



General-purpose incremental covariance update and efficient belief space planning via a factor-graph propagation action tree

The International Journal of
Robotics Research
1–30
© The Author(s) 2019
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0278364919875199
journals.sagepub.com/home/ijr


Dmitry Kopitkov¹ and Vadim Indelman²

Abstract

Fast covariance calculation is required both for simultaneous localization and mapping (SLAM; e.g., in order to solve data association) and for evaluating the information-theoretic term for different candidate actions in belief space planning (BSP). In this article, we make two primary contributions. First, we develop a novel general-purpose incremental covariance update technique, which efficiently recovers specific covariance entries after any change in probabilistic inference, such as the introduction of new observations/variables or relinearization. Our approach is shown to recover them faster than other state-of-the-art methods. Second, we present a computationally efficient approach for BSP in high-dimensional state spaces, leveraging our incremental covariance update method. State-of-the-art BSP approaches perform belief propagation for each candidate action and then evaluate an objective function that typically includes an information-theoretic term, such as entropy or information gain. Yet, candidate actions often have similar parts (e.g., common trajectory parts), which are however evaluated separately for each candidate. Moreover, calculating the information-theoretic term involves a costly determinant computation of the entire information (covariance) matrix, which is $O(n^3)$ with n being dimension of the state or costly Schur complement operations if only marginal posterior covariance of certain variables is of interest. Our approach, rAMDTree, extends our previous BSP method rAMD, by exploiting incremental covariance calculation and performing calculation reuse between common parts of non-myopic candidate actions, such that these parts are evaluated only once, in contrast to existing approaches. To that end, we represent all candidate actions together in a single unified graphical model, which we introduce and call a factor-graph propagation (FGP) action tree. Each arrow (edge) of the FGP action tree represents a sub-action of one (or more) candidate action sequence(s) and in order to evaluate its information impact we require specific covariance entries of an intermediate belief represented by the tree's vertex from which the edge is coming out (e.g., tail of the arrow). Overall, our approach has only a one-time calculation that depends on n , while evaluating action impact does not depend on n . We perform a careful examination of our approaches in simulation, considering the problem of autonomous navigation in unknown environments, where rAMDTree shows superior performance compared with rAMD, while determining the same best actions.

Keywords

Covariance recovery, belief space planning, active SLAM, informative planning, active inference, autonomous navigation

1. Introduction

Autonomous operation in unknown or uncertain environments is a fundamental problem in robotics and is an essential part in numerous applications such as autonomous navigation in unknown environments, target tracking, search-and-rescue scenarios, and autonomous manufacturing. It requires both computationally efficient inference and planning approaches, where the former is responsible for tracking the posterior probability distribution function given available data thus far, and the latter is dealing with finding the optimal action given that distribution and a task-specific objective. Since the state is unknown and only

partially observable, planning is performed in the belief space, where each instance is a distribution over the original state, while accounting for different sources of uncertainty.

¹Technion Autonomous Systems Program (TASP), Technion - Israel Institute of Technology, Haifa, Israel

²Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa, Israel

Corresponding author:

Dmitry Kopitkov, Technion Autonomous Systems Program (TASP), Technion - Israel Institute of Technology, Haifa 32000, Israel.
Email: dimkak@technion.ac.il

Such a problem can be naturally viewed as a partially observable Markov decision process (POMDP), which was shown to be computationally intractable and typically is solved by approximate approaches. The planning and decision-making problems are challenging both theoretically and computationally. First, we need to accurately model future belief as a function of future action while considering probabilistic aspects of state sensing. Second, we need to be able to efficiently evaluate utility of this future belief and to find an optimal action, and to do so on-line.

The utility function in belief space planning (BSP) typically involves an information-theoretic term, which quantifies the posterior uncertainty of various variables within a future belief. This, in turn, requires access to the marginal covariance (information) matrix of appropriate variables whose uncertainty we wish to account for (Kopitkov and Indelman, 2017). Similarly, covariance of specific variables is also required in the inference phase, for example, in the context of data association (Kaess and Dellaert, 2009). However, the recovery of specific covariances is computationally expensive in high-dimensional state spaces: while the belief is typically represented in the (square-root) information form to admit computationally efficient updates (Kaess et al., 2012), retrieving the covariance entries requires an inverse of the corresponding (potentially) high-dimensional information matrix. Although sophisticated methods exist to efficiently perform such inverse by exploiting the sparsity of the square-root information matrix and by reordering state variables for enhancing such sparsity (Kaess et al., 2012), the overall complexity still is at least quadratic with respect to state dimension (Ila et al., 2015). Moreover, in the case of planning, such computation needs to be performed for each candidate action.

The computational efficiency of the covariance recovery and the planning process is the main point of this paper. We develop a novel method to incrementally update covariance entries after any change of the inference problem, as defined next. Moreover, we present a planning algorithm which leverages the key ability of incremental covariance updates and by exploiting action similarity is much faster and yet exact with respect to alternative state-of-the-art techniques.

The inference problem can be represented by a set of currently available observations and state variables whose value we are to infer. For example, in a typical simultaneous localization and mapping (SLAM) scenario these variables are the robot poses along a trajectory and landmarks of the environment, while the observations are motion odometry and projection/range measurements. Covariances of the state variables can change as a result of any change in the inference problem, such as the introduction of new observations or augmentation of the state (e.g., introduction of a new robot pose). Moreover, covariances also depend on the current linearization point (the current estimated mean) of the state vector, which in turn can also change after the introduction of new observations. In this




Fig. 1. Illustration of possible candidate actions during exploration of an unknown environment by an autonomous robot. The robot's current position is marked by a red circle; red lines and green points represent the trajectory and final position of each candidate action, respectively. As can be seen, actions share many mutual parts.

article, we scrupulously analyze each such possible change in the inference problem and show how covariance entries can be appropriately incrementally updated. Such capability to incrementally update covariance entries is important not only for the inference phase but also for efficiently addressing information-theoretic BSP, as we describe next.

BSP is typically solved by determining the best action, given an objective function and a set of (non-myopic) candidate actions, while accounting for different sources of uncertainty. Such an approach requires the evaluation of the utility of each action from a given set of candidate actions. This evaluation is usually done *separately* for each candidate action and typically consists of two stages. First, posterior belief for candidate action is propagated and explicit inference is performed. Second, an application-specific objective function is evaluated given a candidate action and the corresponding posterior belief. Yet, inference over the posterior belief and evaluation of the objective function can be computationally expensive, especially when the original state is high-dimensional, since the complexity of both parts depends on its dimension.

In addition, in many BSP applications candidate (non-myopic) actions are partially overlapping, i.e., have similar parts. For instance, in a building exploration scenario, candidate actions are trajectories to different locations in a building (see Figure 1) that were provided, e.g., by sampling-based motion planning approaches; some of these sampled trajectories will have mutual parts. Typically, these common parts will be evaluated a number of times, as part of the evaluation of each action that shares them. Given that we know what are the similar parts between the different candidate actions, it can significantly reduce runtime complexity if we could handle these similar parts only once.

In this article, we present a technique for reusing computation between the candidate actions and exploitation of actions' similarity, while leveraging the above-mentioned method for incremental covariance updates. We show that such a technique greatly reduces the total decision-making

runtime. Further, we argue that for most cases, explicit inference over the posterior belief is not required. Moreover, evaluation of all candidate actions can be performed very quickly by dividing the overall computation into two computational blocks: first, with complexity dependent on state dimension, yet which is performed only once independently of the number of candidate actions; and second, that computes the objective function for each action efficiently by exploiting candidate actions' similarity, and providing complexity independent of the state dimension. In general, the objective function of BSP can contain multiple terms, such as control cost, distance to goal, and an information-theoretic term (e.g., entropy, information gain (IG), or mutual information). Arguably, in typical settings the control cost and distance to goal can be calculated without explicit inference over the posterior belief, since these terms depend only on the linearization point of the state vector. In this article, we show that the information term also does not require an explicit inference over the posterior belief and that action similarity can be efficiently exploited, concluding that the BSP problem can be solved without performing time-consuming explicit inference over the posterior belief at all.

To that end, we present a new paradigm that represents all candidate (sequence of) actions in a single unified data structure that allows the similarities between candidate actions to be exploited while evaluating the impact of each such action. We refer to this structure as *factor-graph propagation* (FGP) action tree, and show that the developed herein incremental covariance calculation method allows us to compute the IG of the tree's various parts. This, in turn, can be used to efficiently evaluate the information term of different candidate actions while reusing calculations when possible. Combining our recently developed *rAMDL* approach (Kopitkov and Indelman, 2017) with an FGP action tree and incremental covariance update, yields an approach that calculates action impact without explicitly performing inference over the posterior belief, while reusing calculations among different candidate actions.

To summarize, our main contributions in this article are as follows: (a) we develop an incremental covariance update method to calculate specific covariance entries after any change in the inference problem; (b) we introduce the FGP action tree, that represents all candidate actions in a single hierarchical model and allows us to mutually formulate common parts of candidate actions; (c) we apply an incremental covariance update method to calculate covariance entries from intermediate and posterior beliefs within the FGP action tree, where all calculations that depend on state dimension n are accumulated into a single computational block which is performed only once for all actions, with further per-action computations independent of n ; and (d) we combine the FGP action tree graphical model, the incremental covariance update method, and *rAMDL* approach (Kopitkov and Indelman, 2017) to yield a new algorithm, *rAMDL-Tree*, that efficiently solves an information-

theoretic BSP problem while handling candidates' mutual parts only once.

This article is organized as follows. In Section 2, we describe relevant work. In Section 3, we present basic notions and relevant background material on probabilistic inference. Section 4 contains the problem definition. In Section 5, we describe our approaches for incremental covariance recovery (Section 5.1) and information-theoretic BSP problem (Section 5.2). Further, in Section 6 we provide simulation results that emphasize the advantages of the approaches presented herein. Finally, in Section 7, we conclude the discussion about the introduced methods and point out several directions for future research. In addition, we provide an appendix with proofs of several lemmas introduced throughout the paper.

2. Related work

In this section, we discuss the work most relevant to our approach, starting with computationally efficient covariance calculation and then proceeding to state-of-the-art BSP approaches.

2.1. Computationally efficient covariance recovery in high-dimensional state spaces

Fast covariance recovery, under the Gaussian inference setting, is an active research area that has been addressed by several works in recent years. Naïvely calculating an inverse of a high-dimensional information matrix is prohibitively expensive. However, these calculations can be avoided by exploiting sparsity of the square root information matrix, yielding a recursive method to calculate the required entries (Golub and Plemmons, 1980), which has been recently also proposed by Kaess and Dellaert (2009) within their incremental smoothing and mapping solver. Although such a method is faster than a simple inverse of square-root information matrix, the covariances are still calculated from scratch and the complexity depends on state dimension n . Moreover, in order to calculate a specific block of covariance matrix, the recursive approach may still need to calculate the entire covariance matrix (with dimensions $n \times n$), which is undesirable for high-dimensional state spaces.

More recently, Ila et al. (2015) introduced an approach to incrementally update covariances after the inference problem was changed. Given specific prior covariance entries that were calculated in a previous time-step, their approach efficiently calculates covariance deltas to these entries, which comes out to be much faster than the recursive approach from Golub and Plemmons (1980) and Kaess and Dellaert (2009). Although this approach is similar in spirit to our method of incremental covariance update, it is more limited in the following sense. Its theoretical part deals only with the specific scenario where new observations were introduced to the inference problem, without

adding new variables. The latter is a common scenario in SLAM where the state vector is augmented with new variables (e.g., robot poses), and its proper solution is essential for robot navigation missions. Yet, the mathematical formulation in Ila et al. (2015) does not handle such state augmentation case, which actually involves singular matrices that are assumed to be invertible according to the derivation of Ila et al. (2015). Nevertheless, their simulation part suggests that the approach can also be applicable in this case in practice. Moreover, in case of state relinearization, the authors used a recursive method as a fallback and calculate covariances from scratch. In contrast, we present a general approach that is mathematically sound and is capable of dealing with any change in the inference problem, including state augmentation and relinearization. In addition, we show how the presented incremental covariance recovery technique can be used to update conditional covariance entries, which is also not supported in Ila et al. (2015). Finally, even though a limited version of incremental covariance update has been developed (Ila et al., 2015), it was not considered within a BSP problem, which is one of our main contributions in this work.

