Auton Robot (2018) 42:353–373
https://doi.org/10.1007/s10514-017-9620-6

Cooperative multi-robot belief space planning for autonomous navigation in unknown environments

Vadim Indelman[1]

Received: 21 January 2016 / Accepted: 11 January 2017 / Published online: 26 January 2017

Abstract We investigate the problem of cooperative multi-robot planning in unknown environments, which is important in numerous applications in robotics. The research community has been actively developing belief space planning approaches that account for the different sources of uncertainty within planning, recently also considering uncertainty in the environment observed by planning time. We further advance the state of the art by reasoning about future observations of environments that are *unknown at planning time*. The key idea is to incorporate within the belief indirect multi-robot constraints that correspond to these future observations. Such a formulation facilitates a framework for active collaborative state estimation while operating in unknown environments. In particular, it can be used to identify best robot actions or trajectories among given candidates generated by existing motion planning approaches, or to refine nominal trajectories into locally optimal paths using direct trajectory optimization techniques. We demonstrate our approach in a multi-robot autonomous navigation scenario and consider its applicability for autonomous navigation in unknown obstacle-free and obstacle-populated environments. Results indicate that modeling future multi-robot interaction within the belief allows to determine robot actions (paths) that yield significantly improved estimation accuracy.

This is one of several papers published in *Autonomous Robots* comprising the Special Issue on Active Perception.

✉ Vadim Indelman
vadim.indelman@technion.ac.il

[1] Department of Aerospace Engineering, Technion - Israel Institute of Technology, 32000 Haifa, Israel

Keywords Multi-robot belief space planning · Active SLAM · Active perception

1 Introduction

Autonomous operation under uncertainty is essential in numerous problem domains, including autonomous navigation, object manipulation, multi-robot localization and tracking, and robotic surgery. As the robot state is never accurately known due to motion uncertainty and imperfect state estimation obtained from partial and noisy sensor measurements, planning future actions should be performed in the belief space - a probability distribution function (pdf) over robot states and additional states of interest.

Belief space planning has been investigated extensively in the last two decades. While the corresponding problem can be described in the framework of partially observable Markov decision process (POMDP), which is known to be computationally intractable for all but the smallest problems (Papadimitriou and Tsitsiklis 1987), several approaches that tradeoff optimal performance with computational complexity have been recently developed. These approaches can be segmented into several categories: point-based value iteration methods, simulation based approaches, sampling based approaches and direct trajectory optimization approaches.

Point-based value iteration methods (e.g. Pineau et al. 2006; Kurniawati et al. 2008) select a number of representative belief points and calculate a control policy over belief space by iteratively applying value updates to these points. *Simulation-based approaches* (e.g. Stachniss et al. 2005; Valencia et al. 2013) generate a few potential plans and select the best policy according to a given metric. They are referred to as simulation-based approaches, since they

simulate the evolution of the belief for each potential plan to quantify its quality.

Sampling based approaches (e.g. Prentice and Roy 2009; Bry and Roy 2011; Hollinger and Sukhatme 2014) discretize the state space using randomized exploration strategies to explore the belief space in search of an optimal plan. While many of these approaches, including probabilistic roadmap (PRM) (Kavraki et al. 1996), rapidly exploring random trees (RRT) (LaValle and Kuffner 2001) and RRT* (Karaman and Frazzoli 2011), assume perfect knowledge of the state, deterministic control and a known environment, efforts have been devoted in recent years to alleviate these restricting assumptions. These include, for example, the belief roadmap (BRM) (Prentice and Roy 2009) and the rapidly-exploring random belief trees (RRBT) (Bry and Roy 2011), where planning is performed in the belief space, thereby incorporating the predicted uncertainties of future position estimates. Another impressive recent work that incorporates motion and sensing uncertainty within sampling based planning, generalizing the PRM framework, is feedback-based information roadmap (FIRM) (Agha-Mohammadi et al. 2014). We note that similar strategies are used to address also informative planning problems (see, e.g., Hollinger and Sukhatme 2014).

Online planning approaches for POMDP, actively being developed by the AI community, are also related to the research reported herein. These approaches alleviate computational complexity by computing good local policies at each decision step during the execution (Ross et al. 2008). In particular, Monte-Carlo tree search methods (Silver and Veness 2010) have become a dominant class of solution methods for POMDPs because they can scale to large problems and use online planning to incrementally generate solutions. As will be seen in the sequel, these and other online planning approaches are relevant to the proposed method in the context of re-planning upon obstacle detection or upon obtaining new information regarding expected landmark distribution in unknown areas.

Direct trajectory optimization methods (including Platt et al. 2010; Van Den Berg et al. 2012; Indelman et al. 2015a; Patil et al. 2014) calculate locally optimal trajectories and control policies, starting from a given nominal path. Approaches in this category perform planning over a continuous state and action spaces, which is often considered more natural as the robot states (e.g., poses) and controls (e.g., steering angles) are not constrained to few discrete values. For example, Platt et al. (2010) apply linear quadratic regulation (LQR) to compute locally optimal policies, while Van Den Berg et al. (2012) develop a related method using optimization in the belief space and avoiding assuming maximum likelihood observations in predicting the belief evolution. These approaches reduce computational complexity to polynomial at the cost of guaranteeing only locally optimal solutions.

While typically, belief space planning approaches consider the environment is known, in certain scenarios of interest (e.g. navigation in unknown environments) this is not a feasible assumption. In these cases, the environment is either a priori unknown, uncertain or changes dynamically, and therefore should be appropriately modeled as part of the inference and decision making processes. Such a concept was recently developed in Indelman et al. (2013, 2015a), where random variables representing the observed environment have been incorporated into the belief and locally optimal motion plans were calculated using a direct trajectory optimization approach. In Indelman (2015a), the approach was extended to a multi-robot belief space planning centralized framework and was used to facilitate active collaborative estimation in unknown environments. Simulation- and sampling-based approaches that consider a priori unknown environments have also been recently developed in the context of active SLAM (see, e.g. Valencia et al. 2013; Chaves et al. 2014). A limitation of these approaches is that the belief only considers the environment observed by planning time and does not reason, in the context of uncertainty reduction, about new environments to be observed in the future as the robot continues exploration.

In this work we alleviate this limitation, considering the problem of cooperative multi-robot autonomous navigation in unknown environments. While it is well known that collaboration between robots can significantly improve estimation accuracy, existing approaches (e.g. Roumeliotis and Bekey 2002; Carlone et al. 2010; Indelman et al. 2012) typically focus on the inference part, considering robot actions to be determined externally. On the other hand, active multi-robot SLAM approaches (e.g. Burgard et al. 2005) typically focus on coordination aspects and on the trade-off between exploring new regions and reducing uncertainty by re-observing previously mapped areas (performing loop closures). In contrast, in this paper we consider the question - how should the robots act to collaboratively improve state estimation while autonomously navigating to individual goals and operating in unknown environments?

Addressing this question requires incorporating multi-robot collaboration aspects into belief space planning. To that end, we present an approach to evaluate the probability distributions of multiple robot states while modeling future observations of mutual areas that are unknown at planning time (Fig. 1). The key idea is that although the environment may be unknown a priori, or has not been mapped yet, it is still possible to reason in terms of robot actions that will result in the same unknown environments to be observed by multiple robots, possibly at different future time instances. Such observations can be used to formulate non-linear constraints between appropriate robot future states. Importantly, these constraints allow collaborative state estimation without the need for the robots to actually meet each other, in con-

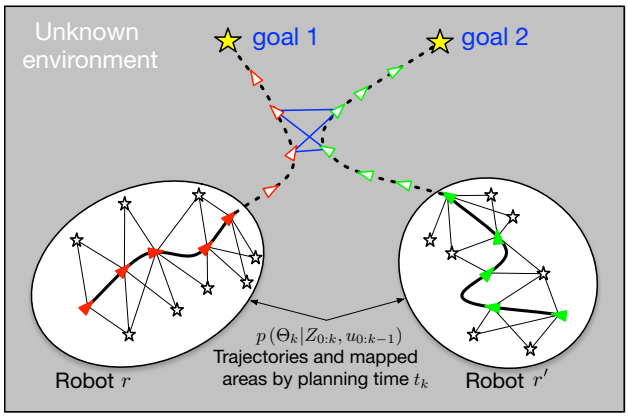

Fig. 1 Illustration of the proposed concept. Multi-robot indirect constraints representing mutual future observations of unknown environments are shown in *blue* (Color figure online)

trast to the commonly used direct relative pose observations that require rendezvous between robots (e.g. Roumeliotis and Bekey 2002). We show how such constraints can be incorporated within a multi-robot belief, given candidate paths that can be generated by any motion planning method. One can then identify the best path with respect to a user-defined objective function (e.g. reaching a goal with minimum uncertainty), and also refine best alternatives using direct trajectory optimization techniques (e.g. Van Den Berg et al. 2012; Patil et al. 2014; Indelman et al. 2015a).

In this paper we also consider a simple variation of the proposed approach to enable operation in the presence of obstacles that are initially unknown and instead detected on the fly. Such a capability is important in variety of applications, such as quadrotor navigation in uncertain environments. Upon obstacle detection, robot paths that were considered previously as optimal, do not necessarily remain as such. Thus, beyond locally modifying robot paths to avoid the detected obstacle, there is a need to replan and to update the robot optimal paths accordingly (Koenig and Likhachev 2005). We adapt the proposed approach to address this aspect by re-planning each time an obstacle is detected, while considering multi-robot collaboration aspects within the belief to yield superior estimation accuracy.

The present paper is an extension of the work presented in Indelman (2015b). As a further contribution, in this manuscript we present a more extensive performance evaluation that includes a statistical study with two different scenarios, and consider applicability of the method to autonomous operation in unknown obstacle-populated environments.

The remaining of this manuscript is structured as follows. Section 2 introduces notations and formulates the problem addressed herein. Section 3 presents the proposed approach; in particular, Sect. 3.2 discusses how to incorporate future multi-robot constraints that correspond to mutual observations of environments unknown at planning time within the multi-robot belief, while Sect. 3.5 considers the mentioned adaptation of our approach to support operation within unknown obstcale-populated environments. Section 4 presents performance evaluation, while Sect. 5 concludes the discussion and suggest some possible directions for future research.

2 Notations and 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 that robot, 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

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

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

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

where $\Theta_i^{ro} \subseteq \Theta_i^r$ are the involved random variables in the measurement likelihood term $\mathbb{P}(Z_i^r|\Theta_i^{ro})$. Assuming data association is given, this term can be further expanded in terms of individual measurements $z_{i,j}^r \in Z_i^r$ representing observations of 3D points l_j : $\mathbb{P}(Z_i^r|\Theta_i^{ro}) = \prod_j \mathbb{P}(z_{i,j}^r|x_i^r, l_j)$.

The motion and observation models in Eq. (2) are assumed to be with additive zero-mean Gaussian white noise,

$$x_{i+1}^r = f\left(x_i^r, u_i^r\right) + w_i^r \ , \ z_{i,j}^r = h\left(x_i^r, l_j\right) + v_i^r \tag{3}$$

where $w_i \sim \mathcal{N}\left(0, \Sigma_w^r\right), v_i \sim \mathcal{N}\left(0, \Sigma_v^r\right)$, with Σ_w^r and Σ_v^r representing process and measurement noise covariance matrices, respectively. Such a model is common in numerous problems in robotics, see e.g. Thrun et al. (2005).

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

$$\Theta_k \doteq X_k \cup L_k \ , \ X_k \doteq \left\{X_k^r\right\}_{r=1}^R \tag{4}$$

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, the *belief* at planning time t_k, can now be written as

$$b(\Theta_k) \doteq \mathbb{P}(\Theta_k | Z_{0:k}, u_{0:k-1}) \propto \prod_{r=1}^{R} \left\{ \mathbb{P}(x_0^r) \prod_{i=1}^{k} \left[\mathbb{P}(x_i^r | x_{i-1}^r, u_{i-1}^r) \mathbb{P}(Z_i^r | \Theta_i^{ro}) \right] \right\}, \quad (5)$$

where $u_{0:k-1}$ represents the controls of all robots and is defined as $u_{0:k-1} \doteq \left\{ u_{0:k-1}^r \right\}_{r=1}^{R}$. Similar to the single-robot case, the above decomposition is obtained by assuming the measurements are not correlated, i.e. $v_i^r \perp\!\!\!\perp v_j^{r'}$ for any $r \neq r' \in [1, R]$ and $i, j \in [0, k]$ (or $i \neq j \in [0, k]$ in case $r = r'$).

The joint belief at a future time t_{k+l} can now be similarly defined as

$$b(\Theta_{k+l}) \doteq \mathbb{P}(\Theta_{k+l} | Z_{0:k+l}, u_{0:k+l-1}), \quad (6)$$

where $u_{k:k+l-1}$ are future actions for a planning horizon of l steps and $Z_{k+1:k+l}$ are the corresponding observations to be obtained. We will discuss in detail how such a belief can be formulated in the sequel (Sects. 3.1 and 3.2).

