Towards Multi-Robot Active Collaborative State Estimation via Belief Space Planning

Vadim Indelman*

***Abstract*— In this paper we address the problem of collaborative *active* state estimation within the framework of multi-robot simultaneous localization and mapping (SLAM). We assume each robot has to autonomously navigate to a pre-specified set of goals in unknown environments and develop an approach that enables the robots to collaborate in order to reduce the uncertainty in their state estimation. We formulate this problem as multi-robot belief space planning, where the belief represents the probability distribution of robot states from the entire group, as well as the mapped environment thus far. Our approach is capable of guiding each robot to reduce its uncertainty by re-observing areas previously observed (only) by other robots. Direct observations between robot states, such as relative-pose measurements, are not required, providing enhanced flexibility for the group as the robots do not have to coordinate rendezvous with each other. Instead, our framework supports indirect constraints between the robots, that are induced by mutual observations of the same area possibly at different time instances, and accounts for these future multi-robot constraints within the planning phase. The proposed approach is evaluated in a simulation study.**

I. Introduction

Accurate and reliable operation in complex, partially unknown or dynamically changing environments is essential in numerous problem domains, including autonomous navigation in GPS-deprived environments, search and rescue scenarios, object manipulation, human-robot interaction, satellite proximity operations and robotic surgery. The corresponding problem can be formulated within the simultaneous localization and mapping (SLAM) paradigm, where inference is performed over the robot states (e.g. pose, navigation state) and the perceived environment, and is also tightly related to vision-aided navigation.

While the SLAM problem has been extensively investigated for over a decade, state of the art approaches often assume robot motion to be externally determined and focus on the inference part. Although treating the two processes separately simplifies the problem, and is a necessity in certain semi-autonomous applications (e.g. hand-held camera, user-driven robots), optimal performance is *not* guaranteed, as the principle of separation [30], [2] does not hold in the general case, often resulting in sub-optimal estimation accuracy.

Addressing the fully coupled problem involves accounting for different sources of uncertainty, such as stochastic dynamics and observations, within motion planning. To do so, we need to reason in terms of the *belief,* a probabilistic distribution over states of interest (e.g. robot poses), and how the belief evolves for different actions.

*Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel.

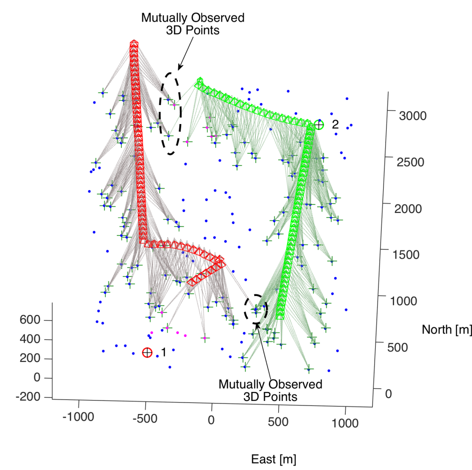

Fig. 1: 3D view of the considered synthetic scenario. Each robot adjusts its trajectory to observe areas mapped by another robot, thereby considerably reducing uncertainty (see also Figure 2).

The corresponding problem can be formulated as a partially observable Markov decision process (POMDP). Calculating a globally optimal solution to POMDP is known to be computationally intractable [22], for all but the smallest problems. Thus, the research community has been extensively investigating approximate approaches to provide better scalability to support real world problems, however, often assuming operation in known environments, as well as assuming robot state is directly observable, see e.g. [18], [24], [23], [29], [7].

In recent years, research focus has also shifted towards *active* SLAM approaches [27], [8], [26], [20], [31], [5], [28], [16] that trade-off between exploration of new areas with improving inference accuracy by actively performing loop closure observations.

In this paper we address the problem of collaborative *active* state estimation within the framework of multi-robot SLAM. Collaboration between multiple possibly heterogenous robots, capable of information sharing, has been investigated both in passive and active multi-robot SLAM contexts. Existing active approaches, however, typically focus mainly on coordination aspects to facilitate efficient exploration of new areas (see, e.g. [3], [6], [31]).

In contrast, here we approach the problem from a state estimation perspective, investigating how can the robots collaboratively improve inference quality. The *passive* problem formulation, where robots motion is externally determined and the objective is to infer states of interest as accurately as possible, has been investigated in recent years. These research efforts (including [25], [1], [17], [13], [4], [14]) explored different aspects of the problem, such as centralized vs distributed architecture, direct and indirect multi-robot constraints, and multi-robot data association.

On the other hand, *active* approaches for cooperative state estimation have been much less investigated, in particular considering operation in unknown or uncertain environments. In this paper we address this problem within the framework of multi-robot SLAM and contribute an approach that allows a group of robots to determine (locally-) optimal motion such that estimation accuracy is substantially improved while operating in unknown environments. Considering a centralized setting, we formulate the problem as multi-robot belief space planning, leveraging our previous research on (single-robot) planning in the generalized belief space [10], [11], [12].

In particular, we incorporate within the planning phase reasoning regarding (future) multi-robot constraints, which correspond to having the same scene observed by different robots, possibly at different time instances. These constraints provide incentive for the robots to properly adjust their trajectories such that appropriate areas are observed if these observations are expected to provide significant information gain (Figure 1). Moreover, our formulation does *not* require the robots to actually meet in order to benefit from each other's information, thereby providing enhanced flexibility.