2.2. BSP

As was mentioned previously, BSP is an instantiation of a POMDP problem. Calculating a globally optimal solution of POMDP is known to be computationally intractable (Kaelbling et al., 1998) in high-dimensional state spaces owing to the curse of dimensionality. Therefore, most of the modern research is focused on approximation methods that solve the planning problem in a sub-optimal form with tractable runtime complexity. These approximation methods can be categorized into those that discretize the state/action/measurement space domains and those that act in continuous spaces. Approaches that perform discretization include sampling (Agha-Mohammadi et al., 2014; Prentice and Roy, 2009), simulation (Stachniss et al., 2005), and point-based value iteration (Pineau et al., 2006) methods. Planning approaches that operate in continuous spaces, often also termed as direct trajectory optimization methods, calculate a locally optimal solution given an initial nominal solution using different optimization techniques such as dynamic programming and gradient descent (Indelman et al., 2015; Patil et al., 2014; Platt et al., 2010; Van Den Berg et al., 2012).

In addition, BSP methods can be separated into those that solve myopic and non-myopic decision making. While myopic approaches, also known as next best view (NBV) approaches in computer vision community (e.g., Dunn and Frahm, 2009; Wenhardt et al., 2007), reason about actions taking the system only one step into the future, non-myopic planning (e.g., He et al., 2011; Indelman et al., 2015; Kim and Eustice, 2014; Platt et al., 2010; Prentice and Roy, 2009; Van Den Berg et al., 2012) deals with sequences of actions taking the system multiple steps into the future. Clearly, for more complex tasks non-myopic methods

will perform better as the time period before receiving the reward can be long. Yet, such methods are typically more computationally expensive as more effort is required to consider different probabilistic outcomes along the long planning horizon. In this article, we consider a non-myopic setting and formulate the problem through factor graphs.

One way to reduce the computational complexity of non-myopic information-theoretic planning is by appropriate factoring of the covariance matrix, as proposed by Prentice and Roy (2009). Such a procedure allows several planning step updates to be combined into a one-stage computation while using *filtering* estimators (e.g., Kalman filter), thus reducing overall computation time. When a *smoothing* formulation is used for planning instead, a non-myopic candidate action can also be evaluated through a one-stage computation, as will be demonstrated in this paper. Further, by exploiting problem's sparsity the presented herein approach solves BSP in a very efficient way.

An information-theoretic BSP problem seeks for an optimal action that maximally reduce the estimation uncertainty of the state vector. Such a problem can be separated into two main cases: *unfocused* BSP tries to reduce uncertainty of all variables inside the state vector, whereas *focused* BSP is only interested in reducing the uncertainty of a predefined subset (termed as *focused* variables) of these variables. Typically, the two cases have different best actions, with optimal action from *unfocused* BSP potentially providing little information about *focused* variables of *focused* BSP (see, e.g., Levine and How, 2013). In both cases, the objective function usually calculates posterior entropy or IG (of all variables from the state vector or of only *focused* variables) and may have high computational complexity that depends on state dimension n . For instance, the calculation of *unfocused* posterior entropy usually requires determinant computation of information (covariance) matrix, which is, in general, $O(n^3)$, and is smaller for sparse matrices as in SLAM problems (Bai et al., 1996). Calculation of *focused* posterior entropy is even more expensive and requires additional Schur complement computation.

Recently, we presented a novel approach, *rAMD*L (Kopitkov and Indelman, 2017), to efficiently calculate entropy and IG for both *focused* and *unfocused* cases. This method requires only one-time calculation that depends on dimension n : computation of specific prior marginal (or conditional) covariances. Given these prior covariances, *rAMD*L evaluates information impact of each candidate action independently of state dimension n . Such a technique was shown to significantly reduce runtime (by orders of magnitude) compared with standard approaches.

Yet, in most BSP approaches, including our own *rAMD*L approach, the similarity between candidate actions is not exploited and each candidate is evaluated from scratch. To the best of the authors' knowledge, only the work by Chaves and Eustice (2016) was done in this

direction. Their approach performs fast explicit inference over the posterior belief, by constraining variable ordering of the Bayes tree data structure to have candidates' common variables eliminated first. Still, this approach has its limitations. It explicitly calculates the posterior belief for each action, and though this calculation is done fast, the method still requires additional memory to store such posterior beliefs. Further, the approach in Chaves and Eustice (2016) does not deal with information-theoretic objective functions whose runtime complexity is usually very expensive, as mentioned above. Moreover, it can only be applied when the SLAM algorithm is implemented using a Bayes tree (Kaess et al., 2012), and it was shown to work only for the case where actions are trajectories constrained to have only a single common part.

In contrast, in this article, we develop a BSP technique that reuses calculations in a general way, by exploiting potentially any number of mutual parts between the candidate actions. It is expressed in terms of factor graphs and can be applied not just for trajectory planning, but for any decision-making problem expressed via factor graphs. Moreover, our technique can be implemented independently of a chosen SLAM factor graph optimization infrastructure. We combine several algorithmic concepts together: a unified graphical model FGP action tree, incremental covariance update, and *rAMD*L approach (Kopitkov and Indelman, 2017), and present a BSP solution that does not require explicit inference over the posterior belief while carefully evaluating information impact of each action in an exact way.

3. Notation and preliminaries

In this article, we address the covariance recovery and BSP problems considering a *smoothing* formulation of probabilistic inference, where the state vector contains all variables of interest without marginalizing old variables out. In general, a probabilistic inference problem can be formulated in terms of factor graphs (Kschischang et al., 2001). Concretely, at time step k it can be represented by a factor graph $G_k = (F_k, X_k, \mathcal{E}_k)$, where X_k and F_k are variable and factor nodes that correspond, respectively, to the estimated state variables and the related probabilistic models. As such, depending on the application, X_k can represent robot configuration and poses (optionally also past and current poses), environment-related variables or any other variables of interest. Further, $F_k = \{f^1(X^1), \dots, f^{n_f}(X^{n_f})\}$ is the set of all factors acquired until time step k (e.g., prior, motion and measurement models), where $X^j \subseteq X_k$ are variables involved in factor f^j . Further, \mathcal{E}_k encodes connectivity according to the variables involved in each factor (see also Figure 2a). Such a formulation is general and can express various inference problems, including SLAM and navigation. The joint probability density function (pdf) over the state vector $X_k \in \mathbb{R}^n$ can be then written in a general form as

$$\mathbb{P}(X_k | \mathcal{H}_k) \propto \prod_{j=1}^{n_f} f^j(X^j), \quad (1)$$

where \mathcal{H}_k is a *history* that contains all the information gathered by the current time (measurements, controls, etc.).

As common in many inference problems, in this work we assume that all factors have a Gaussian form,

$$f^j(X^j) \propto \exp\left(-\frac{1}{2} \|h^j(X^j) - r^j\|_{\Sigma^j}^2\right) \quad (2)$$

with an appropriate model

$$r^j = h^j(X^j) + v^j, \quad v^j \sim \mathcal{N}(0, \Sigma^j) \quad (3)$$

where h^j is a known nonlinear function, v^j is a zero-mean Gaussian noise and r^j is the expected value of h^j , i.e., $r^j = \mathbb{E}[h^j(X^j)]$. Such a factor representation is a general way to express information about the state. In particular, it can represent a measurement model, in which case h^j is the observation model, and r^j and v^j are the actual measurement z and measurement noise, respectively. Similarly, it can also represent a motion model.

As a basic example, consider a typical robot navigation scenario where the joint pdf over robot trajectory can be written as

$$\mathbb{P}(X_k | u_{0:k-1}, z_{1:k}) \propto p(x_0) \prod_{j=1}^k p(x_j | x_{j-1}, u_{j-1}) p(z_j | x_j) \quad (4)$$

where X_k is the state vector at time step k containing all robot poses $\{x_i\}_{i=0}^k$ until now, $p(x_0)$ is a prior over initial pose x_0 , $p(x_j | x_{j-1}, u_{j-1})$ is the motion model and $p(z_j | x_j)$ is the measurement model. As can be seen, Equation (4) is a particular instance of the more general formulation in (1). Note that along time steps the state vector X_k is augmented with new robot poses. Likewise, new measurement and odometry factors are added to the product in (4). Typically, X_k is estimated by maximizing (1), which provides a maximum a posteriori (MAP) estimation of the state vector. We term this MAP solution as *belief* $b[X_k]$. In case all factors are modeled as Gaussian distributions, MAP inference corresponds to solving a nonlinear least squares problem. Its solution then, the belief $b[X_k]$, is itself parametrized by a high-dimensional Gaussian distribution:

$$b[X_k] \doteq \mathbb{P}(X_k | \mathcal{H}_k) = \mathcal{N}(X_k^*, \Sigma_k) = \mathcal{N}^{-1}(\eta_k^*, \Lambda_k) \quad (5)$$

where X_k^* , Σ_k , η_k^* , and Λ_k are the mean vector, covariance matrix, information vector, and information matrix (inverse of covariance matrix), respectively. Recent works developed computationally efficient approaches to calculate the MAP estimate X_k^* incrementally while exploiting sparsity (see, e.g., Dellaert and Kaess, 2006; Ila et al., 2017; Kaess et al., 2012, 2008).

When an inference problem is updated along time (i.e., from time instant k to another time instant $k+1$), new

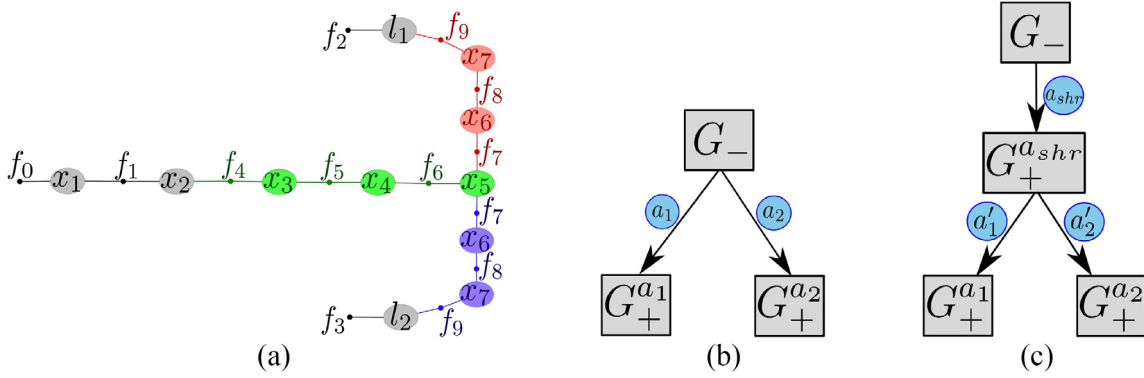


Fig. 2. Illustration of belief propagation in factor graph representation, taken from SLAM application. (a) Prior factor graph G_- (colored in black) contains two robot poses x_1 and x_2 , and two landmarks l_1 and l_2 , as also the prior and motion model factors $\{f_0, f_1, f_2, f_3\}$. Two different actions (trajectories) are considered. The first will take the robot to observe landmark l_1 and will augment the G_- with new factors and state variables colored in green and red. The second will take robot to visit l_2 instead, and will augment the G_- with new factors and state variables colored in green and purple. As can be seen, both candidate actions share some of their new factors/state variables (colored in green). (b), (c) Candidates from (a) represented as an FGP action tree. In an FGP action tree, each vertex represents a specific factor graph (or the belief that is associated with it) and each edge represents a specific action: augmentation of the factor graph with an *increment* of the action, see (12). (b) Posterior factor graphs $G_+^{a_1}$ and $G_+^{a_2}$ are propagated separately for each action. (c) First the prior factor G_- is augmented by a mutual *increment* represented by a_{shr} (colored green in (a)) and the posterior $G_+^{a_{shr}}$ is received. Next, $G_+^{a_{shr}}$ is augmented separately by non-mutual *increments* (colored red and purple in (a)) of each action, providing posterior factor graphs $G_+^{a_1}$ and $G_+^{a_2}$.

