization and mapping dataset. *Submitted to the IJRR*, 2009.