We can now define a general multi-robot objective function

$$J(u_{k:k+L-1}) \doteq \mathbb{E}\left[\sum_{l=0}^{L} c_l \left(b(\Theta_{k+l}), u_{k+l} \right) + c_L \left(b(\Theta_{k+L}) \right) \right], \quad (7)$$

that involves L future steps for all robots, and where c_l is the immediate cost function for the lth step. The expectation operator accounts for all the possible future observations $Z_{k+1:k+l}$. While for notational convenience the same number L of future steps is assumed for all robots in Eq. (7), this assumption can be easily relaxed.

Our objective is to find the optimal controls $u^{\star}_{k:k+L-1}$ for all R robots:

$$u^{\star}_{k:k+L-1} = \arg\min_{u_{k:k+L-1}} J(u_{k:k+L-1}). \quad (8)$$

3 Approach

In this work we show how to incorporate into belief space planning multi-robot collaboration aspects such that estimation accuracy is significantly improved while operating in unknown environments. Our approach extends the state of the art by incorporating into the belief (6) multi-robot constraints induced by multiple robots observing, possibly at different *future* time instances, environments that are *unknown* at planning time. In lack of sources of absolute information (such as reliable GPS, beacons, and known 3D points), these constraints are the *key* for collaboratively improving estimation accuracy.

One can then identify best robot actions or motion plans, according to Eq. (8), among those generated by existing motion planning approaches (e.g. sampling based approaches), or resort to direct optimization techniques to obtain locally optimal solutions in a timely manner. In this work, we focus on the former case, and consider we are given candidate paths for different robots (generated, e.g. by PRM or RRT). A schematic illustration of the proposed approach is shown in Fig. 1.

We start with a recursive formulation of the multi-robot belief (Sect. 3.1) and then discuss in Sect. 3.2 our approach to incorporate into the multi-robot belief future constraints that correspond to mutual observations of unknown scenes. Evaluating the objective function (7) involves simulating belief evolution along candidate robots paths.

3.1 Recursive formulation of a multi-robot belief

We begin with a recursive formulation of the multi-robot belief (6), considering future controls $u_{0:k+l-1}$ for all robots to be given. These are determined from candidate robot paths that are being evaluated, or alternatively in the case of direct trajectory optimization approaches, the controls are determined from either nominal or perturbed robot paths (see, e.g. Indelman et al. 2015a for further details).

Given future controls for all robots, the multi-robot belief $b(\Theta_{k+l})$ at the lth future step can be written recursively as follows (see also Eq. (2)):

$$b(\Theta_{k+l}) \doteq \mathbb{P}(\Theta_{k+l} | Z_{0:k+l}, u_{0:k+l-1}) = \eta b(\Theta_{k+l-1}) \prod_{r=1}^{R} \mathbb{P}(x_{k+l}^r | x_{k+l-1}^r, u_{k+l-1}^r) \mathbb{P}(Z_{k+l}^r | \Theta_{k+l}^{ro}), \quad (9)$$

where η is a normalization constant, and $\mathbb{P}(x_{k+l}^r | x_{k+l-1}^r, u_{k+l-1}^r)$ and $\mathbb{P}(Z_{k+l}^r | \Theta_{k+l}^{ro})$ are respectively the motion model and measurement likelihood terms.

We now focus on the measurement likelihood term $\mathbb{P}(Z_{k+l}^r | \Theta_{k+l}^{ro})$, noting that it appears recursively in Eq. (9), for each look ahead step. As earlier, this term represents sensor observations of the environment (represented e.g. by 3D points), see Eq. (2). However, now, these are future observations of the environment to be made according to robot r's planned motion. It therefore makes sense to distinguish between the following two cases: (a) observation of 3D points from $L_k \subset \Theta_k$ representing environments already mapped by planning time t_k, and (b) observation of new areas that were not previously explored by any of the robots.

The former case allows to plan single- and multi-robot loop closures (e.g. as in Indelman et al. 2015a), i.e. to quantify the expected information gain due to re-observation of previously mapped areas by any of the robots.

We focus on the latter case, which has not been investigated, to the best of our knowledge, in the context of collaborative active state estimation and uncertainty reduction. Since environments that are unknown at planning time t_k are considered, the key question is how to quantify the corresponding measurement likelihood term.

3.2 Incorporating future multi-robot constraints

Despite the fact that the environments (or objects) to be observed are unknown at planning time, it is still possible to reason in terms of mutual observations of these unknown environments to be made by different robots, possibly at different future time instances. We can then formulate constraints relating appropriate robot states while marginalizing out the corresponding random variables representing the unknown environments.

More specifically, let us consider robots r and r' mutually observing at future times t_{k+l} and t_{k+j}, respectively, an unknown environment represented, e.g., by 3D points $L^{r,r'}_{k+l,k+j}$, with $1 \leq j \leq l$. The joint pdf involving the corresponding states and these 3D points is

$$\mathbb{P}(x^r_{k+l}, x^{r'}_{k+j}, L^{r,r'}_{k+l,k+j} | Z^r_{k+l}, Z^{r'}_{k+j}). \tag{10}$$

We can now marginalize out the unknown 3D points $L^{r,r'}_{k+l,k+j}$ to get

$$\mathbb{P}(x^r_{k+l}, x^{r'}_{k+j} | Z^r_{k+l}, Z^{r'}_{k+j}) = \int \mathbb{P}(x^r_{k+l}, x^{r'}_{k+j}, L^{r,r'}_{k+l,k+j} | Z^r_{k+l}, Z^{r'}_{k+j}) dL^{r,r'}_{k+l,k+j}, \tag{11}$$

which corresponds to a multi-robot constraint involving different time instances.

In the passive problem setting, i.e. controls and measurements are given, this constraint is typically a nonlinear function that involves the robot poses, say x^r_i and $x^{r'}_j$, and the measured constraint $\zeta^{r,r'}_{i,j}$ which is obtained by matching the measurements z^r_i and $z^{r'}_j$. Typical examples include matching laser scans or images using standard techniques (e.g. ICP, vision-based motion estimation). The corresponding measurement likelihood term can thus be written as

$$\mathbb{P}(\zeta^{r,r'}_{i,j} | x^r_i, x^{r'}_j) \propto \exp\left(-\frac{1}{2} \| \zeta^{r,r'}_{i,j} - g\left(x^r_i, x^{r'}_j\right) \|^2_{\Sigma^{MR}_v}\right) \tag{12}$$

where Σ^{MR}_v is the corresponding measurement noise covariance matrix, and g is an appropriate measurement function (the specific function used in our implementation is provided in Sect. 4.1). For example, this function could represent a nonlinear relative pose constraint.

We note that in SLAM literature, Eq. (12) is commonly termed a factor, and the product of all existing probabilistic terms (or factors), e.g. such as in Eq. (2), can be represented by a graphical model, a factor graph, which admits computationally efficient sparsity-aware inference (Dellaert 2012; Kaess et al. 2012).

Revisiting Eq. (11), while in our case the *future* observations are *not* given, the reasoning is very similar: we can denote by $\zeta^{r,r'}_{k+l,k+j}$ the measured constraint that would be obtained by matching z^r_{k+l} and $z^{r'}_{k+j}$ if these were known, and considering, as before, the match is successful (i.e. not outlier), it is possible to quantify the measurement likelihood (11) as

$$\mathbb{P}(\zeta^{r,r'}_{k+l,k+j} | x^r_{k+l}, x^{r'}_{k+j}) \propto \exp\left(-\frac{1}{2} \| \zeta^{r,r'}_{k+l,k+j} - g\left(x^r_{k+l}, x^{r'}_{k+j}\right) \|^2_{\Sigma^{MR}_v}\right) \tag{13}$$

Note the above assumes robots r and r' will observe the same unknown scene from future states x^r_{k+l} and $x^{r'}_{k+j}$. How to determine if two future measurements (e.g. images, laser scans), to be captured from robot poses x^r_{k+l} and $x^{r'}_{k+j}$, will be overlapping, i.e. represent a mutually observed a scene? The answer to this question is scenario specific. For example, in an aerial scenario with robots equipped with downward looking cameras, it is possible to assess if the images are overlapping given robot poses and a rough estimate of height above ground. Ground scenarios allow similar reasoning, however here it is more likely that the same (unknown) scene is observed from multiple views (e.g. autonomous driving with a forward looking camera), and moreover, obstacles, that are unknown at planning time, may prevent two adjacent views to observe a mutual scene in practice.

In this paper we assume one is able to predict if two future poses will mutually observe a scene. Specifically, in Sect. 4 we consider aerial robots with downward facing cameras and take a simplified approach, considering two future poses x^r_{k+l} and $x^{r'}_{k+j}$ to overlap if they are "sufficiently" nearby, quantified by a relative distance below a threshold d (in our implementation we use $d = 300$m). Naturally, more advanced approaches can be considered (e.g. account also for viewpoint variation) and be encapsulated by an indicator function as in Levine et al. (2013), Regev and Indelman (2016) - we leave the investigation of these aspects to future research.

Given candidate robot paths it is possible to determine using the above method which future views (poses) will over-

lap and formulate the corresponding multi-robot constraints (13). In particular, multi-robot constraints between robot r at time t_{k+l} and other robots r' at time t_{k+j} with $0 \leq j \leq l$ can be enumerated as

$$\prod_j \mathbb{P}(\zeta^{r,r'}_{k+l,k+j} | x^r_{k+l}, x^{r'}_{k+j}). \tag{14}$$

A similar reasoning we can be also applied considering future viewpoints of any single robot $r \in [1, R]$, given a candidate path: for viewpoints x^r_{k+l} that pass through a currently unknown environment (i.e. previously unobserved areas), it is possible to write the following constraints for any sufficiently overlapping two views

$$\prod_s \mathbb{P}(\zeta^r_{k+l,k+l-s} | x^r_{k+l}, x^r_{k+l-s}), \tag{15}$$

with the same criteria for determining if two views x^r_{k+l} and x^r_{k+l-s} are sufficiently nearby as in the multi-robot case (e.g. in the current implementation, relative distance below a threshold $d = 300$m).

We can now write the measurement likelihood term $\mathbb{P}(Z^r_{k+l}|\Theta^{ro}_{k+l})$ from Eq. (9) as:

$$\begin{aligned}
&\mathbb{P}(Z^r_{k+l}|\Theta^{ro}_{k+l}) \\
&= \prod_{l_j \in \Theta^{ro}_{k+l}} \mathbb{P}(z^r_{k+l,j} | x^r_{k+l}, l_j) \cdot \prod_s \mathbb{P}(\zeta^r_{k+l,k+l-s} | x^r_{k+l}, x^r_{k+l-s}) \cdot \\
&\cdot \prod_j \mathbb{P}(\zeta^{r,r'}_{k+l,k+j} | x^r_{k+l}, x^{r'}_{k+j}).
\end{aligned} \tag{16}$$

Before proceeding further, recall the notations used in the above equation: $z^r_{k+l,j}$ refers to a raw measurement made by robot r from viewpoint x^r_{k+l} of landmark l_j, while $\zeta^r_{k+l,k+l-s}$ refers to a relative pose constraint calculated from raw measurements Z^r_{k+l} and Z^r_{k+l-s} ($\zeta^{r,r'}_{k+l,k+j}$ is defined similarly); Z^r_i denotes all the raw observations obtained by robot r at time instant i.

The first product represents observations of previously mapped 3D points $l_j \in L_k$, with Θ^{ro}_{k+l} including those 3D points that are actually visible from x^r_{k+l}. Determining the latter requires solving data association; while this task can by itself be challenging, especially in perceptually aliased environments, here we assume data association to be externally solved and perfect, as typically done in belief space planning literature. We refer the interested reader to our recent research (Pathak et al. 2016b, a) that relaxes this assumption and incorporates reasoning regarding data association within belief space planning.

The second and the third products in Eq. (16) correspond, respectively, to constraints stemming from mutually observing unknown scenes from different viewpoints to be taken either by a single or multiple robots. See schematic illustration in Fig. 1, where these future constraints are shown in blue.

Substituting Eq. (16) into Eq. (9) yields the final expression for $b(\Theta_{k+l})$:

$$\begin{aligned}
b(\Theta_{k+l}) = \eta b(\Theta_{k+l-1}) \prod_{r=1}^{R} \Bigg[& \mathbb{P}(x^r_{k+l} | x^r_{k+l-1}, u^r_{k+l-1}) \\
& \prod_{l_j \in \Theta^{ro}_{k+l}} \mathbb{P}(z^r_{k+l,j} | x^r_{k+l}, l_j) \\
& \prod_s \mathbb{P}(\zeta^r_{k+l,k+l-s} | x^r_{k+l}, x^r_{k+l-s}) \\
& \cdot \prod_j \mathbb{P}(\zeta^{r,r'}_{k+l,k+j} | x^r_{k+l}, x^{r'}_{k+j}) \Bigg].
\end{aligned} \tag{17}$$

Several remarks are in order at this point. First, observe that direct multi-robot constraints, where a robot measures its pose relative to another robot, are naturally supported in the above formulation by considering the same (future) time index, i.e. $\mathbb{P}(\zeta^{r,r'}_{k+l,k+l} | x^r_{k+l}, x^{r'}_{k+l})$. Of course, being able to formulate constraints involving also different future time instances, as in Eq. (17), provides enhanced flexibility since planning rendezvous between robots is no longer required.