factors and state variables are introduced into (1) and factor graph G_k . As a smoothing formulation is considered, propagating the belief $b[X_k]$ to $b[X_{k+l}]$ means performing inference over the updated (1):

$$b[X_{k+l}] \doteq \mathbb{P}(X_{k+l} | \mathcal{H}_{k+l}) \equiv \mathbb{P}(X_k, X_{new} | \mathcal{H}_{k+l}) \propto \left[\prod_{j=1}^{n_f} f^j(X^j) \right] \cdot \left[\prod_i f_{new}^i(X^i) \right] \quad (6)$$

where X_{new} is a set of new state variables with the entire state vector becoming $X_{k+l} = X_k \cup X_{new}$, and where $\{f_{new}^i\}$ are newly introduced factors. The updated belief $b[X_{k+l}]$ is typically parametrized by the updated mean vector X_{k+l}^* and the updated information matrix Λ_{k+l} , or the square-root information upper-triangular matrix R_{k+l} :

$$b[X_{k+l}] \doteq \mathbb{P}(X_{k+l} | \mathcal{H}_{k+l}) = \mathcal{N}(X_{k+l}^*, (\Lambda_{k+l})^{-1}) = \mathcal{N}(X_{k+l}^*, (R_{k+l}^\top \cdot R_{k+l})^{-1}) \quad (7)$$

Solving for $b[X_{k+l}]$ involves inferring *belief* over both new and old state variables, where the *belief* over the latter in general also changes (i.e., $\mathbb{P}(X_k | \mathcal{H}_k) \neq \mathbb{P}(X_k | \mathcal{H}_{k+l})$). This may happen, for example, when adding newly acquired sensor measurements in inference, or in planning, when considering a specific candidate action and the corresponding future posterior belief.

In this article, we explore how the belief changes after the factor graph G_k is modified as discussed previously. Specifically, we provide methods to incrementally infer information-based properties of the *belief* after G_k changes, such as recovery of posterior covariance entries and computation of an IG.

Furthermore, in this article we also deal with non-myopic planning, where a candidate action may contain an arbitrary number L of sub-actions. The belief $b[X_k]$ can then be propagated to $b[X_{k+L}]$ for some specific action, given the corresponding new factors and variables. Yet, in general, belief propagation can be formulated simply as factor graph augmentation, independently of the number of look ahead steps L . That is, applying any candidate action of any length L can be viewed as the addition of new variable and factor nodes to graph G_k , thus making notation L unnecessary in the context of this article.

Therefore, in this article, we reduce clutter by denoting the belief, *before* the factor graph is changed, as $b[X_-]$. Similarly, we denote the posterior belief, *after* the factor graph is changed, as $b[X_+]$. In inference, $b[X_-]$ corresponds to the belief from the previous time instant, while $b[X_+]$ corresponds to the posterior belief at the current time, after incorporating new factors and variables. In the context of planning, $b[X_-]$ denotes the belief at planning time (i.e., current time) and $b[X_+]$ denotes the propagated belief along a given (non-myopic) candidate action. In both cases (inference and planning), X_- and X_+ represent the corresponding state vector before and after the change in the inference problem. For more details see also Table 1.

4. Problem formulation

We now introduce the two problems this article addresses, along with appropriate notation: general-purpose incremental covariance update and computationally efficient BSP. As shown in Section 5, our approach to address the latter problem builds upon the solution to the first problem.

Table 1. Main notation used within the problem definition (Section 4).

Notation	Description	
	Problem 1	Problem 2
X_-	state vector before a change in inference problem	state vector at planning time
X_+	state vector after a change in inference problem	future state vector after applying a specific candidate action
X_{new}	new state variables introduced after a change in inference problem	new state variables introduced after applying a specific candidate action
F_{new}	new factors set introduced after a change in inference problem	new factors set introduced after applying a specific candidate action
Λ_- and Λ_+	prior and posterior information matrices	
Λ_+^{Aug}	prior information matrix Λ_- augmented with zero rows and columns that represent new state variables X_{new} (see Figure 3)	
Σ_- and Σ_+	prior and posterior covariance matrices	
R_- and R_+	prior and posterior square-root information upper-triangular matrices	
$b[X]$	belief of state vector X	
$\mathcal{H}(b[X])$	differential entropy of belief $b[X]$	
$\Sigma_{M,Y}$	marginal covariance of some state subset Y (partition of covariance matrix Σ with columns\rows belonging to Y)	
$I(a)$	<i>increment</i> of candidate action a , represents new factors and new state variables introduced into inference problem after a is executed	
\mathcal{A}	set of candidate actions considered in BSP	
A	noise-weighted Jacobian of newly introduced factors with respect to all state variables	

4.1. Problem 1: Covariance recovery

As mentioned previously, in many applications it is mandatory to recover covariance entries of belief $b[X_-]$. However, typically this belief's covariance matrix Σ_- is represented by an information matrix $\Lambda_- = (\Sigma_-)^{-1}$, or the square-root information upper-triangular matrix R_- , with $R_-^T \cdot R_- = (\Sigma_-)^{-1}$.

Considering a square-root representation, the covariance matrix is $\Sigma_- = R_-^{-1} \cdot R_-^{-T}$ and its specific covariance entries $\Sigma_- = (\sigma_{ij})$ can be calculated from entries $R_- = (r_{ij})$ as (Golub and Plemmons, 1980)

$$\sigma_{ii} = \frac{1}{r_{ii}} \left(\frac{1}{r_{ii}} - \sum_{j=i+1, r_{ij} \neq 0}^n r_{ij} \sigma_{ji} \right) \quad (8)$$

$$\sigma_{ij} = \frac{1}{r_{ii}} \left(- \sum_{k=i+1, r_{ik} \neq 0}^j r_{ik} \sigma_{kj} - \sum_{k=j+1, r_{ik} \neq 0}^n r_{ik} \sigma_{kj} \right) \quad (9)$$

Note that in order to calculate the upper left covariance entry (σ_{11}), all other covariance entries are required. Therefore, the worst-case computation (and memory) complexity of this recursive approach is still quadratic in state dimension n . Yet, when applied to SLAM problems, the actual complexity is usually better because the matrix R_- is typically sparse in such a setting.

In contrast, an incremental covariance update approach can be applied in order to recover the required covariance entries more efficiently. At each time-step, solving the inference problem for the current belief $b[X_-]$ from (1)

provides a MAP estimate and the corresponding covariance or (square-root) information matrix. However, at the next step the inference problem changes. To see that, consider the belief at the next time-step $b[X_+]$, which was obtained by introducing new state variables X_{new} (e.g., new robot poses in SLAM smoothing formulation), with $X_+ = X_- \cup X_{new}$, and by adding new factors (e.g., new measurements, odometry, etc.) $F_{new} = \{f_{new}^1(X_+^1), \dots, f_{new}^{n_{new}}(X_+^{n_{new}})\}$ where $X_+^j \subseteq X_+$. In addition, consider the set of variables $Y \subset X_+$ whose marginal covariance $\Sigma_+^{M,Y}$ from $b[X_+]$ we are interested in calculating. Note that these variables of interest can contain both old $Y_{old} \subset X_-$ and new $Y_{new} \subset X_{new}$ variables, with $Y = \{Y_{old}, Y_{new}\}$.

Given that we already calculated the required covariance entries $\Sigma_-^{M,Y_{old}}$ from the current belief $b[X_-]$, in incremental covariance update approach we would like to update these entries after the change in the inference problem (from $b[X_-]$ to $b[X_+]$) as

$$\Sigma_+^{M,Y_{old}} = \Sigma_-^{M,Y_{old}} + \Delta^{Y_{old}} \quad (10)$$

where $\Delta^{Y_{old}}$ represents the difference between old and new covariance entries. In addition, in the general case we might be interested in calculating the posterior covariance of new variables of interest $Y_{new} \subset Y$, i.e., $\Sigma_+^{M,Y_{new}}$, as well as also the cross-covariances between Y_{old} and Y_{new} .

Likewise, the conditional covariances, from the conditional pdf of one state subset conditioned on another, are also required for information-theoretic BSP as was shown in Kopitkov and Indelman (2017). Hence, we would also

like to develop a similar approach for incremental conditional covariance update.

A limited technique to perform an incremental update of marginal covariances was presented in Ila et al. (2015). The authors showed how to update the covariance entries by downdating the posterior information matrix. Their derivation can be applied for the case where the state vector was not augmented during the change in the inference problem (X_{new} is empty). However, that derivation is not valid for the case of state augmentation, which involves zero-padding of prior matrices (described in the following); such padding yields singular matrices and requires more delicate handling. Even though their approach is not mathematically sound for the augmentation case, in the simulation part of Ila et al. (2015) it is suggested that the approach can also be applied here in practice. Still, the authors clearly declare that their approach does not handle relinearization of the state vector, which can often happen during the change in the inference problem. Further, Ila et al. (2015) did not consider recovery of conditional covariances. In contrast, we develop a general-purpose method that handles incremental (marginal and conditional) covariance updates in all of the above cases in a mathematically sound way.

In Section 5, we categorize the above general change in the inference problem into different sub-cases. Further, we present an approach that carefully handles each such sub-case and incrementally updates covariances that were already calculated before the change in the inference problem, and also computes covariance of newly introduced state variables. As will be shown, the computational complexity of such a method, when applied to a problem where only the marginal covariances need to be recovered (i.e., block diagonal of Σ_-), is linear in n in the worst case. Furthermore, we will show how our incremental covariance update approach can be also applied to incrementally update conditional covariance entries. Later, this capability will be an essential part in the derivation of our BSP method, *rAMDL-Tree*.

4.2. Problem 2: BSP

Typically in BSP and decision-making problems we have a set of candidate actions $\mathcal{A} = \{a_1, a_2, \dots\}$ from which we need to pick the best action according to a given objective function. As shown in our previous work (Kopitkov and Indelman, 2017), the posterior belief for each action can be viewed as a specific augmentation of the prior factor graph that represents the prior belief $b[X_-]$ (see Figure 2a). In this article, we denote this factor graph by G_- . Each candidate action a can add new information about the state variables in form of new factors. In addition, in specific applications, action a can also introduce new state variables into the factor graph (e.g., new robot poses). Thus, similarly to the change in the inference problem described previously, for each action a we can model the newly introduced state variables denoted by X_{new} , defining the posterior state




Fig. 3. Illustration of Λ_+ 's construction for a given candidate action in the case where new state variables X_{new} were introduced into the state vector. First, Λ_+^{Aug} is created by adding zero rows and columns representing the new state variables. Then, the new information of belief is computed through $\Lambda_+ = \Lambda_+^{Aug} + A^T A$.

vector (after applying the action) as $X_+ = X_- \cup X_{new}$. In a similar manner, we denote the newly introduced factors by $F_{new} = \{f_{new}^1(X_+^1), \dots, f_{new}^{n_{new}}(X_+^{n_{new}})\}$ where $X_+^j \subseteq X_+$.