II. Problem Formulation

Let x_i^r represent the pose of robot r at time t_i and denote by L_i^r the perceived environment by robot r, e.g. represented by 3D points, by that time. We let Z_i^r represent the local observations of robot r at time t_i, i.e. measurements acquired by its onboard sensors, and define the joint state Θ^r over robot past and current poses and observed 3D points as

$$\Theta_k^r \doteq X_k^r \cup L_k^r \ , \ X_k^r \doteq \{x_0^r, \ldots, x_k^r\} . \tag{1}$$

The joint probability distribution function (pdf) over this joint state given local observations $Z_{0:k}^r \doteq \{Z_0^r, \ldots, Z_k^r\}$ and controls $u_{0:k-1}^r \doteq \left\{u_0^r, \ldots, u_{k-1}^r\right\}$ is given by

$$p\left(\Theta_k^r | Z_{0:k}^r, u_{0:k-1}^r\right) \propto p\left(x_0^r\right) \prod_{i=1}^{k} \left[p\left(x_i^r | x_{i-1}^r, u_{i-1}^r\right) p\left(Z_i^r | \Theta_i^{ro}\right)\right] , \tag{2}$$

where $\Theta_i^{ro} \subseteq \Theta_i^r$ are the involved random variables in the measurement likelihood term $p\left(Z_i^r | \Theta_i^{ro}\right)$, which can be further expanded in terms of individual measurements $z_{i,j}^r \in Z_i^r$ representing observations of 3D points l_j:

$$p\left(Z_i^r | \Theta_i^{ro}\right) = \prod_j p\left(z_{i,j}^r | x_i^r, l_j\right) . \tag{3}$$

The motion and observation models in Eqs. (2) and (3) are assumed to be with with additive Gaussian noise,

$$x_{i+1}^r = f\left(x_i^r, u_i^r\right) + w_i^r \quad , \quad w_i \sim N\left(0, \Sigma_w^r\right) \tag{4}$$

$$z_{i,j}^r = h\left(x_i^r, l_j\right) + v_i^r \quad , \quad v_i \sim N\left(0, \Sigma_v^r\right) \tag{5}$$

where Σ_w^r and Σ_v^r are the process and measurement noise covariance matrices, respectively.

We consider now a group of R robots, and denote by Θ_k the corresponding joint state

$$\Theta_k \doteq X_k \cup L_k \ , \ X_k \doteq \{X_k^r\}_{r=1}^R , \tag{6}$$

comprising the past and current poses X_k of all robots, and where L_k represents the perceived environment by the entire group. Assuming a common reference frame between the robots is established, L_k includes all the 3D points in L_k^r for each r, expressed in that reference frame.

The joint pdf over Θ_k can now be written as

$$p\left(\Theta_k | Z_{0:k}, u_{0:k-1}\right) \propto \prod_{r=1}^{R} p\left(\Theta_k^r | Z_{0:k}^r, u_{0:k-1}^r\right) , \tag{7}$$

where $u_{0:k-1}$ represents the controls of all robots and is defined as $u_{0:k-1} \doteq \left\{u_{0:k-1}^r\right\}_{r=1}^R$.

One can now calculate the maximum a posteriori estimate (MAP) of Θ_k in a centralized framework as

$$\Theta_k^\star = \arg\max_{\Theta_k} p\left(\Theta_k | Z_{0:k}, u_{0:k-1}\right) . \tag{8}$$

The above formulation implicitly includes multi-robot constraints: these constraints arise from some robot r observing at the current time t_k landmarks that have been already observed by another robot r' at some time $t_{r'}$, with $t_{r'} \leq t_k$. These multi-robot constraints are the *key* to improving estimation accuracy in a collaborative multi-robot setting.

In this paper we address the problem of how such multi-robot constraints can be planned ahead of time to improve estimation accuracy. To that end, we formulate the problem within the framework of belief space planning and develop an approach capable of autonomously guiding the robots such that the mentioned multi-robot constraints can be generated, if these are expected to provide significant information gain.

While our approach naturally supports also direct multi-robot constraints (e.g. one robot makes relative-pose observations of another robot), we focus on the more general case of indirect multi-robot constraints. The latter provides enhanced flexibility as there is no requirement for the robots to actually meet each other to perform collaborative inference; instead, the latter takes place whenever a common scene is observed by different robots.

III. Multi-Robot Belief Space Planning

We formulate this problem with a belief space planning, where the belief represents the joint pdf of the entire group at an appropriate time. In particular, the belief at the current time t_k is just the joint pdf (7):

$$b\left(\Theta_k\right) \doteq p\left(\Theta_k | Z_{0:k}, u_{0:k-1}\right) , \tag{9}$$

while the belief at a future time t_{k+l} is similarly defined as $b(\Theta_{k+l}) \doteq p(\Theta_{k+l}|Z_{0:k+l}, u_{0:k+l-1})$, where $Z_{k+1:k+l}$ and $u_{l:k+l-1}$, the future observations and controls, are unknown. This belief can be written in terms of the belief $b(\Theta_k)$ from the current time and additional terms describing belief evolution at a future time:

$$b(\Theta_{k+l}) \propto b(\Theta_k) \cdot \prod_{i=1}^{l} \prod_{r=1}^{R} p\left(x_{k+i}^r | x_{k+i-1}^r, u_{k+i-1}^r\right) p\left(Z_{k+i}^r | \Theta_{k+i}^{ro}\right). \quad (10)$$