Second, observe the constraint formulations in Eqs. (14) and (15) is an approximation of the underlying joint pdf of *multiple* views X making observations Z of an unknown scene L, since it only considers *pairwise* potentials. More concretely, marginalizing L out,

$$\mathbb{P}(X|Z) = \int \mathbb{P}(X, L|Z) dL, \tag{18}$$

introduces mutual information between all views in X, i.e. any two views in X become correlated (see related aspects also in structureless bundle adjustment, e.g. Indelman et al. 2015b; Indelman and Dellaert 2015). Thus, a more accurate formulation than (14) would consider all robot poses observing a mutual scene together. We note, however, that using a pairwise potentials formulation is fairly standard in PoseSLAM literature (e.g. Eustice et al. 2006; Kim and Eustice 2014).

Finally, one could also incorporate reasoning regarding (robust) data association, i.e. whether a constraint $\zeta^{r,r'}_{k+l,k+j}$ stemming from matching raw measurements (images, laser scans) Z^r_{k+l} and $Z^{r'}_{k+j}$ is expected to be an inlier, as for example done in Indelman et al. (2014), Indelman et al. (2016) for the passive case, and in Pathak et al. (2016b), Pathak et al. (2016a) for the active case.

3.3 Inference over multi-robot belief given controls

Having described in detail the formulation of a multi-robot belief $b\left(\Theta_{k+l-1}\right)$ at each future time t_{k+l}, this section focuses on simulating belief evolution over time given robot controls or paths. As discussed in Sect. 3, this calculation is required both for sampling based motion planning and direct trajectory optimization approaches.

Thus, we are interested in evaluating the belief $b\left(\Theta_{k+l}\right)$ from Eq. (17). Considering Gaussian prior and noise distributions, and given data association, as assumed herein, this belief is the Gaussian

$$b\left(\Theta_{k+l}\right) \equiv \mathbb{P}(\Theta_{k+l}|Z_{0:k+l}, u_{0:k+l-1}) = \mathcal{N}\left(\Theta^{\star}_{k+l}, \Sigma_{k+l}\right), \tag{19}$$

with appropriate mean $\Theta^{\star}_{k+l} \equiv \hat{\Theta}_{k+l|k+l}$ and covariance $\Sigma_{k+l} = \Lambda^{-1}_{k+l}$, where $\Lambda_{k+l} \equiv \Lambda_{k+l|k+l}$ is the corresponding information matrix.

This process involves maximum a posteriori (MAP) inference

$$\Theta^{\star}_{k+l} = \arg\max_{\Theta_{k+l}} b\left(\Theta_{k+l}\right) = \arg\min_{\Theta_{k+l}} \left[-\log b\left(\Theta_{k+l}\right)\right], \tag{20}$$

which also determines the information matrix Λ_{k+l}.

To perform this inference, recall the recursive formulation (9) and denote the MAP inference of the belief at a previous time by $b\left(\Theta_{k+l-1}\right) = N\left(\Theta^{\star}_{k+l-1}, I_{k+l-1}\right)$. The belief at time t_{k+l} can therefore be written as

$$\begin{aligned} -\log b\left(\Theta_{k+l}\right) = & \left\|\Theta_{k+l-1} - \Theta^{\star}_{k+l-1}\right\|^2_{\Sigma_{k+l-1}} \\ & + \sum_{r=1}^{R}\Big[\left\|x^r_{k+l} - f(x^r_{k+l-1}, u^r_{k+l-1})\right\|^2_{\Sigma_Q} \\ & - \log \mathbb{P}(Z^r_{k+l}|\Theta^{ro}_{k+l})\Big] \end{aligned} \tag{21}$$

We now focus on the term $-\log \mathbb{P}(Z^r_{k+l}|\Theta^{ro}_{k+l})$. Recalling the discussion from Sect. 3.2 and Eq. (16), this term can be written as

$$\begin{aligned} -\log p\left(Z^r_{k+l}|\Theta^{ro}_{k+l}\right) = & \sum_{l_j \in \Theta^{ro}_{k+l}} \left\|z^r_{k+l,j} - h(x^r_{k+l}, l_j)\right\|^2_{\Sigma_v} \\ & + \sum_s \left\|\zeta^r_{k+l,k+l-s} - g(x^r_{k+l}, x^r_{k+l-s})\right\|^2_{\Sigma_v} \\ & + \sum_j \left\|\zeta^{r,r'}_{k+l,k+j} - g(x^r_{k+l}, x^{r'}_{k+j})\right\|^2_{\Sigma^{MR}_v}, \end{aligned} \tag{22}$$

where the motion and measurement models f and h are defined in Sect. 2, and the nonlinear function g was introduced in Eqs. (12) and (13). We note that while here we consider the measurement noise covariance Σ^{MR}_v to be constant, one could go further and model also accuracy deterioration, e.g. as the relative distance between robot poses increases.

We now proceed with the MAP inference (20), which, if the future observations Z^r_{k+l} were known, could be solved using standard iterative non-linear optimization techniques (e.g. Gauss–Newton and Levenberg–Marquardt): in each iteration the system is linearized, the delta vector $\Delta\Theta_{k+l}$ is recovered and used to update the linearization point, and the process is repeated until convergence.

Let us first describe in more detail this fairly standard approach, considering for a moment the future measurements Z^r_{k+l} are known. The linearization point $\bar{\Theta}_{k+l}$ is discussed first. Recalling that we are to evaluate belief evolution given robot paths, these paths can be considered as the linearization point for robot poses. On the other hand, in the case of direct trajectory optimization approaches, the nominal controls over the planning horizon can be used to generate the corresponding nominal trajectories according to (similar to the single robot case, see, e.g. Indelman et al. 2015a)

$$\bar{x}^r_{k+l} = \begin{cases} f(\bar{x}^r_{k+l-1}, u^r_{k+l-1}), & l > 1 \\ f(\hat{x}^r_k, u^r_k), & l = 1 \end{cases} \tag{23}$$

The linearizaiton point for the landmarks $L_k \subset \Theta_{k+l}$ (see Sect. 2) is taken as their most recent MAP estimate. We first linearize Eq. (21)

$$\begin{aligned} -\log b\left(\Theta_{k+l}\right) = & \left\|B_{k+l}\Delta\Theta_{k+l}\right\|^2_{\Sigma_{k+l-1}} \\ & + \sum_{r=1}^{R}\left[\left\|F^r_{k+l}\Delta\Theta_{k+l} - b^r_{k+l}\right\|^2_{\Sigma_Q} - \log \mathbb{P}(Z^r_{k+l}|\Theta^{ro}_{k+l})\right] \end{aligned} \tag{24}$$

and then linearize the term $-\log p\left(Z^r_{k+l}|\Theta^{ro}_{k+l}\right)$ from Eq. (22):

$$\begin{aligned} -\log p\left(Z^r_{k+l}|\Theta^{ro}_{k+l}\right) = & \sum_{l_j \in \Theta^{ro}_{k+l}} \left\|H^r_{k+l,j}\Delta\Theta_{k+l} - b^r_{k+l,j}\right\|^2_{\Sigma_v} \\ & + \sum_s \left\|G^r_{k+l,k+l-s}\Delta\Theta_{k+l} - b^r_{k+l,k+l-s})\right\|^2_{\Sigma_v} \\ & + \sum_j \left\|G^{r,r'}_{k+l,k+j}\Delta\Theta_{k+l} - b^{r,r'}_{k+l,k+j})\right\|^2_{\Sigma^{MR}_v}, \end{aligned} \tag{25}$$

where the matrices F, H and G and the vectors b are the appropriate Jacobians and right-hand-side (rhs) vectors. The binary matrix B_{k+l} in Eq. (24) is conveniently defined such that $B_{k+l}\Delta\Theta_{k+l} = \Delta\Theta_{k+l-1}$.

Using the relation $\Sigma^{-1} \equiv \Sigma^{-\frac{T}{2}}\Sigma^{-\frac{1}{2}}$ to switch from $\|a\|^2_{\Sigma}$ to $\|\Sigma^{-\frac{1}{2}}a\|^2$ and stacking all the Jacobians and rhs vectors into $\mathcal{A}_{k+l}$ and $\breve{b}_{k+l}$, respectively, we get

$$\Delta\Theta^{\star}_{k+l} = \underset{\Delta\Theta_{k+l}}{\arg\min} \left\| \mathcal{A}_{k+l}\Delta\Theta_{k+l} - \check{b}_{k+l} \right\|^2. \tag{26}$$

The a posteriori information matrix Λ_{k+l} of the joint state vector Θ_{k+l} can thus be calculated as

$$\Lambda_{k+l} = \mathcal{A}^T_{k+l}\mathcal{A}_{k+l}. \tag{27}$$

This constitutes the first iteration of the nonlinear optimization. Recalling again that the future observations Z^r_{k+l} are unknown, it is not difficult to show (Indelman et al. 2015a that, while the a posteriori information matrix I_{k+l} is not a function of these observations, the equivalent rhs vector $\check{b}_{k+l}$ from Eq. (26) does depend on Z^r_{k+l}. This presents difficulties in carrying out additional iterations as the linearization point itself becomes a function of the unknown random variables Z^r_{k+l}.

As common in related works (e.g. Platt et al. 2010; Van Den Berg et al. 2012; Patil et al. 2014; Indelman et al. 2015a), we assume a single iteration sufficiently captures the impact of a candidate action(s). Alternatively, to better predict uncertainty evolution, one could resort to using the unscented transformation, as in He et al. (2008), or to particle filtering techniques. Furthermore, for simplicity in this paper we also make the maximum-likelihood measurement assumption, according to which a future measurement z is assumed equal to the predicted measurement using the most recent state estimate. As a result, it can be shown that the rhs vector $\check{b}_{k+l}$ becomes zero and thus $\Theta^{\star}_{k+l} = \bar{\Theta}_{k+l}$. We note one could avoid making this assumption altogether at the cost of more complicated expressions, see, e.g. Van Den Berg et al. (2012), Indelman et al. (2015a).

To summarize, the output of the described inference procedure is a Gaussian that models the multi-robot belief as in Eq. (19):

$$b\left(\Theta_{k+l}\right) = \mathcal{N}\left(\Theta^{\star}_{k+l}, \Lambda_{k+l}\right). \tag{28}$$

3.4 Cooperative belief space planning

Solving optimally[1] the cooperative belief space planning problem involves considering all the combinations between candidate paths of different robots. Specifically, given candidate paths for robots in the group, one can identify the best candidates by evaluating the objective function J from Eq. (7) for different path combinations. Such a process involves simulating belief evolution along the candidate paths of different robots in the group, as discussed in Sect. 3.3, while accounting for multi-robot collaboration in terms of mutual observations of unknown environments (as discussed in Sect. 3.2).

The focus of this paper is primarily given to incorporating within the belief multi-robot factors that represent mutual observations of environments unknown at planning time (see Sect. 3.2), and evaluating the improvement in estimation accuracy as a result of planned collaboration between robots. Thus, while the computational complexity of the above process is exponential with the number of robots, we leave the investigation of the proposed concept within other planning alternatives (e.g. Levine et al. 2013; Amato et al. 2014; Atanasov et al. 2015) to future research. In our recent research (Regev and Indelman 2016) we made a first step in this direction. Additionally, although also outside the scope of this paper, one could resort to direct trajectory optimization approaches (e.g. Indelman et al. 2015a; Patil et al. 2014; Van Den Berg et al. 2012; Indelman 2015a) to calculate locally optimal robot paths given nominal paths.

3.5 Unknown environment with obstacles

While previously the environment was assumed to be unknown and obstacle-free, in this section we show applicability of our approach also to the case when the environment includes obstacles. There are two cases of interest: the obstacles can be either a priori known or unknown.

In the first case, the known obstacles can be efficiently avoided by discarding appropriate candidate paths, as commonly done in sampling based approaches. Moreover, one could envision using the known obstacles as landmarks to localize the robots, if it can be assumed that the obstacles are different in appearance such that one can reliably determine which one is being observed (i.e. data association is solved).

The more challenging case, however, is when the environment includes obstacles which are unknown and should therefore be detected on the fly. In this section we consider a simple modification of our approach to support such a scenario. We assume existence of a black box algorithm capable of obstacle detection. For simplicity, we consider the entire obstacle to be detected once it is within the field of view of the sensor. We leave further investigation of the proposed approach while relaxing this assumption to future research.

Once an obstacle is detected, it triggers re-calculation of robot candidate paths, performing collision checking with respect to the obstacle(s) detected thus far and discarding those paths that collide with any of the obstacles. The objective function is then evaluated for the valid candidate paths and the best path for each robot is selected. In essence, this framework is closely related to model predictive control (MPC) based belief space planning, where only a portion of the optimal actions is executed, followed by re-planning. In the current case, re-planning is simply triggered by detection of a new obstacle.

[1] Optimality here refers to choosing the best actions from the given set of candidate paths.

Uncertainty evolution could also be considered in the context of obstacle collision avoidance, as commonly done within chance-constrained motion planning approaches (see, e.g. Bry and Roy 2011). However, these approaches typically consider known obstacles, which makes it essential to quantify chances of collision given the evolving uncertainty. In contrast, we consider a priori unknown obstacles that are detected on-line. In such a setting, the absolute uncertainty (with respect to some global frame) has less of a meaning in the context of collision avoidance and only relative uncertainty, with respect to the uncertainty from the time the obstacle has been detect could matter. For simplicity, we assume herein this relative uncertainty is small and do not consider it while determining valid robot paths. Further investigation of related aspects is left to future research.