Therefore, similar to (1), after applying candidate action a , the posterior belief $b[X_+]$ can be explicitly written as

$$b[X_+] \propto b[X_-] \prod_{j=1}^{n_{new}} f_{new}^j(X_+^j) \quad (11)$$

Such a formulation is general and supports a non-myopic action a with any planning horizon, that introduces into the factor graph multiple new state variables and multiple factors with any measurement model. Still, in this article we assume factors have a Gaussian form (2).

For the sake of conciseness, in this article the newly introduced factors and state variables that are added when considering action a will be called action a 's *increment* and denoted as

$$I(a) \doteq \{F_{new}, X_{new}\} \quad (12)$$

The posterior information matrix, i.e., the second moment of the belief $b[X_+]$, can be written as

$$\Lambda_+ = \Lambda_- + A^T \cdot A, \quad \Lambda_+ = \Lambda_+^{Aug} + A^T \cdot A \quad (13)$$

where we took the maximum likelihood assumption which considers that the above, a single optimization iteration (e.g., Gauss–Newton), sufficiently captures action impact on the belief. Such an assumption is typical in BSP literature (see, e.g., Indelman et al., 2015; Kim and Eustice, 2014; Platt et al., 2010; Van Den Berg et al., 2012). The left identity in (13) is true when X_{new} is empty, while the right identity is valid for non-empty X_{new} . The matrix A is a noise-weighted Jacobian of newly introduced factors F_{new} with respect to state variables X_+ ; Λ_+^{Aug} is constructed by first augmenting the prior information matrix Λ_- with zero rows and columns representing the new state variables X_{new} , as illustrated in Figure 3 (see, e.g., Kopitkov and Indelman, 2017).

After modeling the posterior information matrix Λ_+ for action a , the *unfocused* IG (uncertainty reduction of the entire state vector X_+) can be computed as

Table 2. Main notation used through the derivation of the incremental covariance recovery approach.

Notation	Description
Y	subset of state variables whose marginal covariance we are interested in updating\computing
Y_{old}	old variables inside Y
Y_{new}	new variables inside Y (that were introduced during the change in the inference problem)
${}^I X$	set of old <i>involved</i> variables in the newly introduced factors F_{new}
W	variable union of sets Y_{old} and ${}^I X$
n	dimension of a prior state vector X_-
N	dimension of a posterior state vector X_+
m	overall dimension of the newly introduced factors F_{new}
${}^I A$	$m \times {}^I X $ matrix that consists of A 's columns belonging to variables in ${}^I X$

$$J_{IG}(a) = \mathcal{H}(b[X_-]) - \mathcal{H}(b[X_+]) = \text{dim.const} + \frac{1}{2} \ln \frac{|\Lambda_+|}{|\Lambda_-|} \quad (14)$$

where $\mathcal{H}(\cdot)$ is the differential entropy function that measures the uncertainty of input belief, and dim.const is a constant that only depends on the dimension of X_+ and, thus, is ignored in this paper. Note that the above *unfocused* IG is typically used in applications where the set of new variables, X_{new} , is empty and so both X_- and X_+ have the same dimension. In cases where X_{new} is not empty (e.g., SLAM smoothing formulation), usually *focused* IG is used (see below). The optimal action a^* is then given by $a^* = \arg\max_{a \in \mathcal{A}} J_{IG}(a)$.

For a *focused* BSP problem we would like to reduce uncertainty of only a subset of state variables $X^F \subseteq X_+$. When X^F consists of old variables X_- , $X^F \subseteq X_-$, we can compute its IG. Such IG is a reduction of X^F 's entropy after applying action a , $\mathcal{H}(b[X_-^F]) - \mathcal{H}(b[X_+^F])$ where $b[X_-^F]$ and $b[X_+^F]$ are prior and posterior beliefs of *focused* variables X^F . In case X^F consists of newly introduced variables X_{new} , $X^F \subseteq X_{new}$, the IG function has no meaning as the prior belief $b[X_-^F]$ does not exist. Instead, we can directly calculate X^F 's posterior entropy $\mathcal{H}(b[X_+^F])$. The IG and entropy functions can be calculated through respectively:

$$J_{IG}^F(a) = \frac{1}{2} \ln \frac{|\Sigma_-^{M,F}|}{|\Sigma_+^{M,F}|},$$

$$J_{\mathcal{H}}^F(a) = \text{dim.const} + \frac{1}{2} \ln |\Sigma_+^{M,F}| \quad (15)$$

where $\Sigma_-^{M,F}$ and $\Sigma_+^{M,F}$ are prior and posterior marginal covariance matrices of X^F , respectively. Note that in *focused* BSP the optimal action will be found through $a^* = \arg\max_{a \in \mathcal{A}} J_{IG}^F(a)$ or $a^* = \arg\min_{a \in \mathcal{A}} J_{\mathcal{H}}^F(a)$.

To summarize, in order to solve an information-theoretic BSP problem, we are required to calculate IG or entropy ((14) and (15)) for each candidate action a , and then choose a candidate action with the maximal gain.

5. Approach

In this section, we present our approaches that efficiently solve the incremental covariance recovery (Section 5.1) and information-theoretic BSP (Section 5.2). The notation is given in Table 2.

5.1. Incremental covariance update

In this section, we present our technique for efficient update of covariance entries (see Problem 1 in Section 4). In Section 5.1.1, we show how to update marginal covariances of specified variables $Y \subset X_+$ after new state variables were introduced into the state vector and new factors were added, yet no state relinearization happened during the change in the inference problem. We show that the information matrix of the entire belief is propagated through a *quadratic* update form, similarly to (13). Assuming such *quadratic* update, we will derive a method to efficiently calculate the change in old covariance entries, to compute the new covariance entries and the cross-covariances between old and new state variables. Further, in Section 5.1.2 we show that also in the relinearization case the information matrix update has an identical *quadratic* update form and conclude that our method, derived in Section 5.1.1, can also be applied when some of the state variables were relinearized. Finally, in Section 5.1.3, we show that also the information matrix of a conditional pdf is updated through a *quadratic* update form and that the same technique from Section 5.1.1 can be applied in order to incrementally update conditional covariance entries. We show that our approach's complexity, given the specific prior covariances, does not depend on state dimension n .

5.1.1. Update of marginal covariance entries. Consider Problem 1 from Section 4. Consider the belief was propagated from $b[X_-]$ to $b[X_+]$ as described. Yet, let us assume for now that no state relinearization happened (we will specifically handle it in the next section). In this section, we show that the posterior covariances of interest $\Sigma_+^{M,Y}$ can be efficiently calculated as

Table 3. Summary of all different variations of change in inference problem. We use n to denote the prior state dimension, N to denote the posterior state dimension, and m to denote the dimension of all new factors within F_{new} .

Case	X_{new}	Information Update	Posterior State Dimension	A's Dimension
<i>Non-augmented</i>	empty	$\Lambda_+ = \Lambda_- + A^T \cdot A$	n	$m \times n$
<i>Rectangular</i>	not empty	$\Lambda_+ = \Lambda_+^{Aug} + A^T \cdot A$	$N = n + X_{new} $	$m \times N$
<i>Squared</i> (subcase of <i>Rectangular</i>)	not empty	$\Lambda_+ = \Lambda_+^{Aug} + A^T \cdot A$	$N = n + X_{new} $	$m \times N, m = X_{new} $

$$\Sigma_+^{M,Y} = f(\Sigma_-^{M,W}) \quad (16)$$

where $\Sigma_-^{M,W}$ is the prior marginal covariance of set $W \doteq \{Y_{old}, {}^I X\}$ and $f(\cdot)$ is a transformation function, with calculation complexity that does not depend on state dimension n . We derive this function in detail in the following. The set Y_{old} contains old state variables inside Y ($Y_{old} \subseteq X_-$) and ${}^I X \subseteq X_-$ is the set of *involved* variables in the newly introduced factors F_{new} : variables that appear in F_{new} 's models (3). Note that the update of old covariances (10) is only one part of this $f(\cdot)$, as also the computation of covariances for new state variables Y_{new} and cross-covariances between Y_{old} and Y_{new} .

Next, let us separate all possible changes in the inference problem into different cases according to augmented state variables X_{new} and the newly introduced factors F_{new} .

If X_{new} is empty, we call such a case as *non-augmented*. This case does not change the state vector ($X_- \equiv X_+$) and only introduces new information through new factors. The information matrix in this case can be updated through $\Lambda_+ = \Lambda_- + A^T \cdot A$, where matrix $A \in \mathbb{R}^{m \times n}$ is a noise-weighted Jacobian of newly introduced factors F_{new} with respect to state variables X_+ , and A 's height m is the dimension of all new factors within F_{new} (see Section 4).

Given X_{new} is not empty, we call such a case as *rectangular*. This case augments the state vector to be $X_+ = \{X_-, X_{new}\}$ and also introduces new information through the new factors. Here the information matrix can be updated through $\Lambda_+ = \Lambda_+^{Aug} + A^T \cdot A$ where $\Lambda_+^{Aug} \in \mathbb{R}^{N \times N}$ is a singular matrix that is constructed by first augmenting the prior information matrix Λ_- with zero rows and columns representing the new state variables X_{new} , as illustrated in Figure 3; $N = |X_+| = n + |X_{new}|$ is the posterior state dimension; A here will be an $m \times N$ matrix.

Finally, for the case when X_{new} is not empty and total dimension of new factors m is equal to the number of newly introduced variables $|X_{new}|$, we call such a case as *squared*. Clearly, the *squared* case is a specific case of the *rectangular* case, which for instance can represent the new robot poses of candidate trajectory and the new motion model factors. The reason for this specific case to be dealt with in a special way is due to the fact that its $f(\cdot)$ function is much simpler than the $f(\cdot)$ function of the more general *rectangular* case, as we show in the following. Thus, when

$m = |X_{new}|$ it would be advisable to use function $f(\cdot)$ of the *squared* case.

A summary of the above cases can be found in Table 3.

Next, we present the function $f(\cdot)$ separately for each one of the *non-augmented*, *rectangular*, and *squared* cases. Although the function $f(\cdot)$ has an intricate form (especially in the *rectangular* case), all matrix terms involved in it have dimensions m , $|X_{new}|$, or ${}^I X$; hence, the overall calculation of posterior $\Sigma_+^{M,Y}$ does not depend on the state dimension n .

Lemma 1. For the *non-augmented* case, the posterior marginal covariance $\Sigma_+^{M,Y}$ can be calculated as

$$\Sigma_+^{M,Y} = \Sigma_-^Y - B \cdot C^{-1} \cdot B^T, \quad B \triangleq \Sigma_-^C \cdot ({}^I A)^T, \quad (17)$$

$$C \triangleq I_m + {}^I A \cdot \Sigma_-^I \cdot ({}^I A)^T$$

where Σ_-^I , Σ_-^Y , and Σ_-^C are parts of prior marginal covariance $\Sigma_-^{M,W}$ partitioned through $W = \{Y, {}^I X\}$:

$$\Sigma_-^{M,W} = \begin{pmatrix} \Sigma_-^Y & \Sigma_-^C \\ (\Sigma_-^C)^T & \Sigma_-^I \end{pmatrix} \quad (18)$$

and where ${}^I A$ consists of A 's columns belonging to involved old variables ${}^I X$.