The belief $b(\Theta_{k+l})$ models the distribution of the joint state at a future time t_{k+l}, while accounting for future controls $u_{l:k+l-1} \equiv \left\{u_{l:k+l-1}^r\right\}_{r=1}^{R}$ and observations $Z_{k+1:k+l} = \left\{Z_{k+1:k+l}^r\right\}_{r=1}^{R}$ of all R robots in the group. Observe that the latter depends on the controls (different areas are observed for different controls). As we shall see next, it allows to model future multi-robot constraints via observation of mutual 3D points. Importantly, these observations that can be acquired at different time instances.

We now define a *general* multi-robot objective function

$$J(u_{k:k+L-1}) \doteq \mathbb{E}\left[\sum_{l=0}^{L} c_l\left(b(\Theta_{k+l}), u_{k+l}\right) + c_L\left(b(\Theta_{k+L})\right)\right], \quad (11)$$

that involves L look-ahead steps for all robots, and where c_l is the immediate cost function for the lth look ahead step. The expectation operator accounts for all the possible future observations $Z_{k+1:k+l}$. We are then interested in finding the optimal controls $u_{k:k+L-1}^{\star}$ for all R robots via

$$u_{k:k+L-1}^{\star} = \underset{u_{k:k+L-1}}{\arg\min}\, J(u_{k:k+L-1}). \quad (12)$$

We note that while we assumed the robots have a single objective function (11), it can still include different terms for each robot (as will be seen next). One could resort also to a different objective function for each robot.

In particular, when c_l includes the second moment of the belief $b(\Theta_{k+l})$, e.g. covariance, calculating (12) facilitates a framework for actively and collaboratively reducing uncertainty in the joint state of the group. The problem, however, is further complicated when the objective function includes additional terms, such as distance to goal and control effort, that require delicate balance of the importance of each term.

IV. Direct Trajectory Optimization

Finding a globally optimal solution $u_{k:k+L-1}^{\star}$ via Eq. (12) involves solving a partially observable Markov decision problem (POMDP), which is known to be computationally intractable [22] for all but the smallest problems. The research community has been extensively investigating approximate approaches, including point-based POMDP, sampling-based and direct trajectory optimization methods. Our approach is of the latter class and is based on the previous work [10], [11], [12], trivially extended to the multi-robot centralized case considered herein.

At each time instant, we calculate locally-optimal controls $u_{k:k+L-1}$ for all R robots, starting from a nominal trajectory or controls that are assumed to be given. The latter can be set according to the controls from the previous time instant, or determined by sampling motion planning techniques, such as RRT [19] or RRT* [15] (see, e.g., [29]).

In particular, as shown in our recent work [9], a similar probabilistic formulation that exploits multi-robot indirect constraints can be used within sampling based techniques, in which case the approach reported herein can be applied to *refine* the identified best trajectories into optimal solutions.

The overall process for calculating the optimal controls $u_{k:k+L-1}^{\star}$ can be described as a *dual-layer inference* (see [10], [11], [12] for details). The *outer layer* constitutes an iterative optimization over the controls, which are updated at each iteration according to

$$u_{k:k+L-1} \leftarrow u_{k:k+L-1} + \Delta u_{k:k+L-1}. \quad (13)$$

Different optimization methods can be used; in this work, similarly to [10], [11], we use a simple gradient method. Thus, the controls of each robot r are updated according to

$$u_{k:k+L-1}^r \leftarrow u_{k:k+L-1}^r - \lambda \nabla_r J, \quad (14)$$

where λ is an appropriate stepsize and $\nabla_r J$ is the gradient of the objective function with respect to the current solution for controls $u_{k:k+L-1}^r$ of robot r.

The *inner layer* performs inference over the belief: Recalling the definition of the objective function $J(u_{k:k+L-1})$ from Eq. (11), each iteration (13) involves inference over the joint belief $b(\Theta_{k+l}) \sim N\left(\Theta_{k+l}^{\star}, \Sigma_{k+l}\right)$, given the current control values, for each look ahead step, i.e. $l = 1, \ldots, L$.

Calculating the gradient captures how the belief, and in particular the covariance Σ_{k+l}, changes for different actions of the R robots and eventually entails the (locally) optimal controls for a given objective function. This provides a natural mechanism to facilitate *active collaborative* state estimation, i.e. adjusting robot trajectories to attain better estimation accuracy.

In the next section we discuss in detail the mentioned inner layer inference, and describe a methodology to account for future multi-robot constraints within this planning phase.

V. Incorporating Future Multi-Robot Constraints into the Joint Belief $b(\Theta_{k+l})$

We begin with the joint belief $b(\Theta_{k+l})$ from Eq. (10) and write it explicitly as

$$b(\Theta_{k+l}) = p(\Theta_{k+l}|Z_{0:k+l}, u_{0:k+l-1}). \quad (15)$$

Recall it is a function of the (unknown) future observations $Z_{k+1:k+l}$ and of the controls $u_{k:k+l-1}$ of all R robots.