4 Simulation results

In this section we demonstrate the proposed approach considering the problem of multi-robot autonomous navigation while operating in unknown GPS-deprived environments. We consider an aerial scenario, where each robot has its own goal and the objective is to reach these goals in minimum time but also with highest accuracy. This can be quantified by the follwing objective function:

$$J = \sum_{r=1}^{R} \left[\kappa_{path}^{r} pathlen_{goal}^{r} + \kappa_{uncert}^{r} \sqrt{tr\left(\Sigma_{goal}^{r}\right)} \right], \tag{29}$$

where Σ_{goal}^{r} and $pathlen_{goal}^{r}$ represent, respectively, the covariance upon reaching the goal and path length for robot r. A robot decides it reached a goal if its estimated location and the goal location coincide or sufficiently small. This means we do not assume availability of known characteristics of the goal (such as visual appearance) that could be used to reset the estimation error; estimation quality is thus of prime importance in such a setting. The design parameters κ_{path}^{r} and κ_{uncert}^{r} weight the importance of each term. In our current implementation, we use $\kappa_{path}^{r} = 0.1$ and $\kappa_{uncert}^{r} = 10$.

As the environment is unknown and there are no beacons, radio sources or any other means to reset estimation error, the robots can only rely on onboard sensing capabilities and collaboration with each other to reduce drift as much as possible. We assume each robot is equipped with camera and range sensors and can observe natural landmarks in the environment, which are used to estimate robot pose within a standard SLAM framework. However, since the environment is unknown ahead of time, these landmarks are discovered on the fly while the planning process has access only to environments observed by planning time (Sect. 3). Initial relative poses between the robots are assumed to be known, such that the robots have a common reference frame—approaches that relax this assumption do exist (e.g. Indelman et al. 2014).

In this basic study we use a state of the art sampling based motion planning approach, a probabilistic roadmap (PRM) Kavraki et al. (1996), to discretize the environment and generate candidate paths for different robots over the generated roadmap. We use uniform sampling within PRM, which corresponds to assuming uninformative prior regarding the landmark distributions in unknown environments. The sampling process could be accordingly biased in case another distribution is available, or learned on-line. The approach should be also applicable to other alternatives, such as the recently introduced RRT* and RRG. As discussed in Sect. 3.4, in this paper we consider a centralized framework and exhaustively evaluate the objective function for all candidate combinations between different robots.

The next section provides some implementation details and then Sect. 4.2 examines the performance of the proposed approach and compare it to an alternative approach that does not incorporate multi-robot factors within belief space planning. Section 4.2 assumes an obstacle-free unknown environment and considers only a few candidate paths to demonstrate key aspects. Section 4.3 then reports the results of a statistical study, where each run involves numerous candidate paths, still assuming an obstacle-free environment. Results for two scenarios are presented. Finally, in Sect. 4.4 we consider unknown obstacle-populated environments, trivially adjust the method to an incremental setting and re-plan each time an obstacle is detected (see Sect. 3.5).

4.1 Implementation details

The developed approach was implemented in MATLAB using the GTSAM toolbox Dellaert (2012). In this implementation, the mentioned multi-robot constraints, possibly involving different future time instances, are formulated between any two poses with relative distance closer than d meters. We use $d = 300$ m for this threshold parameter (in the considered scenario the aerial robots height is about 500 m). More advanced methods could be implemented of course, considering also viewpoint variation, uncertainty associated with involved estimation of viewpoints and incorporating statistical knowledge. Investigating this direction is left to future research.

In all cases, a downward facing camera and a range sensor were simulated, with the nominal (ideal) measurements corrupted by a zero-mean white noise drawn from a Gaussian with standard deviation (std) of $\sigma_v = 0.5$ pixels (for each axis) and $\sigma_v = 1$ meters for image and range observations, respectively. The camera calibration matrix was

$$K = \begin{bmatrix} 2080 & 0 & 1200 \\ 0 & 2570 & 900 \\ 0 & 0 & 1 \end{bmatrix} \tag{30}$$

The motion model (3) was represented by odometry corrupted with Gaussian noise with std of 0.5° in rotation and 1 m in translation, in each axis. All robots shared the same motion and observation models.

The corresponding measurement model functions h (from Eq. (3)) to the camera and range sensors are, respectively

$$h_{\text{cam}}(x_i, l_j) \doteq \pi(x_i, l_j) = K\,[x_i.R, x_i.t] \begin{pmatrix} l_j \\ 1 \end{pmatrix}, \tag{31}$$

with $\pi(.)$ denoting a projection operator Hartley and Zisserman (2004), and

$$h_{\text{range}}(x_i, l_j) \doteq \|x_i.t - l_j\|. \tag{32}$$

Here, $x_i.R$ and $x_i.t$ denote the rotation and translation of pose x_i.

The multi-robot future constraints $\zeta_{i,j}^{r,r'}$ (see Eq. (13)) represent, in our current implementation, relative pose measurements; thus the measurement function g is defined as

$$g(x_i^r, x_j^{r'}) \doteq x_i^r \ominus x_j^{r'}. \tag{33}$$

Single robot constraints are similarly defined. Here, we follow Lu and Milios (1997) and use the notation $\ominus$ in $a \ominus b$ to express b locally in the frame of a for any two poses a and b.

Finally, the implemented criteria for considering if two robot views x_i^r and $x_j^{r'}$ overlap (see Eqs. 14–15) is given by

$$\|x_i.t - x_j.t\| < d \tag{34}$$

with $d = 300$ m.

4.2 Obstacle-free unknown environment

Figure 3 shows some of these candidate paths considering a scenario of two robots starting operating from different locations. In each case we also show the belief evolution (in terms of uncertainty covariance) along each path, calculated as described in Sect. 3.3, and the multi-robot constraints that have been incorporated into the appropriate beliefs (denoted by cyan color).

As seen in Fig. 3, only in two of the considered cases (Fig. 3b, c), robot paths were sufficiently close to facilitate multi-robot constraints within belief space planning. In practice, however, only in the latter case numerous informative constraints have been incorporated. Figure 4 compares between the two terms in the considered objective function

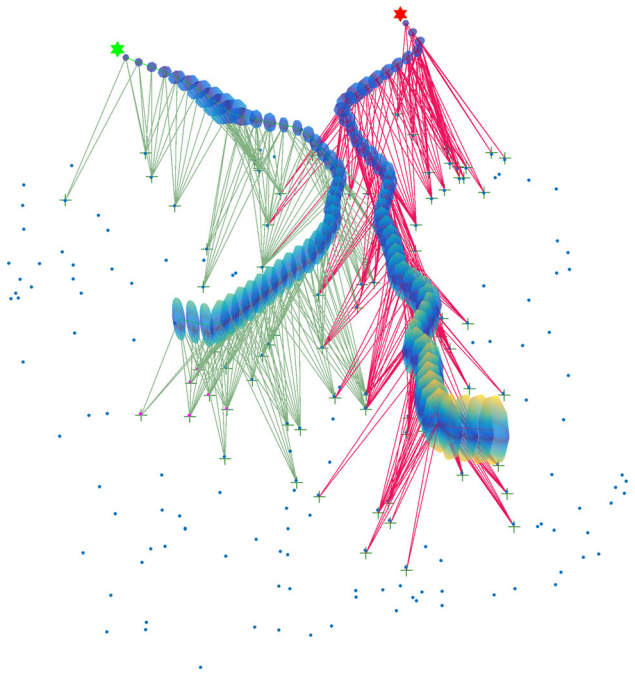

Fig. 2 3D view of the scenario from Fig. 5b: Robots operate in an unknown environment and follow paths generated by PRM that have been identified by the proposed approach to provide the best estimation accuracy upon reaching the goals. One can observe the mutually-observed 3D points that induce indirect multi-robot constraints involving different time instances; these constraints have been accounted for in the planning phase. Robot initial positions are denoted by *asterisks marks* (at the *top* of the figure); uncertainty covariances of robot poses are represented by *ellipsoids*

(29), path length and uncertainty upon reaching the goal, for the candidate paths shown in Fig. 3.

The lowest predicted uncertainty covariances are obtained for candidate paths with identified multi-robot constraints as shown in Fig. 4b. In particular, the predicted uncertainty is reduced by about 40% from 35 m to below 20 m for the first (red) robot. There is a price to pay, however, in terms of path lengths (or time of arrival): as shown in Fig. 4a, to attain these levels of uncertainty, the path of the second (green) robot is not the shortest among the considered candidate paths. The actual decision regarding the best path for each robot depends on the specific values of the weights κ_{path} and κ_{uncert} in the objective function J from Eq. (29), which are to be defined by the user.

Next, we consider actual performance while navigating to pre-defined goals in unknown environments using as controls the identified robot paths in the planning phase described above. The results are shown in Fig. 5 for two alternatives from Fig. 3a, c. Only the latter included multi-robot constraints within planning. One can observe that also in practice, using controls from Configuration C drives the robots sufficiently close to make mutual observations of 3D points (that were unknown at planning time) and as a result

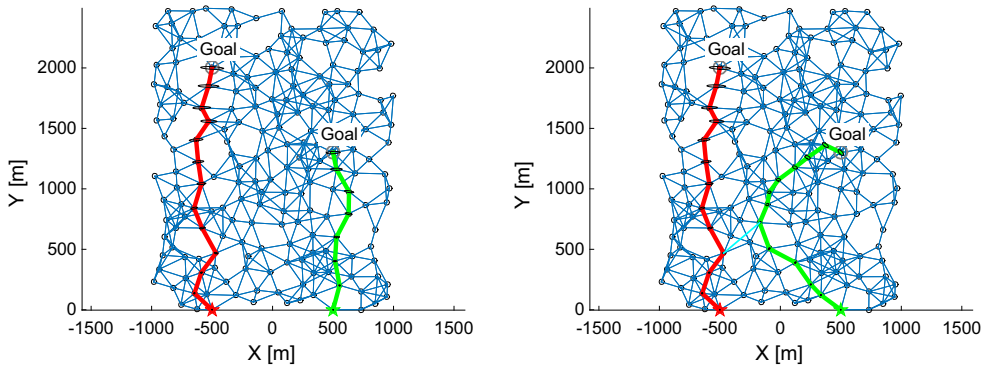

(a) Configuration A **(b)** Configuration B

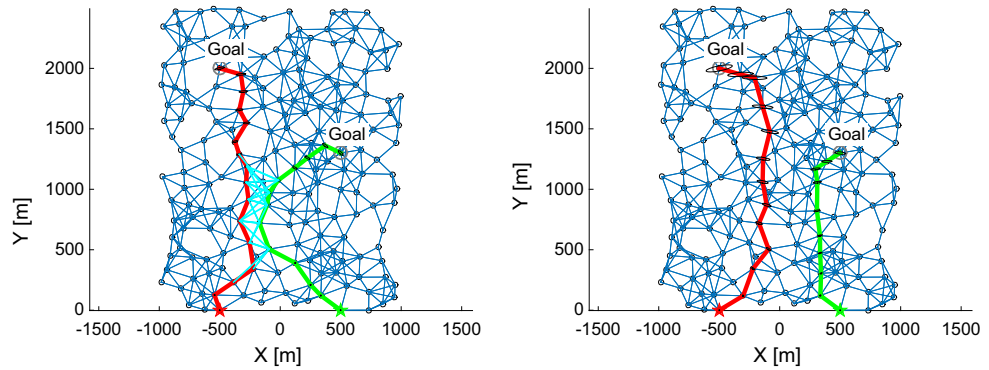

(c) Configuration C **(d)** Configuration D

Fig. 3 Different candidate paths for *red* and *green* robots calculated over a PRM. Robot initial positions are denoted by *asterisks marks*; each robot has to navigate to a different goal, while operating in an unknown environment. The figures show the covariance evolution along each path. Multi-robot constraints have been incorporated (denoted by *cyan color*) whenever robot poses are sufficiently close, which happens mainly in **c**; as a result, uncertainty covariances are drastically reduced. Note these constraints involve different *future* time instances. Covariances were artificially inflated by a constant factor for visualization—actual values are shown in Fig. 4. **a** Configuration A. **b** Configuration B. **c** Configuration C. **d** Configuration D. (Color figure online)

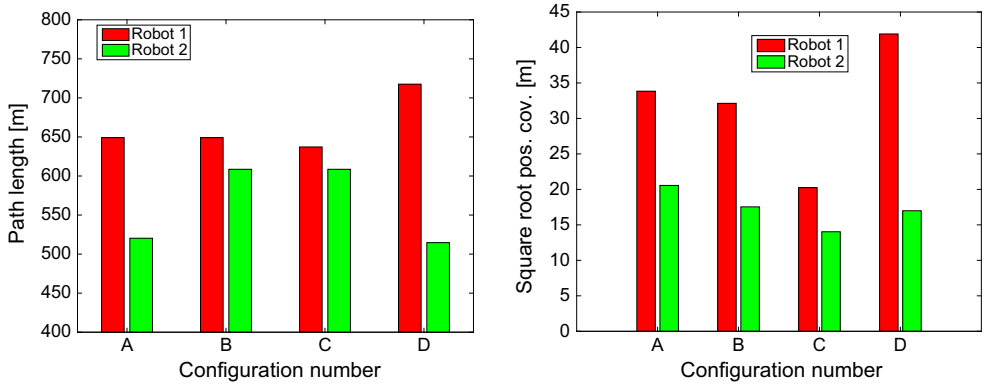

(a) Path length **(b)** Trace of square root final position covariance

Fig. 4 Quantitative comparison between the four alternatives shown in Fig. 3: **a** path length; **b** covariance upon reaching the goals. Multi-robot constraints lead to lowest predicted uncertainty represented by configuration C from Fig. 3c