The proof of Lemma 1 is given in Appendix A. Note that sets Y and ${}^I X$ are not always disjoint. In case these sets have mutual variables, the cross-covariance matrix Σ_-^C can be seen just as $\Sigma_-^{(Y, {}^I X)}$, partition of prior covariance matrix Σ_- with rows belonging to Y and columns belonging to ${}^I X$. Also note that the columns inside information matrices do not have to be ordered in any particular way, and that the previous lemma is correct for any ordering whatsoever. Typically, a symmetric matrix $\Sigma_-^{M,W}$ is already known since it was already calculated at a previous time step as part of covariance update. In case some rows/columns from $\Sigma_-^{M,W}$ are unknown, we can calculate them by computing specific columns of Λ_- 's inverse (see Section 5.1.4).

Lemma 2. For the *rectangular* case the prior marginal covariance $\Sigma_-^{M,W}$ and the posterior marginal covariance $\Sigma_+^{M,Y}$ have the forms

$$\Sigma_-^{M,W} = \begin{pmatrix} \Sigma_-^{Y_{old}} & \Sigma_-^C \\ (\Sigma_-^C)^T & \Sigma_-^I \end{pmatrix} \quad (19)$$

$$\Sigma_+^{M,Y} = \begin{pmatrix} \Sigma_+^{M,Y_{old}} & \Sigma_+^{(Y_{old}, Y_{new})} \\ (\Sigma_+^{(Y_{old}, Y_{new})})^T & \Sigma_+^{M,Y_{new}} \end{pmatrix} \quad (20)$$

where we partition Y variables into two subsets $Y_{old} \doteq X_- \cap Y$ and $Y_{new} \doteq X_{new} \cap Y$, and where $W = \{Y_{old}, {}^I X\}$.

Using parts of $\Sigma_-^{M,W}$ we can calculate parts of $\Sigma_+^{M,Y}$ as:

$$\Sigma_+^{M,Y_{old}} = \Sigma_-^{Y_{old}} - B \cdot G^{-1} \cdot B^T \quad (21)$$

$$\Sigma_+^{M,Y_{new}} = P^{(Y_{new}, \cdot)} \quad (22)$$

$$C \triangleq I_m + {}^I A \cdot \Sigma_-^I \cdot ({}^I A)^T \quad (23)$$

$$P \triangleq [(A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1}]^{(:, Y_{new})} \quad (24)$$

$$F \triangleq (A_{new}^T \cdot A_{new})^{-1} \quad (25)$$

$$K \triangleq I_m - A_{new} \cdot F \cdot A_{new}^T \quad (26)$$

$$K_1 \triangleq K \cdot {}^I A \quad (27)$$

$$B \triangleq \Sigma_-^C \cdot K_1^T \quad (28)$$

$$G \triangleq I_m + K_1 \cdot \Sigma_-^I \cdot K_1^T \quad (29)$$

where A_{new} consists of A 's columns belonging to newly introduced variables X_{new} . In addition, we use matrix slicing operator (e.g., $P^{(Y_{new}, \cdot)}$) as is accustomed in Matlab syntax.

Further, there are two methods to calculate $\Sigma_+^{(Y_{old}, Y_{new})}$ from (20).

Method 1:

$$\Sigma_+^{(Y_{old}, Y_{new})} = \Sigma_-^C \cdot ({}^I A)^T \cdot [C^{-1} \cdot {}^I A \cdot \Sigma_-^I \cdot ({}^I A)^T - I_m] \cdot A_{new} \cdot P \quad (30)$$

Method 2:

$$\Sigma_+^{(Y_{old}, Y_{new})} = \Sigma_-^C \cdot [K_1^T \cdot G^{-1} \cdot K_1 \cdot \Sigma_-^I - I_k] \cdot ({}^I A)^T \cdot A_{new} \cdot F^{(:, Y_{new})} \quad (31)$$

Empirically we found that method 2 is the fastest option. The proof of Lemma 2 is given in Appendix B.

Lemma 3. For the squared case the prior marginal covariance $\Sigma_-^{M,W}$ and the posterior marginal covariance $\Sigma_+^{M,Y}$ have the forms

$$\Sigma_-^{M,W} = \begin{pmatrix} \Sigma_-^{Y_{old}} & \Sigma_-^C \\ (\Sigma_-^C)^T & \Sigma_-^I \end{pmatrix} \quad (32)$$

$$\Sigma_+^{M,Y} = \begin{pmatrix} \Sigma_+^{M,Y_{old}} & \Sigma_+^{(Y_{old}, Y_{new})} \\ (\Sigma_+^{(Y_{old}, Y_{new})})^T & \Sigma_+^{M,Y_{new}} \end{pmatrix} \quad (33)$$

where we partition Y variables into two subsets $Y_{old} \doteq X_- \cap Y$ and $Y_{new} \doteq X_{new} \cap Y$, and where $W = \{Y_{old}, {}^I X\}$.

Using parts of $\Sigma_-^{M,W}$ we can calculate parts of $\Sigma_+^{M,Y}$ as

$$\Sigma_+^{M,Y_{old}} = \Sigma_-^{Y_{old}} \quad (34)$$

$$\Sigma_+^{M,Y_{new}} = A_{iv} \cdot C \cdot A_{iv}^T \quad (35)$$

$$\Sigma_+^{(Y_{old}, Y_{new})} = -\Sigma_-^C \cdot ({}^I A)^T \cdot (A_{iv})^T \quad (36)$$

$$A_{iv} \triangleq [A_{new}^{-1}]^{(Y_{new}, \cdot)} \quad (37)$$

$$C \triangleq I_m + {}^I A \cdot \Sigma_-^I \cdot ({}^I A)^T \quad (38)$$

We can see that in case of a *squared* alteration, the covariances of old variables X_- do not change. The proof of Lemma 3 is given in Appendix C.

Note that in some applications the inner structure of Jacobian partitions ${}^I A$ and A_{new} can be known a priori. In these cases, such knowledge can be exploited and the runtime complexity of the above equations can be reduced even more.

5.1.2. Incremental covariance update after relinearization. Up to now we have explored scenarios where new information is introduced into our estimation system in a *quadratic* form via (13). Such information update is appropriate for planning problems where we take the linearization point of existing variables X_- (their current mean vector) and assume to know the linearization point of newly introduced variables X_{new} . However, during the inference process itself, state relinearization can happen and such a *quadratic* update form is no longer valid. This is because relinearization can be viewed as the removal of a factors subset and the addition of a new factors subset, yet the *quadratic* update form in (13) is correct only when new factors are solely added. Specifically, relinearization can be seen as a modification of a factor graph where factors, linearized with an old linearization point, are removed from the system and their relinearized versions are then introduced. As we show in (39)–(40), such inference update introduces new information that can be expressed via a *complex quadratic* form instead of regular *quadratic* form in (13). This *complex quadratic* form in its turn is required to apply the approach derived in Section 5.1.1 to incrementally update posterior covariances in relinearization case. In this section, we describe the entire process to incrementally update covariance entries after a relinearization event, which is more efficient than state-of-the-art approaches that calculate specific posterior covariances from posterior information matrix from scratch.

Relinearization may happen when a significantly new piece of information was added into the inference problem and current linearization point of state vector X_- does not optimally explain it anymore. In such cases, iterative optimization algorithms, such as Gauss–Newton, are responsible for updating the current linearization point, i.e., to find a more optimal linearization point that better explains the collected so far measurement/motion/prior factors.

Conventional approaches relinearize the entire state vector when new data comes in. On the other hand, incremental optimizer ISAM2 (Kaess et al., 2012) tracks instead the validity of a linearization point of each state variable and relinearizes only those variables whose change in the linearization point was above a predefined threshold. At each iteration of the nonlinear optimization and for each state variable x_i , ISAM2 finds δ_i and, given it is too big (norm of δ_i is bigger than the threshold), updates the current estimate of x_i to $x_i^* = x_i^* + \delta_i$. In such case, factors involving this state variable need to be relinearized. Clearly, the frequency of such a relinearization event during the inference process depends on the value of the threshold, and can be especially high during, for example, loop closures in a SLAM scenario. Still, in our simulations we have seen that even with a relatively high threshold and small number of loop closures, relinearization of some small state subset $R \subseteq X_-$ happens almost every second time-step. Thus, in order to accurately track covariances in the general case, while using conventional approaches that relinearize each time or ISAM2 which relinearizes only when it is needed, it is very important to know how to incrementally update covariance entries also after the state was relinearized. In the following, we show that information update of such a relinearization event can be also expressed in a *quadratic* form by using *complex* numbers; thus, the methods from Section 5.1.1 that incrementally update specific covariance terms can be used also here after applying *complex* numbers appropriately.

Denote by F_R the factors that involve any of the variables in R . In order to update information of the estimation after relinearization, we would want to remove F_R 's information with respect to old linearization point and to add F_R 's information with respect to the new one. It is not hard to see that the posterior information matrix (after relinearization of subset R) can be calculated through

$$\Lambda_+ = \Lambda_- - A_-^T \cdot A_- + A_+^T \cdot A_+ \quad (39)$$

where matrix A_- is a noise-weighted Jacobian of factors F_R with respect to old linearization point, and matrix A_+ is a noise-weighted Jacobian of factors F_R with respect to new linearization point.

Next, using complex numbers the above equation becomes

$$\begin{aligned} \Lambda_+ &= \Lambda_- + \begin{pmatrix} iA_-^T & A_+^T \end{pmatrix} \cdot \begin{pmatrix} iA_- \\ A_+ \end{pmatrix} = \Lambda_- \\ &+ \begin{pmatrix} iA_- \\ A_+ \end{pmatrix}^T \cdot \begin{pmatrix} iA_- \\ A_+ \end{pmatrix} = \Lambda_- + B^T \cdot B, \quad B \triangleq \begin{pmatrix} iA_- \\ A_+ \end{pmatrix} \end{aligned} \quad (40)$$

Note that the T operator is transpose and not conjugate transpose. In the above, we see that also here the information update is *quadratic* and the *update* matrix B contains terms of old and new Jacobians of factors F_R that were

affected by the relinearization event. Therefore, the incremental covariance update described in Section 5.1.1 is also applicable here, making the update of specific covariances much more efficient compared with computation of the covariances from scratch (e.g., through (8)–(9)).

More specifically, the update in (40) is an instance of the *non-augmented* case from Section 5.1.1. By exploiting the specifics of matrix B 's structure, Lemma 1 can be reduced to the following.

Lemma 4. *For the relinearization case (40), the posterior marginal covariance $\Sigma_+^{M,Y}$ can be calculated as*

$$\Sigma_+^{M,Y} = \Sigma_-^Y - U \cdot U^T, \quad U \triangleq \Sigma_-^C \cdot M, \quad (41)$$

$$M \triangleq (iM_1 \quad M_2),$$

$$M_2 \triangleq ({}^l A_+)^T / \text{chol}[I + {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_+)^T], \quad (42)$$

$$M_1 \triangleq [({}^l A_-)^T - M_2 \cdot G] / \text{chol} \quad (43)$$

$$[I - {}^l A_- \cdot \Sigma_-^l \cdot ({}^l A_-)^T + G^T \cdot G],$$

$$G \triangleq M_2^T \cdot \Sigma_-^l \cdot ({}^l A_-)^T \quad (44)$$

where Σ_-^l , Σ_-^Y , and Σ_-^C are parts of the prior marginal covariance $\Sigma_-^{M,W}$ partitioned through $W = \{Y, {}^l X\}$:

$$\Sigma_-^{M,W} = \begin{pmatrix} \Sigma_-^Y & \Sigma_-^C \\ (\Sigma_-^C)^T & \Sigma_-^l \end{pmatrix} \quad (45)$$

and where ${}^l A_-$ consists of A_- 's columns belonging to the involved variables ${}^l X$; ${}^l A_+$ contains columns of A_+ that belong to ${}^l X$; I is the identity matrix of an appropriate dimension; $\text{chol}(\cdot)$ represents Cholesky decomposition which returns an upper triangular matrix; and “/” is the backslash operator from Matlab syntax ($A/B = A \cdot B^{-1}$).