While in inference the existence of observations is a given fact (either the measurement is obtained or not), when planning future actions, we can only model probabilistically whether or not *future* measurements will be acquired. Intuitively, if a 3D point is outside the camera field of view, or is too far (e.g. outside sensing range), it will not be observed by the robot sensors. In a previous work [11] we introduced binary latent variables to represent the probability of acquiring future measurements.

Here, we go one step further and note that a similar reasoning allows to model future multi-robot constraints, i.e. observation of 3D points by different robots, not necessarily at the same time. This observation is the *key element* that facilitates our approach for opportunistic active collaborative state estimation.

Specifically, for each robot r we use two types of binary latent variables, denoted by $\gamma^r_{k+l,j}$ and $\psi^r_{k+l,m}$. Variables of the first type, $\gamma^r_{k+l,j}$, represent a loop closure event to happen at time t_{k+l} (as in [11]), i.e. robot r re-observes 3D points l_j that were previously seen by robot r and possibly by other robots. In contrast, the variables $\psi^r_{k+l,m}$ represent observations, to be made by robot r, of 3D points l_m that were *only* observed by *other* robots. These latter observations induce additional (indirect) multi-robot constraints which, being part of the gradient calculation (14), attract the robots to locations that admit such constraints to take place.

Collecting these variables into corresponding sets $\mathcal{J}^r_{k+l} \doteq \left\{\gamma^r_{k+l,j}\right\}$ and $\Psi^r_{k+l} \doteq \left\{\psi^r_{k+l,m}\right\}$, and then letting

$$\Gamma_{k+l} \doteq \cup_{r=1}^{R}\Gamma^r_{k+l}\ ,\ \Psi_{k+l} \doteq \cup_{r=1}^{R}\Psi^r_{k+l}, \tag{16}$$

the joint pdf of the entire group at the lth look ahead step is given by

$$p\left(\Theta_{k+l},\Gamma_{k+l},\Psi_{k+l}|Z_{0:k+l},u_{0:k+l-1}\right). \tag{17}$$

The belief $b\left(\Theta_{k+l}\right)$, see Eq. (15), can be calculated by marginalizing out the latent variables as in

$$b(\Theta_{k+l}) = \sum_{\Gamma_{k+l}}\sum_{\Psi_{k+l}} p(\Theta_{k+l},\Gamma_{k+l},\Psi_{k+l}|Z_{0:k+l},u_{0:k+l-1}) \tag{18}$$

Let us now focus on the joint pdf (17) and write it recursively in terms of the belief from the previous step, $b\left(\Theta_{k+l-1}\right)$:

$$p\left(\Theta_{k+l},\Gamma_{k+l},\Psi_{k+l}|Z_{0:k+l},u_{0:k+l-1}\right) \propto b\left(\Theta_{k+1-1}\right)\cdot \prod_{r=1}^{R} p\left(x^r_{k+1}|x^r_k,u^r_k\right)p\left(Z^r_{k+l},\Gamma^r_{k+l},\Psi^r_{k+l}|\Theta^{ro}_{k+l}\right), \tag{19}$$

and where the last term can be expressed using the binary latent variables as

$$\begin{aligned} p\left(Z^r_{k+l},\Gamma^r_{k+l},\Psi^r_{k+l}|\Theta^{ro}_{k+l}\right) = \\ \prod_{l_j\in L^r_k} p\left(z^r_{k+l,j}|x^r_{k+l},l_j,\gamma^r_{k+l,j}\right)p\left(\gamma^r_{k+l,j}|x^r_{k+l},l_j\right)\cdot \\ \prod_{l_m\in L_k\setminus L^r_k} p\left(z^r_{k+l,m}|x^r_{k+l},l_m,\psi^r_{k+l,m}\right)p\left(\psi^r_{k+l,m}|x^r_{k+l},l_m\right). \end{aligned} \tag{20}$$

In the above equation, one can see the mentioned binary variables of both types: the first product involves the variables $\gamma^r_{k+l,j}$, accounting for the observations to be made by robot r of 3D points in L^r_k, i.e. representing areas previously observed by robot r and possibly also by other robots. The second product includes variables $\psi^r_{k+l,m}$, modeling robot r's observations of 3D points in $L_k\setminus L^r_k$, i.e. areas that were previously observed by other robots but *not* by robot r.

Intuitively, if the corresponding 3D points in $L_k\setminus L^r_k$ are estimated with high confidence (small uncertainty covariance), there is much to gain for robot r by observing (some of) these points. In such cases, we would like each robot r to be autonomously guided towards appropriate areas so that its estimation quality can be performed (see Figures 1 and 2a). In the next section we discuss in detail a mechanism to accomplish this objective.

Remark: While Eq. (20) accounts for the areas observed by the current time t_k, as represented by the 3D points L_k, one could also consider, for each lth look ahead step, *new* 3D points that will be observed for the first time during the time interval $[t_{k+1},t_{k+l}]$. As we show in [9], this facilitates a framework for active collaborative state estimation while operating in unknown environments, and can lead to significantly improved estimation accuracy.

VI. Inference over the Belief

Having discussed in detail the joint belief $b\left(\Theta_{k+l}\right)$, we now turn our attention towards the corresponding maximum a posteriori (MAP) inference, which involves calculating the first two moments of that belief. This is the inner layer inference, which is performed for each of the L look ahead steps as part of the gradient calculation in Eq. (14).