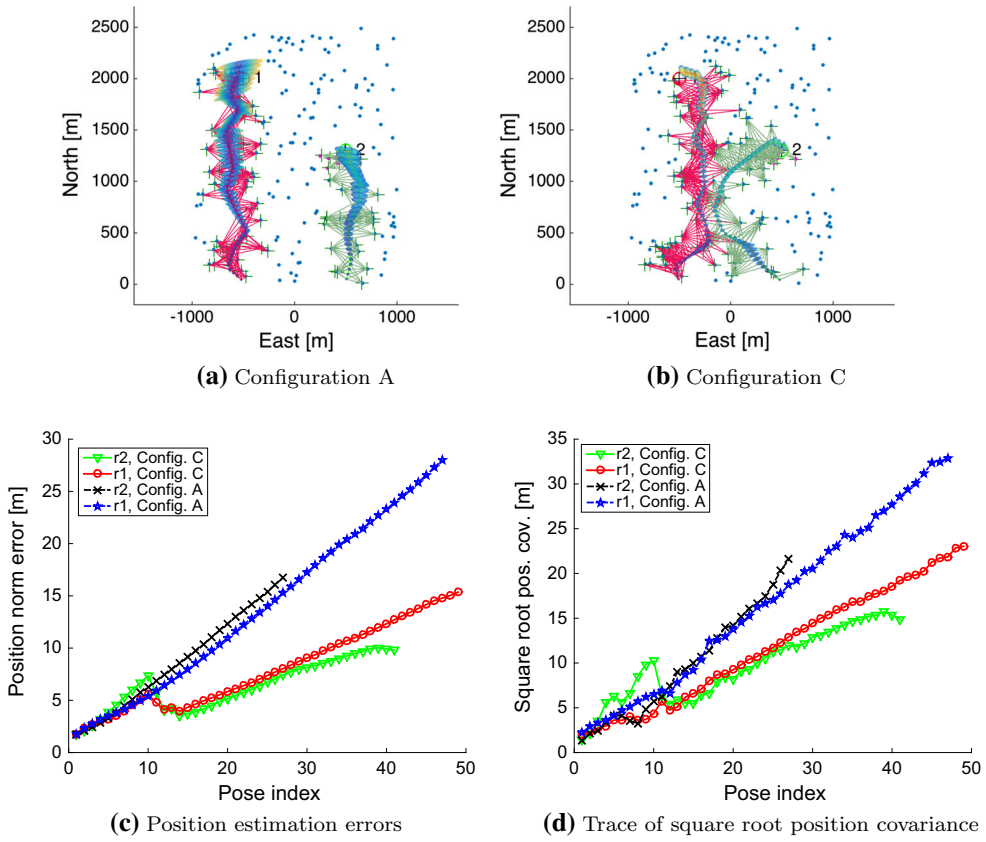

(a) Configuration A **(b)** Configuration C

(c) Position estimation errors **(d)** Trace of square root position covariance

Fig. 5 Autonomous navigation to goals according to identified robot paths in the planning phase. The environment, represented by a sparse set of landmarks, is initially unknown and only gradually discovered. Figures **a**, **b** show robot trajectories and landmark observations using paths defined, respectively, by Configuration A and C (see Fig. 3). The latter involves numerous mutual observations of landmarks, that induce indirectly multi-robot constraints. A 3D view is also shown in Fig. 2. **c**, **d** Shows the corresponding estimation errors and developing covariance over time, which, in overall, agree with the predicted belief evolution from Fig. 4b.

significantly improve estimation accuracy for both robots (see Figs. 5c, d, and 2 for a 3D view).

4.3 Obstacle-free unknown environment: statistical study

In this section we present a statistical study comprising 50 runs, considering two scenarios in an obstacle-free unknown environment: `Scenario A` and `Scenario B`. These scenarios only differ in the starting position and goal of each robot, as shown in Fig. 6. The scenario in Fig. 6a is identical to the one considered in Sect. 4.2. `Scenario B` (Fig. 6b) represents opportunistic active collaborative estimation, along the lines of Indelman (2015a).

Figure 6 also shows, for both scenarios, 25 candidate paths for each robot. These paths are randomly generated and thus *change* from one run to another. The process and measurement noise in SLAM are randomly drawn from a Gaussian distribution for each run as well, see Sect. 4.1.

First scenario

The results for `Scenario A` (Fig. 6a) are reported in Figs. 7 and 8. These figures compare between the proposed approach that incorporates multi-robot factors within belief space planning to an equivalent approach that does not do so. In particular, Fig. 7a shows for each run the identified best paths by the belief space planning algorithm considering the appropriate set of candidate paths (since the latter change from one run to another). It can be seen that the proposed approach tends to prioritize paths along which multi-robot factors can be generated, since the latter contribute to accuracy at the cost of longer path. In contrast, as seen in the second column of Fig. 7a, without incorporating multi-robot factors within belief space planning generally does not motivate collaborative estimation between the robots—in such a setting, the best path for each given run, are typically far away from each other and thus there will be no mutual observations made by the robots.

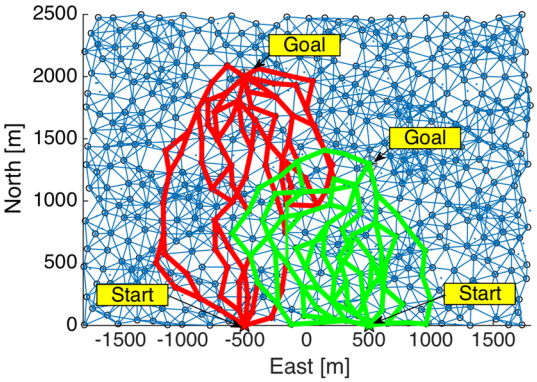

(a) `Scenario A`

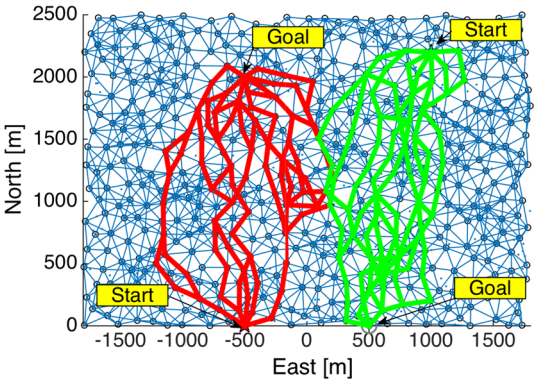

(b) `Scenario B`

Fig. 6 Two considered scenarios in statistical study. Note candidate paths change from one run to another. Shown are candidate paths for one of the runs

The weight parameters $\kappa_{uncertain}$ and κ_{path} (see Eq. (29)), and in particular the ratio $\kappa_{uncertain}/\kappa_{path}$ can be considered as tuning parameters, with higher values of $\kappa_{uncertain}$ assigning greater importance to reducing uncertainty. The latter would typically result in determined best paths of corresponding runs to be closer to each other such that more multi-robot factors can be generated.

Figure 7b, c show the corresponding position estimation errors and square root position covariances as a function of time, as obtained from SLAM (i.e. inference) by executing the best paths identified for each run. As expected, it can be clearly seen that the errors in the first column are typically lower then those in the second column.

The above insights are quantified in Fig. 8, which shows statistics (in terms of box plots) upon reaching the goal. In particular, Fig. 8a, b show statistics within the planning algorithm: position covariance that corresponds to belief at the goal for determined best paths and the corresponding path lengths. The trade-off between accuracy and path length can be clearly seen here—incorporating multi-robot factors (denoted as MR in the figures) within belief space planning improves accuracy at the cost of somewhat longer path length. Figure 8c reports statistics regarding estimation error at the goal as obtained from SLAM; the beneficial impact of the proposed approach can be seen here. For example, the median estimation error for the first and second robot is reduced from about 30 and 25 to about 20 and 12 m respectively.

Table 1 provides running time statistics for the performed 50 runs in this statistical study. The reported running time refers to a single planning session which is performed at the beginning of each run, and the entire multi-robot SLAM session along the chosen paths (see Fig. 7a). Similar SLAM and per-planning session timing results are also obtained in the obstacle-populated unknown environments considered in Sect. 4.4, which involve multiple planning sessions.

Second scenario

Similar aspects can be also observed in `Scenario B` (Fig. 6b). The results are shown in Figs. 9 and 10. Here, however, the beneficial impact of incorporating multi-robot factors within belief space planning is even more significant. This can be clearly seen by comparing the first and second columns of Fig. 9c; for example, estimation errors of the second robot (green) drop to around 10 m in the proposed approach but are much higher in the alternative approach.

Figure 10 quantifies the above in terms of box plots depicting position covariance, path length and estimation error upon reaching the goals (see more detailed explanation in `Scenario A`). In particular, the median estimation error without using multi-robot factors is about 30 and 45 m for the red and green robots; these are reduced to below 10 m by the propose approach.

4.4 Unknown environment with obstacles

In this section we examine the proposed method considering an unknown obstacle-populated environment. We assume a black-box object detector and consider for simplicity the entire obstacle is detected once it is within the camera field of view. Upon obstacle detection, there is a need to re-calculate the best paths since those paths that were considered as optimal before obstacle detection do not necessarily remain as such once an obstacle is detected.

Clearly, there are numerous efficiency-related aspects one could consider at this point. However, since in this work we focus mainly on the impact of future multi-robot factors on the estimation performance, a naïve approach is considered: in each re-planning session we recalculate the candidate paths for each robot, re-evaluate the objective function J from (29) for all combinations between different robots, and choose the best paths accordingly. In the results reported herein, only two robots are considered, each generating 50

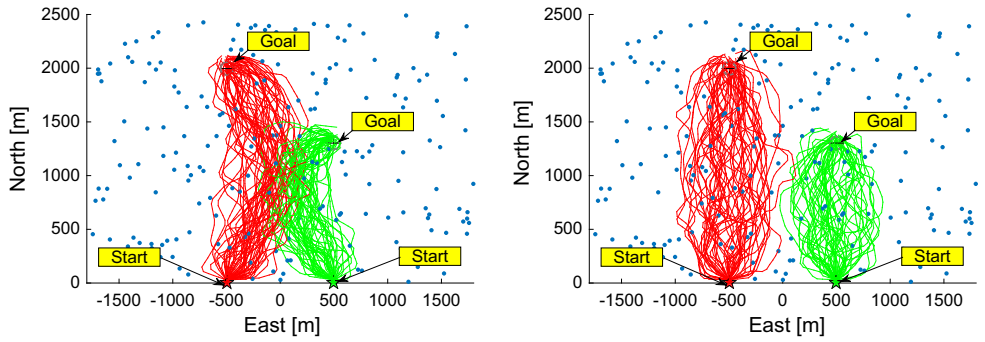

(a) Determined best paths for different runs

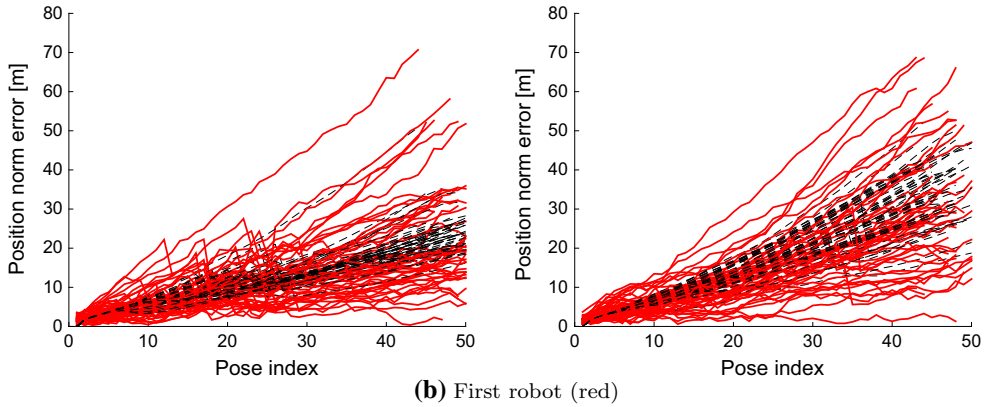

(b) First robot (red)

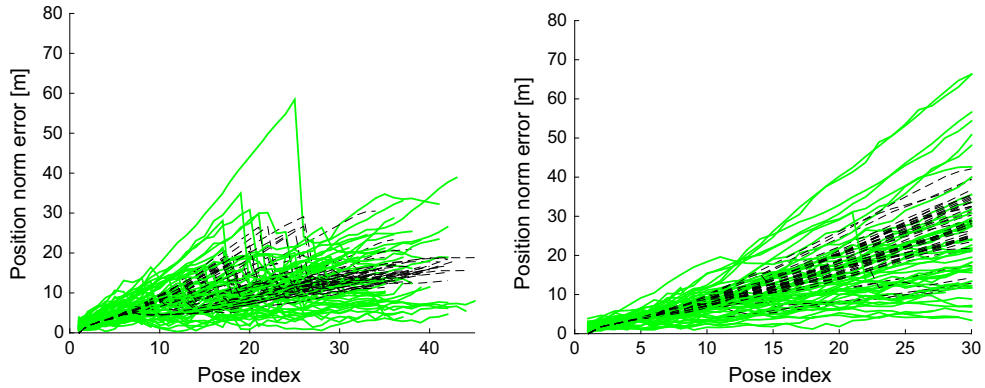

(c) Second robot (green)

Fig. 7 Statistical study of 50 runs for `Scenario A` from Fig. 6a, with (*first column*) and without (*second column*) incorporating multi-robot factors within belief space planning: **a** Determined best paths by the planning approach for each run from appropriate candidate paths (each run has a different set of candidate paths); **b–c** Position estimation error norm (*solid*) and position square root covariance norm (*dashed*) of the first and second robot, as obtained from SLAM upon executing the paths from **a**

candidate paths. Running time for each planning session is typically similar to the results reported in Table 1.