The proof of Lemma 4 is given in Appendix D. While it is mathematically equivalent to Lemma 1, empirically we found that such a formulation is faster and more numerically stable in the case of relinearization.

5.1.3. Incremental conditional covariance update. Previously, we have seen how to update specific prior marginal covariances given that state's information update has a *quadratic* form $\Lambda_+ = \Lambda_- + A^T \cdot A$ or $\Lambda_+ = \Lambda_-^{Aug} + A^T \cdot A$. Similarly, we can derive such a method that incrementally updates *specific* conditional covariances since, as we show in the following, the update of the conditional information matrix from the conditional pdf has a similar form.

To prove this statement, let us focus on the *non-augmented* case where X_{new} is empty. Define a set of variables Y , whose posterior conditional covariance $\Sigma_+^{Y|F}$, conditioned on an arbitrary disjoint variable set F (with $\{Y \cup F\} = \emptyset$), needs to be updated. Next, let U be the set of all state variables that are not in F , and note that $Y \subseteq U$. The prior information matrix $\Lambda_-^{U|F}$ of the prior conditional probability distribution $U|F$ is just a partition of the entire prior

information matrix Λ_- that belongs to columns/rows of variables in U . Similarly, the posterior $\Lambda_+^{U|V}$ is a partition of Λ_+ . It can be easily shown that

$$\Lambda_+^{U|V} = \Lambda_-^{U|V} + (A^U)^T \cdot A^U \quad (46)$$

where A^U is a partition of the noise-weighted Jacobian matrix A that belongs to columns of variables in U .

Equation (46) shows that the conditional probability distribution $U|F$ has a *quadratic* update, similar to the marginal probability distribution of the entire state vector X . In addition, note that the required posterior conditional matrix $\Sigma_+^{Y|F}$ is a partition of the posterior conditional covariance matrix $\Sigma_+^{U|V} = (\Lambda_+^{U|V})^{-1}$. For better intuition, it can be seen similar to the posterior marginal matrix $\Sigma_+^{M,Y}$ being a partition of the posterior marginal covariance matrix $\Sigma_+ = (\Lambda_+)^{-1}$ in the *non-augmented* case (see Lemma 1). Thus, there exists a function $f^C(\cdot)$ that calculates $\Sigma_+^{Y|F}$ from $\Sigma_-^{W|F}$, where $\Sigma_-^{W|F}$ is the prior conditional covariance matrix of set $W \doteq \{Y, {}^I X^U\}$, conditioned on the set F ; here, ${}^I X^U$ are the *involved* variables that are in U . Derivation of such a function $f^C(\cdot)$ is trivial, by following the steps to derive function $f(\cdot)$ in Section 5.1.1, and is left out of this paper in order to not obscure it with additional complex notation.

A similar exposition can be also shown in the augmented case (i.e., X_{new} is not empty), where information update of the conditional distribution also has the augmented *quadratic* form. To summarize, the derived function $f(\cdot)$ in Sections 5.1.1 and 5.1.2 can also be used to incrementally update the specific conditional covariances by replacing the prior marginal covariance terms in it with appropriate prior conditional covariances.

5.1.4. Application of incremental covariance update to SLAM. In order to apply our incremental update method in a SLAM setting, we model each change in the inference problem in the form of two separate changes as follows. We consider a specific scenario where at each time step, new robot pose x_k (k is index of time step) and new landmarks L_{new} are introduced into the state vector X . Further, new factors are introduced into the inference system; these factors include one odometry factor f_O between poses x_{k-1} and x_k , projection and range factors F_{new}^L between the new pose x_k and new landmarks L_{new} , and finally projection and range factors F_{old}^L between x_k and old landmarks. In addition, in general a subset of old factors (denoted by F_R) was relinearized as a result of a linearization point change of some old state variables during the inference stage. In case no linearization point change was performed, this set of factors F_R is empty. Note that although we assume above only range and visual measurements, our approach would work for other sensors as well, e.g., in a purely monocular case.

In the first modeled change, we introduce into the inference system all the new state variables (x_k and L_{new}) and

their constraining factors (f_O and F_{new}^L), denoted by $X_{new}^S = \{x_k, L_{new}\}$ and $F_{new}^S = \{f_O, F_{new}^L\}$, respectively. It can be shown for this change that the dimension of its newly introduced state variables X_{new}^S is equal to the dimension of newly introduced factors F_{new}^S . Thus, such change has a form of the *squared* case (see Table 3) and the updated covariance entries owing to this change can be calculated by applying Lemma 3. Also note that after this change all the new state variables are properly constrained, which is essential for the information matrix to remain invertible. Denote this information matrix, i.e., after applying the first change, by Λ_M :

$$\Lambda_M = \Lambda_k^{Aug} + A_S^T \cdot A_S \quad (47)$$

where Λ_k^{Aug} is the prior information matrix Λ_{k-1} augmented with zero rows/columns for new state variables X_{new}^S and A_S is noise-weighted Jacobian of factors F_{new}^S .

The remaining parts of the original change in the inference problem are represented by the second change. The posterior information matrix can be updated due to this second change as

$$\Lambda_k = \Lambda_M + A_O^T \cdot A_O - A_-^T \cdot A_- + A_+^T \cdot A_+ \quad (48)$$

where A_O is the noise-weighted Jacobian of factors F_{old}^S ; A_- and A_+ are noise-weighted Jacobians of factors F_R with respect to old and new linearization points, respectively. The above equation can be rewritten as

$$\Lambda_k = \Lambda_M + B^T \cdot B, \quad B \triangleq \begin{pmatrix} iA_- \\ A_+ \\ A_O \end{pmatrix} \quad (49)$$

and the corresponding covariance matrix can be calculated through Lemma 4, or through Lemma 1 in case there was no relinearization at the current time step, i.e., $B \equiv A_O$.

To summarize, any change in the inference problem of our SLAM scenario can be represented as a combination of two fundamental changes: *squared* (Eq. (47)) followed by (relinearized) *non-augmented* (Eq. (49)); the information matrix is updated as

$$\Lambda_{k-1} \Rightarrow \Lambda_M \Rightarrow \Lambda_k \quad (50)$$

where M can be seen as a logical time step of middle point.

Covariances after the first change can be updated very fast through Lemma 3, since as we saw in Section 5.1.1, the marginal covariances of old variables do not change and only marginal covariances of new variables need to be computed in this case. To do this, we require marginal covariances of *involved* variables ${}^I X$ from Σ_{k-1} . Notice that ${}^I X$ of the first change contains only x_{k-1} , whose marginal covariance is available since it was already calculated in the previous time step. Thus, Lemma 3 can be easily applied and the marginal covariances of all state variables at middle point M can be efficiently evaluated.

Algorithm 1: `rAMDLInformationEvaluation` evaluates information impact of candidates through *rAMDL* approach and picks the one with the biggest impact.

1. **Inputs:**
 2. $\{^I A_i\}$: non-zero columns of noise-weighted Jacobians of action candidates $\{a_i\}$
 3. $\{^I X_i\}$: variables that are *involved* in new factors F_{new} of each action a_i
 4. **Outputs:** a^* : optimal action
 5. **begin:**
 6. Calculate prior covariances of variables $X_{All} = \{\cup^I X_i\}$
 7. **for** a_i **do**
 8. Calculate information impact (IG or posterior entropy, *unfocused* or *focused*), using $^I A_i$ and the required prior covariances calculated in line 6
 9. **end**
 10. Select candidate a^* with maximal information impact
 11. **end**
-

To update all marginal covariances after the second change (through Lemma 1 or Lemma 4) we require marginal covariance of *involved* variables $^I X$ (in factors F_{old}^S and F_R of the second change) from covariance matrix $\Sigma_M = \Lambda_M^{-1}$. Moreover, we will require cross-covariances from Σ_M between variables $^I X$ and the rest of the variables, as can be seen from the equations of the lemmas. Thus, we require entire columns from Σ_M that belong to $^I X$. These columns can be easily calculated at time $k-1$ (from prior covariance matrix Σ_{k-1}) and propagated to middle point M by applying Lemma 3. The specific columns (belonging to some state subset Y) of matrix Σ_{k-1} can be efficiently calculated through two back-substitution operations:

$$V \triangleq R_{k-1}^T \setminus I_Y, \quad \Sigma_{k-1}^{(:,Y)} = R_{k-1} \setminus V \quad (51)$$

where I_Y are columns from the identity matrix I that belong to variables in Y ; R_{k-1} is the square-root information upper-triangular matrix at time step $k-1$; and “ \setminus ” is Matlab’s back-substitution operator with $x = A \setminus B$ being identical to solving linear equations $Ax = B$ for x .

The other alternative for this two-stage incremental covariance update is to use Lemma 2 for the *rectangular* inference change as follows. The posterior information matrix Λ_k can be calculated in one step as

$$\begin{aligned} \Lambda_k &= \Lambda_k^{Aug} + A_S^T \cdot A_S + A_O^T \cdot A_O - A_-^T \cdot A_- + A_+^T \cdot A_+ \\ &= \Lambda_k^{Aug} + B^T \cdot B, \quad B \triangleq \begin{pmatrix} A_S \\ iA_- \\ A_+ \\ A_O \end{pmatrix} \end{aligned} \quad (52)$$

Such a change has a form of the *rectangular* case (see Table 3); therefore, the updated covariance entries and the marginal covariances of new state variables can be calculated by applying Lemma 2. Note that the T operator within the lemma is transpose and not conjugate transpose. Similarly to the above two-stage method, the *rectangular* case will also require entire columns from Σ_{k-1} that belong

to old *involved* variables. This can be done here in the same way through (51).

We evaluate the above methods in our SLAM simulation in Section 6.1 and show their superiority over other state-of-the-art alternatives.

5.1.5. Incremental covariance update by Ila et al. (2015). In this section, we briefly describe an alternative approach to incrementally update covariance entries as proposed by Ila et al. (2015). Later, in Section 6.1 we will use it as one of the state-of-the-art baseline methods and compare its performance with our method presented in Section 5.1.4. After the estimation system at current time step was updated and new information matrix Λ_k (or square-root information matrix R_k) was inferred, we check whether state relinearization happened. If any variable was relinearized, the method from Ila et al. (2015) calculates all the required covariances from scratch using the recursive method from (8)–(9).

In case no relinearization was done, columns of *involved* variables $^I X$ from matrix $\Sigma_k = (\Lambda_k)^{-1}$, denoted as $\Sigma_k^{(:,^I X)}$, are calculated similarly to (51) via

$$V \triangleq R_k^T \setminus I_{^I X}, \quad \Sigma_k^{(:,^I X)} = R_k \setminus V \quad (53)$$

Next, the posterior marginal covariances of $^I X$ are retrieved from the calculated columns as $\hat{\Sigma} = \Sigma_k^{(:,^I X)}$. Further, covariances of new variables (which are a subset of $^I X$) can be just retrieved from appropriate partitions of the calculated above columns $\Sigma_k^{(:,^I X)}$. In order to update prior covariances $\Sigma_{k-1}^{(Y,Y)}$ of specific old variables Y , the following scheme is applied:

$$U = I - ^I A \cdot \hat{\Sigma} \cdot (^I A)^T, \quad (54)$$

$$\hat{B}_Y = \Sigma_k^{(Y,^I X)} \cdot (^I A)^T, \quad (55)$$

$$\Delta \Sigma_Y = \hat{B}_Y \cdot U^{-1} \cdot (\hat{B}_Y)^T, \quad (56)$$

$$\Sigma_k^{(Y,Y)} = \Sigma_{k-1}^{(Y,Y)} + \Delta \Sigma_Y \quad (57)$$

where ${}^l A$ is a noise-weighted Jacobian of all new factors with respect to ${}^l X$ and $\Sigma_k^{(Y, {}^l X)}$ is a partition of previously calculated columns $\Sigma_k^{(, {}^l X)}$. For more details please refer to the original paper (Ila et al., 2015). The above method has a similarity with our Lemma 1. Note, however, that unlike matrix C in Lemma 1, the matrix U in (54) can be singular since it is calculated as the difference between two positive-definite matrices. In our evaluation this was indeed the case: matrix U had very small eigenvalues which in turn caused U^{-1} to have extra-large entries. Yet, these entries did not affect the overall covariance update accuracy. Further, the above method is derived by applying Woodbury matrix identity to inverse the $\Lambda' = \Lambda_k - A^T \cdot A$, where A is a noise-weighted Jacobian of all new factors with respect to the entire posterior state vector X_+ . In the augmentation case scenario, such matrix Λ' will have zero-padding rows and columns and therefore cannot be inverted. This can be another reason why matrix U was almost singular in our simulations.