In other words, we are looking for $\Theta^\star_{k+l},\Sigma_{k+l}$ that represent the belief at the lth look ahead step:

$$b\left(\Theta_{k+l}\right) = N\left(\Theta^\star_{k+l},\Sigma_{k+l}\right). \tag{21}$$

We note that, alternatively, an information form could be used to attain better computational efficiency (as in [12]).

Calculating the MAP estimate of Θ_{k+l} involves marginalization of all latent binary variables Γ_{k+l},Ψ_{k+l} (see Eq. (18))

$$\Theta^\star_{k+l} = \underset{\Theta_{k+l}}{\arg\max}\sum_{\Gamma_{k+l}}\sum_{\Psi_{k+l}} p\left(\Theta_{k+l},\Gamma_{k+l},\Psi_{k+l}|Z_{0:k+l},u_{0:k+l-1}\right)$$

which, similarly to the single robot case, is computationally intractable. Instead, we resort to expectation maximization (EM) [21] approach and write

$$\Theta^\star_{k+l} = \underset{\Theta_{k+l}}{\arg\min}\,\mathbb{E}\left[-\log p\left(\Theta_{k+l},\Gamma_{k+l},\Psi_{k+l}|Z_{0:k+l},u_{0:k+l-1}\right)\right]$$

where the expectation is taken with respect to $p\left(\Gamma_{k+l},\Psi_{k+l}|\hat{\Theta}_{k+l},Z_{0:k+l},u_{0:k+l-1}\right)$, and where $\hat{\Theta}_{k+l}$ is the estimate of Θ_{k+l} at the current iteration.

Following a similar process as in [11], [12], the above calculation of $\Theta^\star_{k+l}$ can be written recursively as in Eq. (22) that appears at the top of next page. Linearizing about $\hat{\Theta}_{k+l}$ and employing algebraic manipulation, one can express this equation compactly as

$$\Theta^\star_{k+l} = \underset{\Theta_{k+l}}{\arg\min}\left\|A_{k+l}\Delta\Theta_{k+l}+b_{k+l}\right\|^2, \tag{23}$$

with the corresponding covariance (or, alternatively, information matrix) calculated as

$$\Sigma_{k+l} = \left(A^T_{k+l}A_{k+l}\right)^{-1}. \tag{24}$$

$$\Theta^{\star}_{k+l} = \underset{\Theta_{k+l}}{\arg\min} \left\| \Theta_{k+l-1} - \Theta^{\star}_{k+l-1} \right\|^2_{\Sigma_{k+l-1}} + \sum_{r=1}^{R} \left\| x^r_{k+l} - f\left(x^r_{k+l-1}, u^r_{k+l-1}\right) \right\|^2_{\Sigma^r_w} + \tag{22a}$$

$$+ \sum_{r=1}^{R} \sum_{l_j \in L^r_k} p\left(\gamma^r_{k+l,j} = 1 | \hat{x}^r_{k+l}, \hat{l}_j, z^r_{k+l,j}\right) \left\| z^r_{k+l,j} - h\left(x^r_{k+l}, l_j\right) \right\|^2_{\Sigma_v} + \tag{22b}$$

$$+ \sum_{r=1}^{R} \sum_{l_m \in L_k \setminus L^r_k} p\left(\psi^r_{k+l,m} = 1 | \hat{x}^r_{k+l}, \hat{l}_m, z^r_{k+l,m}\right) \left\| z^r_{k+l,m} - h\left(x^r_{k+l}, l_m\right) \right\|^2_{\Sigma_v} \tag{22c}$$

The terms $p\left(\gamma^r_{k+l,j} = 1 | \hat{x}^r_{k+l}, \hat{l}_j, z^r_{k+l,j}\right)$ and $p\left(\psi^r_{k+l,m} = 1 | \hat{x}^r_{k+l}, \hat{l}_m, z^r_{k+l,m}\right)$ in Eq. (22) represent the probabilities of observing 3D points l_j and l_m at a *future* time t_{k+l}. In standard EM inference, one can evaluate both of these terms based on the estimates $\hat{x}^r_{k+l}, \hat{l}_j$ and $\hat{l}_m$ at the current iteration and the *given* observations $z^r_{k+l,j}$ and $z^r_{k+l,m}$ (see, e.g. [14]). However, this is not possible in the planning phase, since these *future* observations are unknown (we are trying to assess whether these will be acquired).

Nevertheless, these two terms play a key role in the described dual-layer inference as they provide a mechanism to *attract* the robot to observe informative 3D points, i.e. observations that are expected to significantly impact the uncertainty covariance (24). Crucially, it allows to do so even for 3D points *outside* the robots current sensing range (or field of view). By appropriately modeling these terms, e.g. letting the probability of observing a 3D point decrease with robot distance from the latter, it is possible to induce *non-zero* contributions to the gradient $\nabla_r J$ from Eq. (14). These contributions become dominant if information gain is significant and as a result the robot will be guided towards the corresponding informative 3D points.