Figure 13 shows a representative scenario where two robots are assigned to navigate to pre-defined goals in an unknown environment populated with landmarks and randomly generated obstacles. In all figures in this section, we will use the following color convention to ease results interpretation: an obstacles is shown in gray if it is not yet detected and in black after detection. Obviously, in the beginning all obstacles are undetected and hence appear as gray.

The figure shows the candidate paths for each re-planning session (first column), the identified best paths for each robot with both multi-robot factors and factors to previously mapped landmarks shown in cyan (second column), see Eq. (17). Additionally, in the third column the figure shows the estimated robot trajectories as obtained from SLAM. To

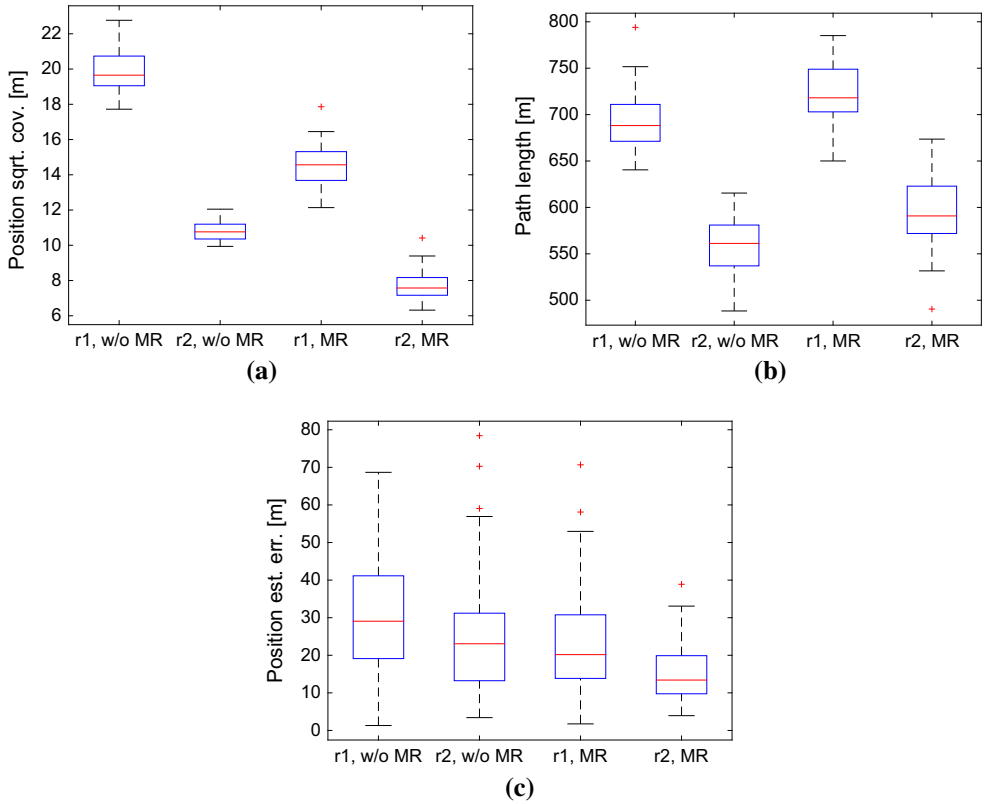

Fig. 8 Statistical study of 50 runs for `Scenario B` from Fig. 6a, with and without incorporating multi-robot (MR) factors within belief space planning: **a** box plots of the square root position covariance, calculated within belief space planning, upon reaching the goal of the chosen best paths; **b** box plots of the corresponding path lengths; **c** box plots of SLAM estimation errors at the goal of each robot

reduce clutter, landmark estimations are not shown. One can see that some of the candidate paths go through the undetected obstacles (gray color), while once an obstacle is detected the newly calculated paths avoid the obstacle.

As expected, overall, the planner typically chooses as best paths those path combinations that admit multi-robot factors (see, e.g., Fig. 11b). The corresponding belief uncertainty in such combinations is smaller compared to those path candidates without multi-robot factors. Similarly to the discussion in Sect. 4.2, this often comes with a price of larger path lengths, since the robots typically need to deviate from the individual shortest paths to facilitate the mentioned multi-robot factors. The actual decision regarding the best path for each robot depends on the specific values of the weights κ_{path} and κ_{uncert}.

To evaluate the impact of incorporating multi-robot factors within belief space planning, as proposed herein, we also consider performance without using such a multi-robot collaboration within belief space planning. Figure 12 shows the corresponding results in this case. As seen, the determined best paths are different—since there is no incentive to move the robots into positions from which the same (unknown) environment could be mutually observed, the chosen paths typically correspond to shortest paths since the uncertainty develops rapidly with path length. Although the same is also true in the proposed approach, multi-robot factors constrain the uncertainty development, often making longer paths with multi-robot factors more attractive than shorter paths without multi-robot factors. Thus, without incorporating multi-robot factors the evolving uncertainty at the goal is significantly higher—compare, for example, the results in the $3rd$ re-planning session (Fig. 11h vs. Fig. 12h).

Figure 13 presents a quantitative comparison between the two methods (with and without incorporating multi-robot factors). In particular, Fig. 13a shows the uncertainty at the

Table 1 Running time in a statistical study of 50 runs for `Scenario A`.

	Mean [sec]	Min [sec]	Max [sec]
BSP	78.5	73.1	89.4
SLAM	25.9	21.9	30.3

Multi-robot belief space planning (the proposed approach) is performed once at the beginning. The presented SLAM running time refers to performing multi-robot SLAM along chosen paths for each run

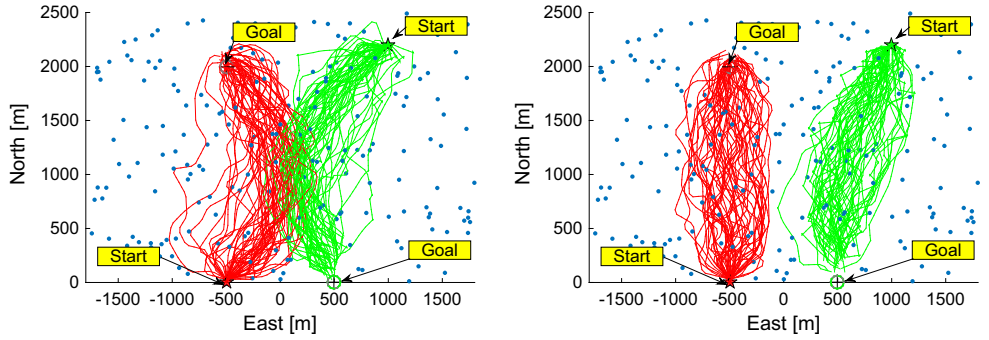

(a) Determined best paths for different runs

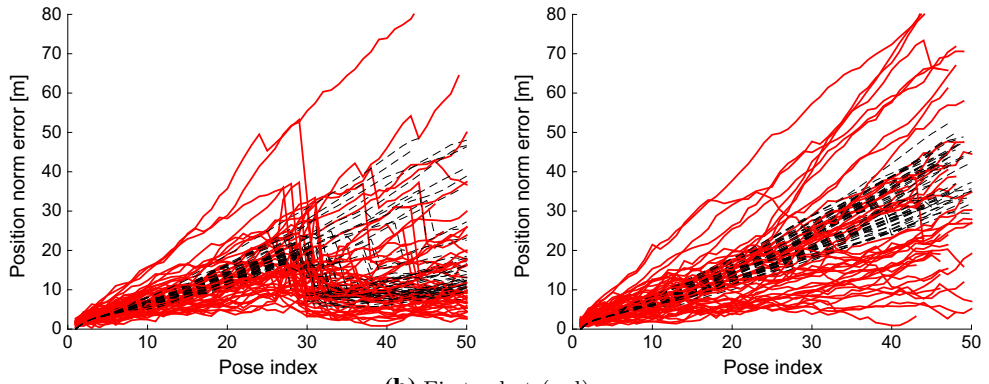

(b) First robot (red)

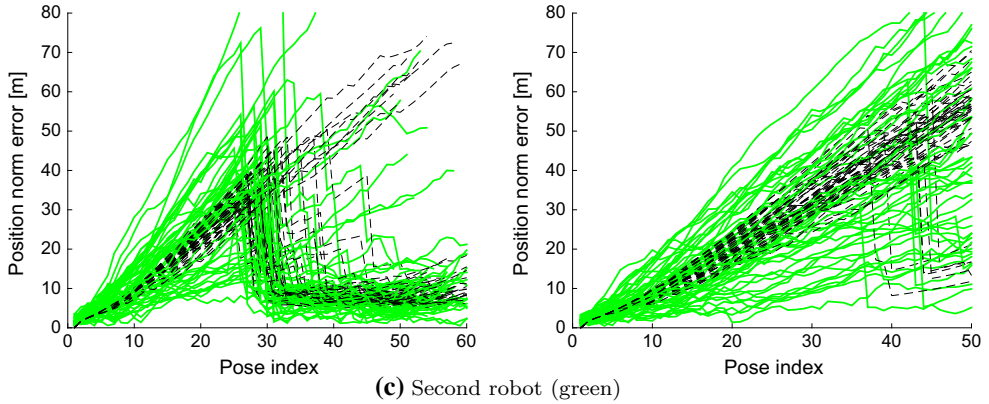

(c) Second robot (green)

Fig. 9 Statistical study of 50 runs for `Scenario B` from Fig. 6b, with and without incorporating multi-robot factors within belief space planning: **a** determined best paths by the planning approach for each run from appropriate candidate paths (each run has a different set of candidate paths); **b**, **c** Position estimation error norm (*solid*) and position square root covariance norm (*dashed*) of the first and second robot, as obtained from SLAM upon executing the paths from **a**

goal for the identified best paths for each robot in each re-planning session. As mentioned above, one can clearly see uncertainty is significantly reduced, often by a factor of two, when using the proposed formulation with multi-robot factors—for example, the uncertainty at the goal for r_1 and r_2 in the final re-planning phase is about 25 and 20 m without using multi-robot factors, and is reduced, respectively, to only about 15 and 10 m.

Since the proposed method facilitates better paths in terms of estimation quality, it is not surprise that the inference process (SLAM) indeed performs better when using such paths over those generated without incorporating multi-robot factors within belief space planning. This can be clearly seen from Fig. 13b that provides the uncertainty covariance in both cases for each robot. Thus, the covariance of r_1 upon reaching the goal is reduced from about 25 to 15 m, at the cost of a longer path.

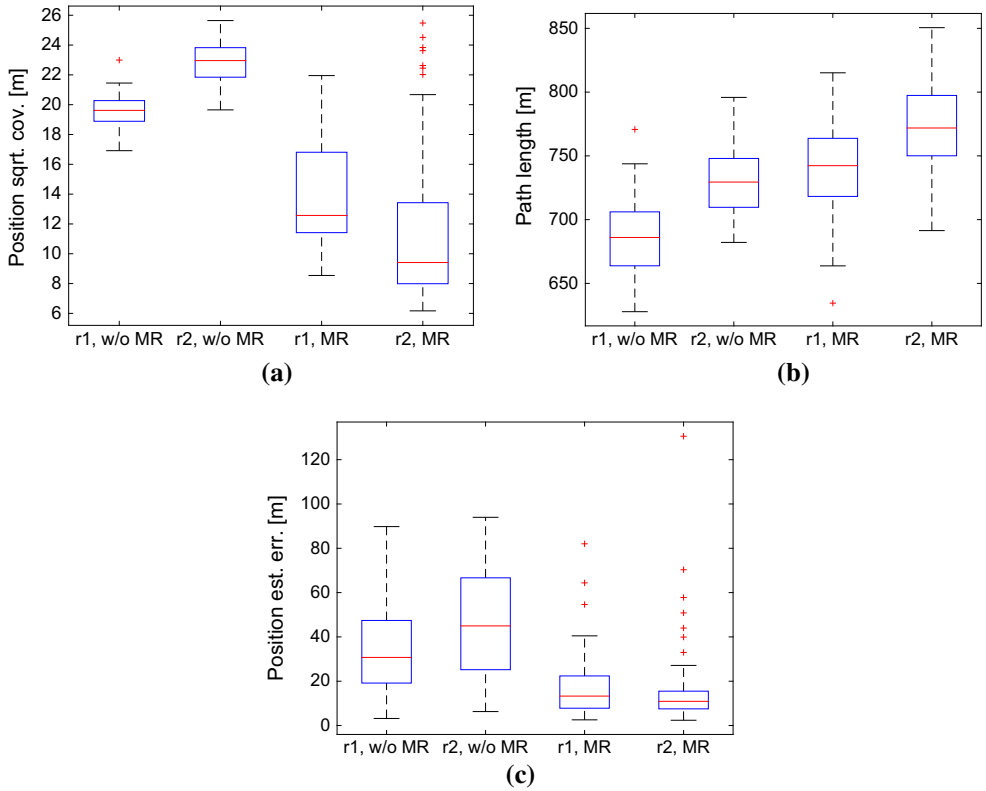

Fig. 10 Statistical study of 50 runs for `Scenario B` from Fig. 6b, with and without incorporating multi-robot (MR) factors within belief space planning: **a** box plots of the square root position covariance, calculated within belief space planning, upon reaching the goal of the chosen best paths; **b** box plots of the corresponding path lengths; **c** box plots of SLAM estimation errors at the goal of each robot

As an additional "sanity" check, one can also see there is a good agreement between the uncertainty predicted in the planning phase and the actual evolving uncertainty, even though the former does not incorporate actual landmark observations from unmapped areas and instead abstracts these using the mentioned multi-robot factors (see Eq. (17)).