In addition, as can be seen from the equations, when no relinearization is done both our approach and that proposed by Ila et al. (2015) contain the same basic computations of the same dimension and thus are expected to have similar runtime complexity. In Section 6.1, we demonstrate it empirically.

5.2. Information-theoretic BSP

In this section, we develop a new approach that, based on the derived above incremental covariance update method, efficiently solves the information-theoretic BSP problem defined in Problem 2 from Section 4. Given a set of candidate actions, the proposed paradigm exploits common aspects among different actions for efficient BSP in high-dimensional state spaces. Each (non-myopic) action gives rise to a posterior belief that can be represented by an appropriate factor graph. In many applications different candidate actions will share some newly introduced factors and state variables (their factor graph *increments*). For example, two trajectory candidates that partially share their navigation path, will introduce the same factors for this mutual trajectory part (see Figure 2a). The posterior factor graphs of these candidate actions therefore have common parts, in terms of factor and variable nodes, and in addition all of these factor graphs *start* from the belief at the current time.

Our proposed paradigm saves computation time by identifying the common parts in these posterior factor graphs, and switching to a unified graphical model that we introduce, the FGP action tree, which represents gradual construction of posterior factor graphs from the current factor graph. For instance, in Figures 2b and 2c two different FGP action trees are depicted. Both lead to the same posterior beliefs of candidate actions, yet one of them can be evaluated more efficiently, as will be explained in Section 5.2.2. Given such a graphical model, we develop efficient

method to evaluate information impact of each candidate action in unified way. As we show, this method requires specific covariance entries for the intermediate beliefs that are represented by the tree's vertices, which we calculate by our incremental covariance recovery method (see Section 5.1) with computational complexity that does not depend on state dimension n (see Section 5.2.3). Further, we avoid posterior belief propagation and calculation of determinants of huge matrices for each candidate action by using the aforementioned incremental covariance update and the *rAMD*L method from Kopitkov and Indelman (2017). Moreover, we evaluate candidates' common parts only once instead of considering these parts separately for each of the candidates.

Determining the *best* topology of the FGP action tree, given the individual factor graphs for different candidate actions, is by itself a challenge that requires further research. In this article, we consider one specific realization of this concept, by examining the problem of motion planning under uncertainty and using the structure of the candidate trajectories for FGP action tree construction (see Section 5.2.2). In the results reported in Section 6 we consider scenario of autonomous exploration in unknown environment where such tree topology allows us to reduce computation time twice compared with baseline approaches.

5.2.1. *rAMD*L Approach. In our recently developed approach, *rAMD*L (Kopitkov and Indelman, 2017), the information-theoretic costs (14) and (15) are evaluated efficiently, without explicit inference over posterior beliefs for different actions and without calculating determinants of large matrices. As *rAMD*L is an essential part of our approach presented herein, in the following we provide a concise summary for the sake of completeness of the current article. For a more detailed review of *rAMD*L the reader is referred to Kopitkov and Indelman (2017).

In Kopitkov and Indelman (2017) we showed that the information impact of action a ((14) and (15)) is a function of prior covariances for the subset ${}^l X \subseteq X_-$ that contains variables *involved* in new factors F_{new} of a , and of matrix ${}^l A$ that contains non-zero columns of the noise-weighted Jacobian matrix A . Given the prior covariances of ${}^l X$, such a function can be calculated very fast, with complexity independent of state dimension. Thus, in *rAMD*L we first calculate the required prior covariances for all candidate actions as a one-time, yet still expensive, calculation, after which we efficiently evaluate the information impact of each candidate action. The main structure of the *rAMD*L approach is shown in Algorithm 1.

In particular, for the case where X_{new} is empty, the *unfocused* IG from (14) can be calculated as

$$J_{IG}(a) = \frac{1}{2} \ln \left| I_m + {}^l A \cdot \Sigma_-^{M, {}^l X} \cdot ({}^l A)^T \right| \quad (58)$$

Algorithm 2. EvaluateFGPTree evaluates the information impact of candidates and picks the one with the biggest impact.

1. **Inputs:** T : FGP tree
 2. **Outputs:** a^* : optimal action
 3. **begin:**
 4. **for** v : vertex of T **do**
 5. Determine the set of variables Y whose covariances are required by $rAMD L$ in order to calculate IG of action between v 's parent and v (variables *involved* in factors that were introduced by augmenting v 's parent to acquire v)
 6. **end**
 7. Calculate these covariances (see Section 5.1)
 8. Calculate IG of each edge of the tree through $rAMD L$ (e.g., through (58))
 9. Calculate IG of each candidate by summarizing IG along the candidate's trajectory
 10. Select candidate a^* with maximal IG
 11. **end**
-

where $\Sigma_-^{M,IX}$ is the prior marginal covariance of IX variables.

In case X_{new} is empty and we want to calculate *focused* IG of *focused* variables in $X^F \subseteq X_-$ (see (15), left), it can be calculated through

$$J_{IG}^F(a) = \frac{1}{2} \ln \left| I_m + {}^I A \cdot \Sigma_-^{M,IX} \cdot ({}^I A)^T \right| - \frac{1}{2} \ln \left| I_m + {}^I A^U \cdot \Sigma_-^{IX^U|X^F} \cdot ({}^I A^U)^T \right| \quad (59)$$

where $IX^U \equiv IX \setminus X^F$ denotes the *involved* variables that are *unfocused*, $\Sigma_-^{IX^U|X^F}$ is the prior conditional covariance of IX^U conditioned on X^F , and ${}^I A^U$ is a partition of ${}^I A$ with columns that belong to variables in IX^U .

In order to efficiently evaluate all candidates in the *unfocused* case, $rAMD L$ first calculates the prior marginal covariance $\Sigma_-^{M,X_{All}}$ of variables $X_{All} \subseteq X_-$, where X_{All} is the union of *involved* variables IX of all candidate actions. Further, evaluation of IG for each action is done by first retrieving $\Sigma_-^{M,IX}$ from $\Sigma_-^{M,X_{All}}$ and then calculating $J_{IG}(a)$ via (58). Overall, such a process consists of only one-time calculation that depends on state dimension n , i.e., calculation of $\Sigma_-^{M,X_{All}}$. Other cases of interest (where X_{new} is non-empty or for *focused* BSP objective functions) are also addressed by $rAMD L$. Note that in case of *focused* BSP the prior conditional covariances $\Sigma_-^{IX^U|X^F}$ are additionally required (see (59)) and can also be calculated for all candidates in one-block computation.

Yet, the $rAMD L$ method does not fully exploit similarities between candidate actions. The mutual *increment* part of the actions is expressed as identical block-rows in the matrix ${}^I A$ of these actions and, thus, is evaluated multiple times. In the next section, we present a novel approach to perform planning under uncertainty where mutual parts of the actions can be evaluated only once, further decreasing the CPU demand of the overall planning task.

5.2.2 FGP action tree. The FGP action tree describes the concept of belief propagation through a factor graph representation. Each vertex in this tree (see Figures 2b and 2c)

encodes a factor graph that represents a specific belief. For example, the root represents the prior belief $b[X_-]$ and leafs represent the posterior factor graphs of different candidate actions. Each edge $e_{v \rightarrow u}$, between vertices v and u , represents an action a with an appropriate *increment* $I(a)$, see (12). Thus, the factor graph encoded by vertex u is obtained by applying the increment $I(a)$ to the factor graph that is encoded by vertex v . In the following, we show how such a graphical model can be used to efficiently reason, while exploiting common parts, about posterior beliefs of different actions.

Let us consider a simple case as a running example, where two candidate actions a_1 and a_2 share some of their *increments* (see Figure 2a). As can be seen both trajectories have a mutual part that is colored in green. One way to evaluate the action impact for actions a_1 and a_2 is to handle each case separately (see Figure 2b). Indeed, existing approaches typically perform inference over the posterior belief for each of the actions and then evaluate the information-theoretic cost. However, this can be done by far more efficiently using our recently developed $rAMD L$ approach (Kopitkov and Indelman, 2017), where first we perform a one-time calculation of *specific* prior covariance entries that are required by both candidates, followed by information impact evaluation of each candidate (see Section 5.2.1). While the one-time covariance computation depends on state dimension n , the candidates evaluation does not. Yet, in such an approach, although we significantly reduce runtime by gathering the expensive computation of prior covariances from all candidates into a single computational block, we still waste computational resources related to the mutual *increment*, which is calculated separately for each candidate action (e.g., twice in the considered example).

In this article, we propose another alternative. Referring to the running example, we split each of the two actions into $a_1 = \{a_{shr}, a'_1\}$ and $a_2 = \{a_{shr}, a'_2\}$ and present them through a multiple-layered FGP tree (see Figure 2c), where a_{shr} represents the shared part of actions' *increments*, and where a'_1 and a'_2 represent parts of the original actions that are not shared. It is not difficult to show that IG of each

Algorithm 3. CalculateCovariances incrementally calculates specific covariances in the beliefs represented by vertices of the FGP action tree.

1. **Inputs:**
 2. T : FGP tree
 3. $\{Y_u\}$: set of variables whose covariances from vertex u 's belief we are interested in, for each vertex u in T
 4. **Outputs:** $\{\Sigma_u^{M, Y_u}\}$: the calculated covariances
 5. **begin:**
 6. **for** u : vertex of T , in bottom–top ordering**do**
 7. Message u 's parent v that we are interested in covariances in v 's belief for variables $\{Y_u, {}^I X, Y_{ch}\}$, where Y_u are variables required by the main algorithm (Algorithm 2), ${}^I X$ are variables *involved* in factors that were introduced by augmenting v to acquire u , and Y_{ch} are variables that were required by children of u
 8. **end**
 9. Define the set of variables $Y_T \doteq \{Y_u, Y_{ch}\}$ for each vertex u
 10. Calculate marginal covariances of Y_T at T 's root G_- , Σ_-^{M, Y_T} from prior belief $b[X_-]$ (e.g., through the Schur complement)
 11. **for** u : vertex of T , in top–bottom ordering**do**
 12. Calculate Σ_u^{M, Y_T} through function $f(\cdot)$ (see Section 5.1.1), by using required covariances $\Sigma_v^{M, W}$ from u 's parent v
 13. Retrieve the required Σ_u^{M, Y_u} from calculated Σ_u^{M, Y_T}
 14. **end**
 15. **end**
-

candidate a_i is equal to sum of IGs of its sub-actions a_{shr} and a'_i , i.e., $IG(a_i) = IG(a_{shr}) + IG(a'_i)$ (see the proof in Appendix E). Thus, in this specific example in order to select the best action it is enough to calculate IG of a'_1 and a'_2 . This IG can be efficiently calculated through the *rAMD*L technique, but this time we will require *specific* covariance entries of the intermediate belief associated with $G_+^{a_{shr}}$ (factor graph obtained after execution of action a_{shr} , see Figure 2c). For example, *unfocused* IG of a_1 (see (58)) can be calculated as

$$J_{IG}(a'_1) = \frac{1}{2} \ln \left| I_m + {}^I A_1 \cdot \Sigma_{shr}^{M, {}^I X_1} \cdot ({}^I A_1)^T \right| \quad (60)$$

where ${}^I X_1$ are variables *involved* in new factors of a'_1 , ${}^I A_1$ are non-zero columns from noise-weighted Jacobian of these new factors, and $\Sigma_{shr}^{M, {}^I X_1}$ is the marginal covariance of ${}^I X_1$ from intermediate belief represented by $G_+^{a_{shr}}$. Assuming there is an efficient way to calculate *specific* covariance entries for each vertex within the FGP tree (see Section 5.2.3), we can apply the *rAMD*L method, calculate the required information impacts and make a decision between a_1 and a_2 while handling mutual *increment* a_{shr} only once, and not twice as would be done by existing approaches.