It is exactly this mechanism that allows to plan the mentioned indirect multi-robot constraints, as represented by the summation in Eq. (22c). In particular, the gradient $\nabla_r J$ calculated by each robot r accounts for the impact of observing 3D points $L_k \setminus L^r_k$ that were previously observed only by other robots.

In practical terms, while the summation in Eq. (22c) involves all the 3D points observed by the entire group of R robots excluding robot r, one could prefer to only account for those areas that are within a certain neighborhood of robot r. One approach to attain this is by appropriately nullifying the gradient of $p\left(\psi^r_{k+l,m} = 1 | \hat{x}^r_{k+l}, \hat{l}_m, z^r_{k+l,m}\right)$ such that it becomes zero outside a user-defined radius, centered at $\hat{x}^r_{k+l}$.

VII. Experiments

We evaluate the proposed approach for opportunistic collaborative active state estimation in a simulative environment, considering the problem of uncertainty-constrained aerial autonomous navigation in unknown and GPS-deprived environments. In the next section we describe the considered objective function J from Eq. (11) and then in Section VII-B present the results.

A. Scenario and Objective Function

In this work we consider a multi-robot uncertainty-constrained autonomous navigation scenario, where each robot r has to navigate to a pre-defined goal x^r_g while operating in unknown environments and keeping its estimation uncertainty below a given (soft) threshold β. We use a single objective function for the entire group:

$$J\left(u_{k:k+L-1}\right) \doteq \sum_{r=1}^{R} \left[\left(1-\alpha^r\right) \left\| \hat{x}^r_{k+L} - x^r_g \right\| + \alpha^r tr\left(\Sigma^r_{k+L}\right) \right], \tag{25}$$

where Σ^r_{k+L} is the covariance of robot r at the last look ahead step. For simplicity we do not include penalty over control usage and assume maximum likelihood observations, which allows to omit the expectation operator in Eq. (25). The parameter α^r represents an adaptive weight to trade-off uncertainty reduction and goal-attainment. Given an uncertainty threshold β, this parameter is calculated for each robot r based on its pose covariance Σ^r_{k+L} as (see further details in [10], [11], [12]):

$$\alpha^r \doteq min\left(\frac{tr\left(\Sigma^r_{k+L}\right)}{\beta}, 1 \right). \tag{26}$$

Since the objective function J involves the belief of all R robots (in terms of the first two moments $\hat{x}^r_{k+L}$ and Σ^r_{k+L}), calculating the actions $u^r_{k:k+L-1}$ for each robot r via Eq. (14) takes into account the impact of the latter on robot r and other involved robots. In particular, when robot r makes an observation of 3D points previously observed by other robots, the uncertainty covariance of these and possibly additional robots will be impacted. For each robot r, the gradient $\nabla_r J$ from Eq. (14) will therefore include contributions from all these robots, properly quantifying the effect of robot rth candidate controls $u^r_{k:k+L-1}$ on the entire group, a process that implicitly also involves the controls $u^{r'}_{k:k+L-1}$ of other robots $r' \in \{1, \ldots R\} \setminus \{r\}$.

We note that while the goals $\left\{x^r_g\right\}_{r=1}^{R}$ are pre-defined, an interesting question that deserves further research, is how to choose these goals online. Addressing this task allocation problem could, for example, lead to improved estimation accuracy by enhancing collaboration between nearby robots.

B. Results

We assume each robot r starts operating from a different location and needs to reach its own goal x^r_g, while localizing itself and mapping the environment perceived with its camera and range sensors. We use a soft uncertainty threshold of $\beta = 25$ meters for all robots. The proposed approach is demonstrated in two scenarios as described below.

1) First Scenario: In the first scenario, shown in Figure 2a, we consider the case of two aerial robots passing by in opposite directions, each robot being guided towards its own goal. We assume the robots have accurate knowledge regarding their starting point (i.e. initial pose) and that the robots share a common reference frame. Furthermore, the robot trajectories are assumed to be sufficiently distant from each other such that the robots will not observe any common areas, without properly adjusting their motion. As a result, without collaboration, each robot performs its own inference process (i.e. SLAM).

The result of our approach is shown in Figure 2: Figure 2a shows the robot trajectories and displays also uncertainty covariances as ellipsoids, while Figure 2b depicts the covariance evolution over time. Initially, each of the two robots proceeds directly to its goal as its covariance is small compared to the uncertainty threshold β. However, despite using SLAM as the inference engine, uncertainty does develop over time, in particular due to short feature tracks, and eventually reaches β.

At this point, which happens in general at different time instances for the two robots, the parameter α^r becomes 1 and the goal-attainment term in Eq. (25) vanishes for robot r, and as a result, robot r attempts to reduce its uncertainty. A single-robot framework would guide each robot r to re-observe informative 3D points previously mapped by robot r [10], [11], [12]. In contrast, the proposed multi-robot centralized approach instead drives robot r to observe informative areas previously seen by the other robot, r', thereby implicitly forming multi-robot constraints spanning different time instances.

As a result, uncertainty covariance is significantly reduced, as shown in Figure 2b, and the robots reach the goal with higher estimation accuracy. Figure 1 illustrates the corresponding robot trajectories and the observed 3D points by each robot. Note the 3D points mutually observed by the two robots (at different time instances) as the result of our multi-robot planning approach.