5 Conclusions

We presented an approach for collaborative multi-robot belief space planning while operating in unknown environments. Our approach advances the state of the art in belief space planning by reasoning about observations of environments that are unknown at planning time. The key idea is to incorporate within the belief constraints that represent multi-robot observations of unknown mutual environments. These constraints can involve different future time instances, thereby providing enhanced flexibility to the group as rendezvous are no longer necessary. The corresponding formulation facilitates an active collaborative state estimation framework. Given candidate robot actions or paths, it allows to determine best paths according to a user-defined objective function, while modeling future multi-robot interaction and its impact on the belief evolution. Candidate robot paths can be generated by existing motion planning algorithms, and most promising candidates could be further refined into locally optimal solutions using direct trajectory optimization approaches. The approach has been also trivially adapted to support autonomous operation in unknown obstacle-populated environments, where upon detection of an obstacle, the proposed approach is triggered to re-calculate the best paths since previously identified best paths not necessarily remain optimal given the newly detected obstacle. We examined the performance of our approach in simulation, considering as application autonomous navigation to pre-defined goals within unknown obstacle-free and obstacle-populated environments. Simulation results demonstrate estimation performance is significantly improved as a consequence of identifying more informative robot actions. While scenario-dependent, these results show uncertainty covariance could be reduced by more than 50%, while operating in GPS-deprived unknown environments.

There are a number of possible directions for future research to improve the proposed approach. The current formulation involves exhaustive evaluation of all candidate

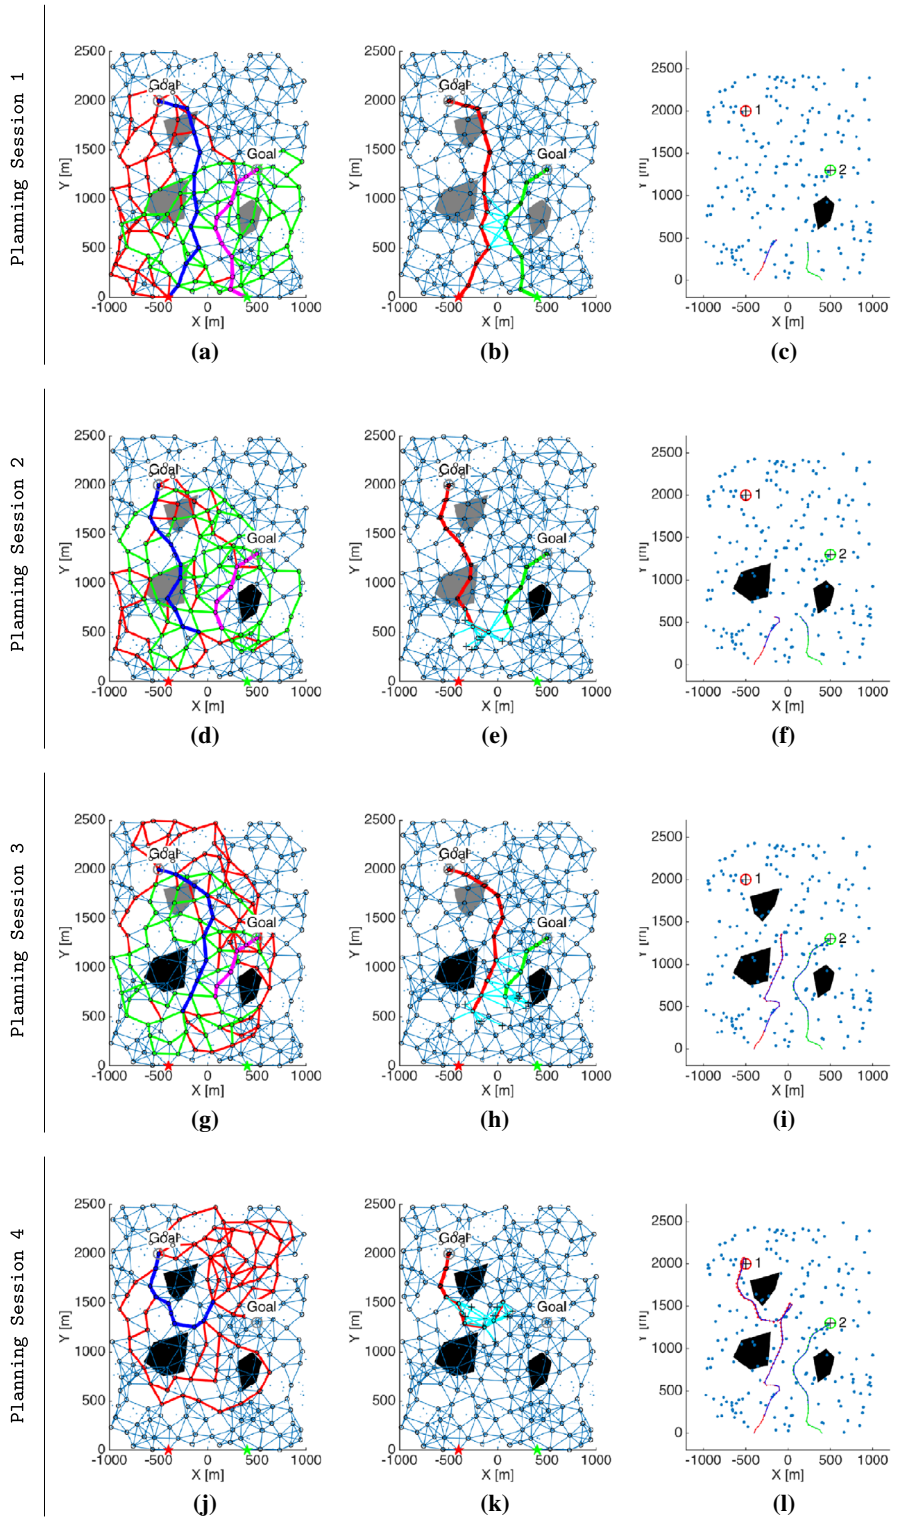

Fig. 11 Performance *with* incorporating multi-robot factors within belief space planning. Each row represents a different re-planning session with detected and undetected obstacles shown in *black* and *gray*, respectively. *First column* shows all candidate paths for both robots on top of PRM, and the identified best paths. *Second column* shows belief evolution for determined best paths, along with multi-robot factors and factors to previously mapped landmarks (e.g. igure **e**). *Third column* shows SLAM solution before the next re-planning phase that is triggered by obstacle detection. Figure (l) shows the final SLAM solution upon reaching both goals. To reduce clutter, mapped landmarks are not explicitly shown. See quantitative evaluation in Fig. 13

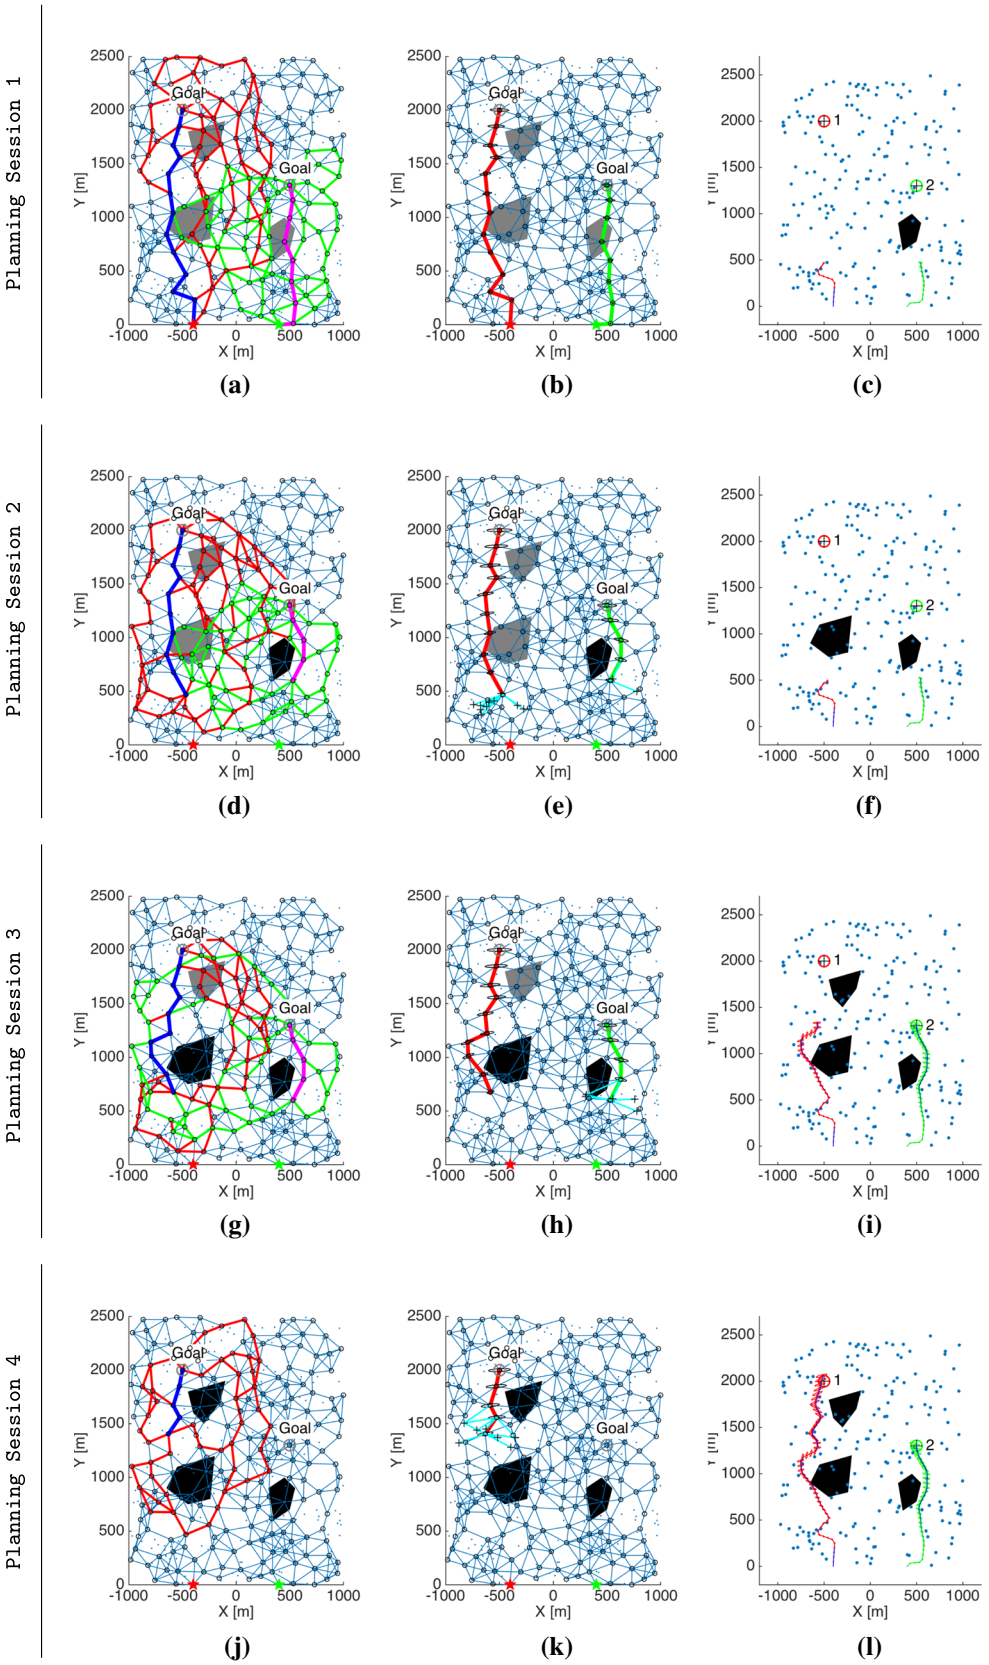

Fig. 12 Performance *without* incorporating multi-robot factors within belief space planning. Refer to caption of Fig. 11 for an explanation regarding the shown sub-figures. See quantitative evaluation in Fig. 13