Finally, in Figure 3 we present statistical study results comprising 50 runs for the two-robot case, where each run used different noise realizations (but the same goal and robot initial positions). One can observe that overall, robot trajectories planned by our approach are consistent and smooth, and the estimation errors and uncertainty covariances are significantly reduced.

2) Second Scenario: In the second scenario we demonstrate an advantage of using the proposed multi-robot active collaborative estimation as opposed to single-robot belief planning. The same objective function (25) and uncertainty threshold β are considered, however the scenario is a bit different, as shown in Figure 4: One of the robots starts operating in an area without any 3D points and as a result its uncertainty and estimation errors develop much faster,

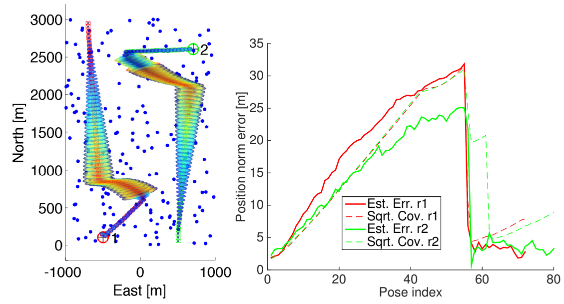

(a) (b)

Fig. 2: (a) Top view on robot trajectories. Ellipses represent uncertainty covariances. (b) Corresponding position estimation errors and uncertainty covariances. Goals for each robot are denoted by a circle and a number.

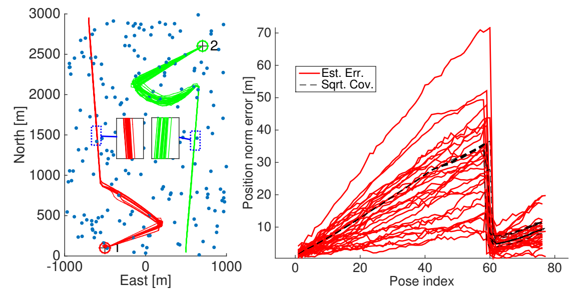

Fig. 3: Statistical study results. (a) Robot trajectories. (b) Corresponding position estimation errors (red) and square root covariances for the red robot.

reaching the uncertainty threshold β before reaching the goal.

In the single robot case, this triggers uncertainty-reduction motion planning that guides the robot to perform loop closures. However, in the considered case the goal is simply beyond the *uncertainty budget* [12]; reaching the goal with uncertainty below the threshold β is *infeasible*. This can be seen in Figure 5 that shows the robot is endlessly stuck in a cycle of going towards the goal, reaching the uncertainty threshold β, going back to reduce uncertainty via loop closures, and doing the same process forever.

In contrast, assuming the availability of another robot that happened to be operating in the same area and employing our approach for active opportunistic collaborative state estimation allows to properly adjust the robot motion such that multi-robot constraints are created and, as a result, the uncertainty is significantly reduced (Figure 4). Interestingly, this can be considered as *extending* the mentioned single-robot uncertainty budget.

VIII. Conclusions

We presented an approach for active collaborative state estimation assuming the robots operate in unknown or uncertain environments. Our belief space planning framework, where the belief accounts for the uncertainty in robot poses and in the observed environment, allows the robots to calculate locally optimal trajectories such that their state estimation is improved. We showed scenarios where this corresponds to each robot being guided towards areas previously

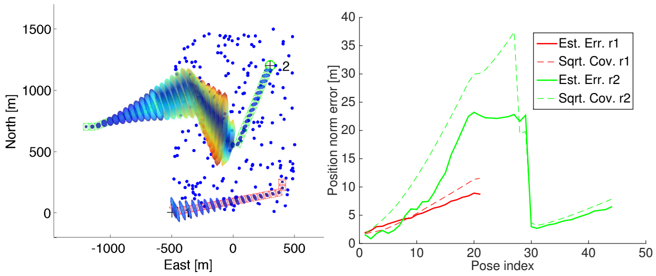

(a) (b)

Fig. 4: A different scenario, where the green robot starts operating in an area without distinctive landmarks and as a result develops significant errors. Our approach guides the robot to observe areas previously mapped by the red robot, thereby drastically reducing estimation errors and covariances. (a) top view; (b) position estimation errors and square root of uncertainty covariances.

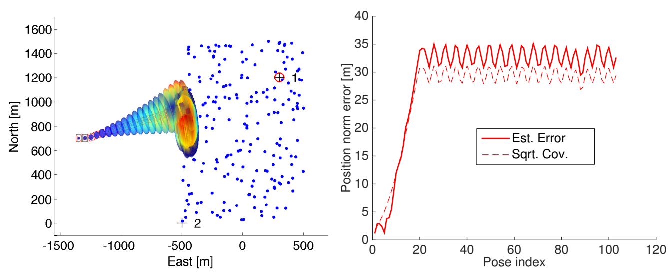

(a) (b)

Fig. 5: The same scenario as in Figure 4, however considering only a single robot. The goal is outside the uncertainty budget of a single robot. Without assistance from another robot, as in Figure 4, the robot is unable to reach the goal with the pre-defined uncertainty bound.

mapped by other robots, thereby facilitating an opportunistic collaborative state estimation framework.

Acknowledgments

This work was partially supported by the Technion Autonomous Systems Program (TASP).