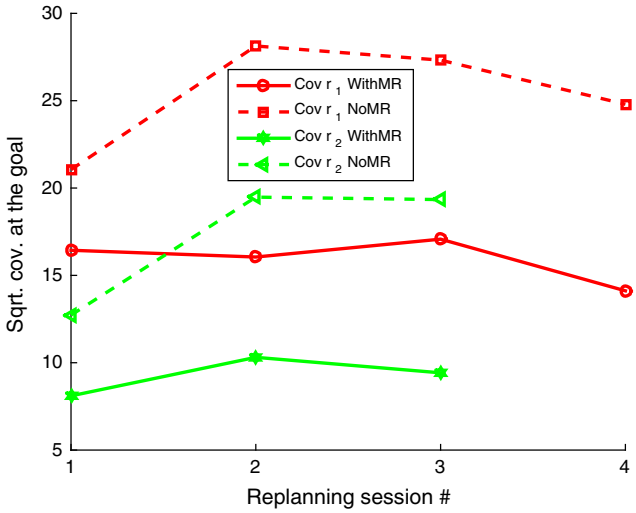

(a) Belief space planning performance

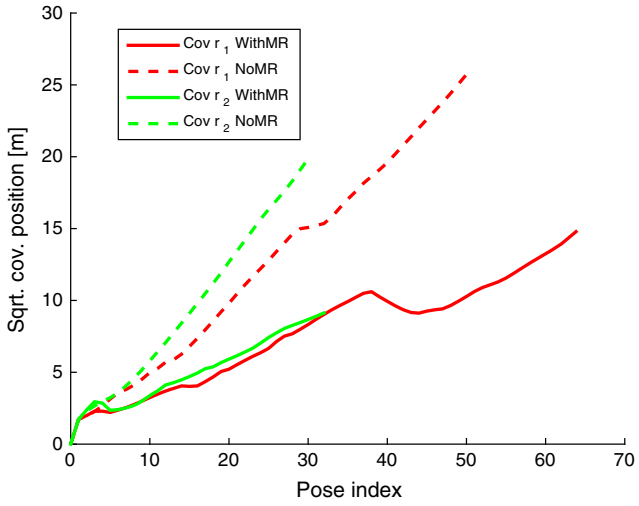

(b) Inference (SLAM) performance

Fig. 13 Quantitative comparison of the proposed method with an alternative that does not incorporate multi-robot collaboration aspects within belief space planning, considering in both cases operation in obstacle-populated unknown environments. See Figs. 11 and 12 for further details. **a** Predicted covariance of the *belief* at the goal Σ^r_{goal} for each robot r using determined best paths for each re-planning phase. **b** *Inference* performance using paths determined by belief space planning with and without incorporating multi-robot collaboration. It is evident that considering multi-robot collaboration within belief space planning, as proposed herein, significantly improves overall estimation performance

paths for all robots, and as such scales poorly. Thus, while the concept of reasoning about collaborative aspects also within unknown environments is general, additional, possibly approximate, approaches will have to be considered for improved scalability. We note such approaches addressing related problems have been actively developed in the last few years (e.g. Levine et al. 2013; Atanasov et al. 2015), including our recent work (Regev and Indelman 2016). Another interesting direction for future research is to consider decentralized and distributed multi-robot frameworks, as well as biasing the sampling according to on-line learned landmarks distribution. Finally, evaluating the method in more realistic simulations and in real-world experiments is clearly also an important avenue for future research.

Acknowledgements This work was partially supported by the Technion Autonomous Systems Program.

References

Agha-Mohammadi, A.-A., Chakravorty, S., & Amato, N. M. (2014). Firm: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements. *The International Journal of Robotics Research.*, *33*(2), 268–304.

Amato, C., Konidaris, G. D., Cruz, G., Maynor, C. A., How, J. P., & Kaelbling, L. P. (2014). *Planning for decentralized control of multiple robots under uncertainty*. arXiv preprint arXiv:1402.2871.

Atanasov, N., Le Ny, J., Daniilidis, K., & Pappas, G. J. (2015). Decentralized active information acquisition: Theory and application to multi-robot slam. In *IEEE international conference on robotics and automation (ICRA)*.

Bry, A., & Roy, N. (2011). Rapidly-exploring random belief trees for motion planning under uncertainty. In *IEEE international conference on robotics and automation (ICRA)*, pp. 723–730.

Burgard, W., Moors, M., Stachniss, C., & Schneider, F. (2005). Coordinated multi-robot exploration. *IEEE Transactions on Robotics.*, *21*(3), 376–386.

Carlone, L., Kaouk Ng, M., Du, J., Bona, B., & Indri, M. (2010). Rao-blackwellized particle filters multi robot SLAM with unknown initial correspondences and limited communication. In *IEEE international conference on robotics and automation (ICRA)*, pp. 243–249. doi:10.1109/ROBOT.2010.5509307.

Chaves, S. M., Kim, A., & Eustice, R. M. (2014). Opportunistic sampling-based planning for active visual slam. In *IEEE/RSJ interantional conference on intelligent robots and systems (IROS)*, IEEE, pp. 3073–3080.

Dellaert, F. (September 2012). Factor graphs and GTSAM: A hands-on introduction. *Technical Report GT-RIM-CP&R-2012-002*, Georgia Institute of Technology.

Eustice, R. M., Singh, H., & Leonard, J. J. (2006). Exactly sparse delayed-state filters for view-based SLAM. *IEEE Transactions on Robotics*, *22*(6), 1100–1114.

Hartley, R. I., & Zisserman, A. (2004). *Multiple view geometry in computer vision* (2nd ed.). Cambridge: Cambridge University Press.

He, R., Prentice, S., & Roy, N. (2008). Planning in information space for a quadrotor helicopter in a gps-denied environment. In *IEEE international conference on robotics and automation (ICRA)*, pp. 1814–1820.

Hollinger, G. A., & Sukhatme, G. S. (2014). Sampling-based robotic information gathering algorithms. *International Journal of Robotics Research*, *33*(9), 1271–1287.

Indelman, V. (September 2015a). Towards multi-robot active collaborative state estimation via belief space planning. In *IEEE/RSJ international conference on intelligent robots and systems (IROS)*.

Indelman, V. (September 2015b). Towards cooperative multi-robot belief space planning in unknown environments. In *Proceedings of the international symposium of robotics research (ISRR)*.

Indelman, V., & Dellaert, F. (2015). Incremental light bundle adjustment: Probabilistic analysis and application to robotic navigation. In *New development in robot vision, cognitive systems monographs* (Vol. 23, pp. 111–136). Springer: Berlin Heidelberg. ISBN 978-3-662-43858-9. doi:10.1007/978-3-662-43859-6_7.

Indelman, V., Gurfil, P., Rivlin, E., & Rotstein, H. (2012). Distributed vision-aided cooperative localization and navigation based on three-view geometry. *Robotics and Autonomous Systems*, *60*(6), 822–840.

Indelman, V., Carlone, L., & Dellaert, F. (December 2013). Towards planning in generalized belief space. In *The 16th international symposium on robotics research*, Singapore.

Indelman, V., Nelson, E., Michael, N., & Dellaert, F. (2014). Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization. In *IEEE international conference on robotics and automation (ICRA)*.

Indelman, V., Carlone, L., & Dellaert, F. (2015a). Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments. *Intnernational Journal of Robotics Research*, *34*(7), 849–882.

Indelman, V., Roberts, R., & Dellaert, F. (2015b). Incremental light bundle adjustment for structure from motion and robotics. *Robotics and Autonomous Systems*, *70*, 63–82.

Indelman, V., Nelson, E., Dong, J., Michael, N., & Dellaert, F. (2016). Incremental distributed inference from arbitrary poses and unknown data association: Using collaborating robots to establish a common reference. *IEEE Control Systems Magazine (CSM), Special Issue on Distributed Control and Estimation for Robotic Vehicle Networks*, *36*(2), 41–74.

Kaess, M., Johannsson, H., Roberts, R., Ila, V., Leonard, J., & Dellaert, F. (2012). iSAM2: Incremental smoothing and mapping using the Bayes tree. *International Journal of Robotics Research*, *31*, 217–236.

Karaman, S., & Frazzoli, E. (2011). Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, *30*(7), 846–894.

Kavraki, L. E., Svestka, P., Latombe, J.-C., & Overmars, M. H. (1996). Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, *12*(4), 566–580.

Kim, A., & Eustice, R. M. (2014). Active visual slam for robotic area coverage: Theory and experiment. *International Journal of Robotics Research.*, *34*(4–5), 457–475.

Koenig, Sven, & Likhachev, Maxim. (2005). Fast replanning for navigation in unknown terrain. *IEEE Transactions on Robotics*, *21*(3), 354–363.

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

LaValle, S. M., & Kuffner, J. J. (2001). Randomized kinodynamic planning. *International Journal of Robotics Research*, *20*(5), 378–400.

Levine, D., Luders, B., & How, J. P. (2013). Information-theoretic motion planning for constrained sensor networks. *Journal of Aerospace Information Systems*, *10*(10), 476–496.

Lu, F., & Milios, E. (Apr 1997). Globally consistent range scan alignment for environment mapping. *Autonomous robots*, pp. 333–349.

Papadimitriou, C., & Tsitsiklis, J. (1987). The complexity of markov decision processes. *Mathematics of Operations Research*, *12*(3), 441–450.

Pathak, S., Thomas, A., Feniger, A., & Indelman, V. (2016a). Robust active perception via data-association aware belief space planning. *arXiv preprint*: arXiv.1606.05124.

Pathak, S., Thomas, A., Feniger, A., & Indelman, V. (September 2016b). Da-bsp: Towards data association aware belief space planning for robust active perception. In *Europian conference on AI (ECAI)*. Accepted.

Patil, S., Kahn, G., Laskey, M., Schulman, J., Goldberg, K., & Abbeel, P. (2014). *Scaling up gaussian belief space planning through covariance-free trajectory optimization and automatic differentiation*. In *International workshop on the algorithmic foundations of robotics*.

Pineau, J., Gordon, G. J., & Thrun, S. (2006). Anytime point-based approximations for large pomdps. *Journal of Artificial Intelligence Research*, *27*, 335–380.

Platt, R., Tedrake, R., Kaelbling, L. P., & Lozano-Pérez, T. (2010). Belief space planning assuming maximum likelihood observations. *Robotics: science and systems (RSS)* (pp. 587–593). Spain: Zaragoza.

Prentice, S., & Roy, N. (2009). The belief roadmap: Efficient planning in belief space by factoring the covariance. *International Journal of Robotics Research.*, *28*(11–12), 1448–1465.

Regev, T., & Indelman, V. (2016). Multi-robot decentralized belief space planning in unknown environments via efficient re-evaluation of impacted paths. In *IEEE/RSJ international conference on intelligent robots and systems (IROS)*, accepted.

Ross, Stéphane, Pineau, Joelle, Paquet, Sébastien, & Chaib-Draa, Brahim. (2008). Online planning algorithms for pomdps. *Journal of Artificial Intelligence Research*, *32*, 663–704.

Roumeliotis, S. I., & Bekey, G. A. (August 2002). *Distributed multi-robot localization*. In *IEEE transactions on robotics and automation*.

Silver, David, & Veness, Joel (2010). Monte-carlo planning in large pomdps. In *Advances in neural information processing systems (NIPS)*, pp. 2164–2172.

Stachniss, C., Grisetti, G., & Burgard, W. (2005). *Information gain-based exploration using rao-blackwellized particle filters*. In *Robotics: science and Systems (RSS)*, pp. 65–72.

Thrun, S., Burgard, W., & Fox, D. (2005). *Probabilistic Robotics*. Cambridge: The MIT press.

Valencia, R., Morta, M., Andrade-Cetto, J., & Porta, J. M. (2013). Planning reliable paths with pose SLAM. *IEEE Transactions on Robotics*, *29*(4), 1050–1059.

Van Den Berg, J., Patil, S., & Alterovitz, R. (2012). Motion planning under uncertainty using iterative local optimization in belief space. *International Journal of Robotics Research*, *31*(11), 1263–1278.

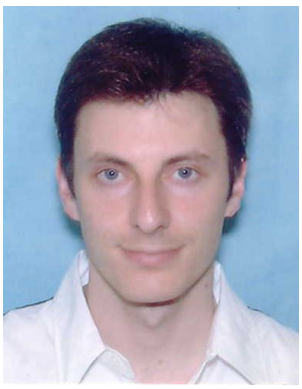

Vadim Indelman obtained his Ph.D. degree in Aerospace Engineering at the Technion - Israel Institute of Technology in 2011. He also holds B.A. and B.Sc. degrees in Computer Science and Aerospace Engineering, respectively, both awarded by the Technion in 2002. Between 2012 and 2014, Dr. Indelman was a postdoctoral fellow in the Institute of Robotics and Intelligent Machines (IRIM) at the Georgia Institute of Technology. In July 2014 he assumed his present position of an Assistant Professor in Aerospace Engineering at the Technion. Dr. Indelman is also a member of the Technion Autonomous Systems Program (TASP). His research interests include autonomous navigation and mapping, distributed perception and information fusion, belief-space planning and active sensing, visionaided navigation (VAN) and simultaneous localization and mapping (SLAM).