References

[1] A. Bahr, M.R. Walter, and J.J. Leonard. Consistent cooperative localization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3415–3422, May 2009.

[2] Y. Bar-Shalom and E. Tse. Dual effect, certainty equivalence, and separation in stochastic control. *IEEE Transactions on Automatic Control*, 19(5):494–500, 1974.

[3] W. Burgard, M. Moors, C. Stachniss, and F. Schneider. Coordinated multi-robot exploration. *IEEE Trans. Robotics*, 2005.

[4] A. Cunningham, K. Wurm, W. Burgard, and F. Dellaert. Fully distributed scalable smoothing and mapping with robust multi-robot data association. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, St. Paul, MN, 2012.

[5] J. Du, L. Carlone, M. Kaouk Ng, B. Bona, and M. Indri. A comparative study on active SLAM and autonomous exploration with particle filters. In *Proc. of the IEEE/ASME Int. Conf. on Advanced Intelligent Mechatronics*, pages 916–923, 2011.

[6] D. Fox, J. Ko, K. Konolige, B. Limketkai, D. Schulz, and B. Stewart. Distributed multi-robot exploration and mapping. *Proceedings of the IEEE - Special Issue on Multi-robot Systems*, 94(9), Jul 2006.

[7] G. Hollinger and G. Sukhatme. Stochastic motion planning for robotic information gathering. In *Robotics: Science and Systems (RSS)*, 2013.

[8] S. Huang, N. Kwok, G. Dissanayake, Q. Ha, and G. Fang. Multi-step look-ahead trajectory planning in SLAM: Possibility and necessity. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1091–1096, 2005.

[9] V. Indelman. Towards cooperative multi-robot belief space planning in unknown environments. In *Proc. of the Intl. Symp. of Robotics Research (ISRR)*, September 2015.

[10] V. Indelman, L. Carlone, and F. Dellaert. Towards planning in generalized belief space. In *The 16th International Symposium on Robotics Research*, Singapore, December 2013.

[11] V. Indelman, L. Carlone, and F. Dellaert. Planning under uncertainty in the continuous domain: a generalized belief space approach. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014.

[12] V. Indelman, L. Carlone, and F. Dellaert. Planning in the continuous domain: a generalized belief space approach for autonomous navigation in unknown environments. *Intl. J. of Robotics Research*, 34(7):849–882, 2015.

[13] V. Indelman, P. Gurfil, E. Rivlin, and H. Rotstein. Graph-based distributed cooperative navigation for a general multi-robot measurement model. *Intl. J. of Robotics Research*, 31(9), August 2012.

[14] V. Indelman, E. Nelson, N. Michael, and F. Dellaert. Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014.

[15] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *Intl. J. of Robotics Research*, 30(7):846–894, 2011.

[16] A. Kim and R.M. Eustice. Perception-driven navigation: Active visual SLAM for robotic area coverage. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2013.

[17] B. Kim, M. Kaess, L. Fletcher, J. Leonard, A. Bachrach, N. Roy, and S. Teller. Multiple relative pose graphs for robust cooperative mapping. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 3185–3192, Anchorage, Alaska, May 2010.

[18] H. Kurniawati, D. Hsu, and W. S. Lee. Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces. In *Robotics: Science and Systems (RSS)*, volume 2008, 2008.

[19] S. M. LaValle and J. J. Kuffner. Randomized kinodynamic planning. *Intl. J. of Robotics Research*, 20(5):378–400, 2001.

[20] R. Martinez-Cantin, N. De Freitas, A. Doucet, and J.A. Castellanos. Active policy learning for robot planning and exploration under uncertainty. In *Robotics: Science and Systems (RSS)*, 2007.

[21] T.P. Minka. Expectation-Maximization as lower bound maximization. November 1998.

[22] C. Papadimitriou and J. Tsitsiklis. The complexity of markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.

[23] R. Platt, R. Tedrake, L.P. Kaelbling, and T. Lozano-Pérez. Belief space planning assuming maximum likelihood observations. In *Robotics: Science and Systems (RSS)*, pages 587–593, 2010.

[24] S. Prentice and N. Roy. The belief roadmap: Efficient planning in belief space by factoring the covariance. *Intl. J. of Robotics Research*, 2009.

[25] S.I. Roumeliotis and G.A. Bekey. Distributed multi-robot localization. *IEEE Trans. Robot. Automat.*, August 2002.

[26] C. Stachniss, G. Grisetti, and W. Burgard. Information gain-based exploration using rao-blackwellized particle filters. In *Robotics: Science and Systems (RSS)*, pages 65–72, 2005.

[27] C. Stachniss, D. Haehnel, and W. Burgard. Exploration with active loop-closing for FastSLAM. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2004.

[28] R. Valencia, J. Valls Miró, G. Dissanayake, and J. Andrade-Cetto. Active pose SLAM. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 1885–1891, 2011.

[29] J. Van Den Berg, S. Patil, and R. Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *Intl. J. of Robotics Research*, 31(11):1263–1278, 2012.

[30] W. Wonham. On the separation theorem of stochastic control. *SIAM Journal on Control*, 6(2):312–326, 1968.

[31] K.M. Wurm, C. Stachniss, and W. Burgard. Coordinated multi-robot exploration using a segmentation of the environment. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Nice, France, September 2008.