The above concept applies also to more general problem settings, with numerous candidate actions with mutual parts in their *increments*. An excellent example for this is BSP for autonomous navigation. Here, the set of trajectory candidates can be naturally represented as a tree of possible paths, and the FGP action tree can be constructed in such a way that each of its intermediate vertices will represent a belief at a specific splitting waypoint of the trajectories (see Figure 4). In such a general case, in order to pick up the optimal action we will need to calculate IG for each one of the tree's edges. This can be done again by applying a *rAMD*L technique but will require us to know the *specific*

covariance entries for each intermediate vertex within the FGP tree. An efficient calculation of these entries is presented in Section 5.2.3, while the overall algorithm to evaluate the FGP tree is summarized in Algorithm 2.

Note that although in this article we create an FGP action tree with a structure similar to the tree of candidate navigation paths, in general, different structures can be used. For example, if candidates share their trajectories' terminal part, this part can be represented as first action under root G_- . As long as tree's root represents the prior belief $b[X_-]$ and the tree has a vertex for posterior belief of each candidate action, it represents the same decision problem. An interesting question that arises is how to find the tree's structure that provides the biggest calculation reuse between the candidates and can be evaluated most efficiently. We will leave this question for future research.

In addition, note that the proposed method can also be applied to the scenario where a similar candidate trajectory is evaluated at sequential time steps. Such a candidate trajectory, taking the robot to some location, at each time step may have a different starting section due to robot's movement since the previous time step, but will have the same terminal section that brings the robot to the aforementioned location (see also Chaves and Eustice (2016)). Thus, this candidate trajectory will have similar posterior factor graphs each time it is evaluated. This similarity between posterior factor graphs can be naturally represented through our FGP tree and hereof it is just another application for our BSP approach.

5.2.3. Incremental covariance update within FGP action tree. In order to reason about different actions inside an FGP action tree, we have to know *specific* covariance entries for each intermediate vertex in the tree. We can calculate these entries by first propagating the beliefs through (13) followed by appropriate Schur complement and inverse




Fig. 4. Different candidate trajectories and their FGP action tree representation. Here A and B are splitting waypoints; C , D , and E are the final waypoints of three trajectory candidates. Each waypoint has an associated factor graph vertex within the tree. Each action a_x^y is an augmenting factor graph with factors/variables gathered by passing path $x \rightarrow y$.

operations. However, such a procedure will depend on a potentially huge state dimension n , which we would like to avoid. Here we propose an alternative method to calculate specific covariance entries at each one of the beliefs inside the tree which is based on our incremental covariance update technique (see Section 5.1) and does *not* depend on n . Moreover, the proposed method can be applied to calculate both specific marginal and conditional covariance entries, where the former are required for the *unfocused* information objective function (see (58)) and the latter are required for the *focused* information objective function (see (59)).

First, let us focus on a specific edge $e_{v \rightarrow u}$ in the FGP tree that represents some action a with *increment* $I(a) = \{F_{new}, X_{new}\}$. In other words, the factor graph represented by v is augmented by $I(a)$ in order to receive the factor graph that is represented by u . In addition, let us denote state vectors of beliefs of v and of u by X_u and X_v , respectively. Note that $X_v \subseteq X_u$, and that X_u will sometimes contain state variables that are not present yet in X_v (the variable set X_{new}).

Now, consider the set of variables $Y \subseteq X_u$ whose marginal covariance $\Sigma_u^{M,Y}$ from u 's belief we would like to calculate. As was shown in Section 5.1, $\Sigma_u^{M,Y}$ can be calculated efficiently and independently of state dimension n , given that we have the marginal covariance of the set $W \doteq \{Y_{old}, I^X\}$ from v 's belief, where $I^X \subseteq X_v$ is the set of *involved* variables in action a and Y_{old} is the intersection between Y and X_v . It is important to note that, in a general case, the marginal covariance of Y_{old} is modified after applying some action a , i.e., $\Sigma_u^{M,Y_{old}} \neq \Sigma_v^{M,Y_{old}}$. Similarly to Section 5.1 we can separate all possible actions in the FGP tree into different categories depending on their *increments*, i.e., *non-augmented*, *rectangular*, and *squared*. Consequently, for each action type we can use an appropriate covariance update method in order to calculate $\Sigma_u^{M,Y} = f(\Sigma_v^{M,W})$.

Next, we can use the above-mentioned function $\Sigma_u^{M,Y} = f(\Sigma_v^{M,W})$ to calculate the required *specific*




Fig. 5. Incremental covariance update within an FGP action tree, illustrated on the tree from Figure 4. Each vertex in the tree represents a specific factor graph and belief associated with it. Each edge in tree (black arrows) represents action that augments parent's factor graph in order to obtain the child's factor graph. The covariance update consists of two stages. First (left drawing): from bottom to top each vertex notifies (green arrows) its parent what are the marginal covariance entries that it needs from parent's belief. Second (right drawing): from top to bottom each vertex calculates the required marginal covariance entries of its belief and notifies its children (red arrows) to proceed with their covariance calculations. This covariance update process contains one-time calculation depending on state dimension n , the computation of required marginal covariance entries at root G_- . The rest of the calculation is incremental and does not depend on n , as described in Section 5.1.

covariances for each one of the vertices in the tree recursively (see also Figure 5). First, for each vertex u we define by Y the variables of interest whose marginal covariances $\Sigma_u^{M,Y}$ at the belief associated with u we would like to calculate. In our case Y are the variables required by *rAMD*L in order to evaluate the impact of actions that are performed on u (see Section 5.2.2). Next, for each leaf vertex u we message its parent v that we require v 's marginal covariances for $\{Y, I^X\}$. Then in recursive form from bottom to top each vertex v will message its parent that it requires its parent's covariances for $\{Y, I^X, Y_{ch}\}$ where Y_{ch} is the set of variables that were required by v 's children. Eventually, for each vertex v in the tree we will have a total set of variables $Y_T \doteq \{Y, Y_{ch}\}$ whose covariances we need to compute for this specific vertex.

Finally, we start to propagate these covariances in top to bottom order. Using the equations from Section 5.1.1, for each vertex u we can calculate Σ_u^{M,Y_T} using $\Sigma_v^{M,W}$ from its parent vertex v . Note that when following top to bottom order, when we get to node u , its parent's covariances $\Sigma_v^{M,W}$ will already be computed. Also note that the required prior covariance entries of the root G_- should be calculated first. This is done only once and its complexity depends on state dimension n , similarly to the *rAMD*L technique. But once calculated, the rest of the covariance updates do not depend on n .

To summarize, the described algorithm consists of two parts: detecting variables set Y_T for each vertex and propagating *specific* covariances from top to bottom. See Algorithm 3 and a schematic illustration of the incremental

covariance update in Figure 5. The runtime complexity of the algorithm mainly depends on its second part, because variable detection does not require any matrix manipulations and can be done quickly. The second part handles each edge of the FGP tree only once, thus again allowing us to evaluate mutual *increment* of actions only once. The runtime to propagate *specific* covariances along each edge depends on a number of parameters such as dimension of required covariances and the size of an action’s *increment*.

In a similar way, we can also incrementally propagate *specific* conditional covariances along the FGP tree. Such covariances may also be required in order to perform informative-theoretic decision making when we want to reduce uncertainty of a subset of old variables $X^F \subseteq X_-$, see (Kopitkov and Indelman, 2017).

6. Results

We evaluate the proposed approaches for incremental covariance update and BSP in simulation considering the problem of autonomous navigation in unknown environments. The robot has to autonomously visit a set of predefined goals while localizing itself and mapping the environment using its onboard sensors. In our simulation, we currently consider a monocular camera and a range sensor. The code is implemented in Matlab and uses the GTSAM library (Dellaert, 2012; Kaess et al., 2012). All scenarios were executed on a Linux machine with an Intel i7 2.40 GHz processor and 32 GB of memory. All compared approaches were implemented in single thread to provide better visualization of their runtime complexity. In addition, we provide our implementation of the FGP action tree as an open-source library at <http://goo.gl/dmNenc>.

6.1. Covariance recovery

Here we consider the passive setting where at each time step the robot moves toward the next predefined goal (see Figure 6), updates the inference problem with new pose/landmarks and motion/measurement factors, and calculates/updates marginal covariance of each variable inside the state vector.

We apply our incremental covariance update methods (*2-stage* and *Rectangular*) as described in Section 5.1.4. Their performance is compared with two baseline approaches. First, *Recursive* uses a recursive formulation (see (8)–(9)) to calculate the covariance matrix Σ_k (k is index of time step) from a square-root information matrix R_k . It is done for each k and entire Σ_k is calculated at each time step from scratch. Note that to calculate the marginal covariance of each state variable (block-diagonal of Σ_k) the *Recursive* method requires the entire covariance matrix Σ_k to be calculated as was explained in Problem 1 from Section 4.

The second approach, *Backsubstitution*, calculates Σ_k through the back-substitution operation:

$$V \stackrel{\Delta}{=} R_k \setminus I, \quad \Sigma_k = V \cdot V^T \quad (61)$$

where I is an identity matrix of appropriate dimensions and “ \setminus ” is the Matlab’s back-substitution operator with $x = A \setminus B$ being identical to solving linear equations $Ax = B$ for x . Such back-substitution can be done very efficiently since the matrix R_k is upper triangular and sparse. Still, similar to *Recursive*, the *Backsubstitution* method calculates covariances from scratch for each time step and needs to calculate the entire Σ_k matrix before fetching its diagonal blocks.

As can be seen in Figure 7, in general both of our incremental approaches have very similar runtime, and the both are significantly faster than the baseline alternatives. Towards the end of the scenario, while the fastest alternative (*Backsubstitution*) needs almost 400 ms to recover marginal covariance for a 3,054-dimensional state vector, our incremental method does it in only 20 ms.

The only time our methods are slower than the alternatives is around pose 150, at which point a loop-closure event occurs: the robot reaches a predefined goal 6 (see Figure 6a) and observes old landmarks from the beginning of the scenario. As expected for such a relatively big loop closure, the number of relinearized state variables and the affected factors is very large (see Figures 6b–6c). Thus, m (overall dimension of new/relinearized factors) and $|^I X|$ (dimension of *involved* variables) are huge and increase the runtime complexity of our incremental method. However, such results are expected: it is a known fact that incremental techniques become slower in the presence of big loop closures. For example, the incremental optimization algorithm iSAM2 (Kaess et al., 2012), which calculates incrementally the MAP estimate of the state but not its covariance matrix, takes significantly more time during loop-closure events. It is reasonable to expect a similar situation also in the context of incremental covariance recovery. Also note that during a loop-closure event, the *Rectangular* technique is significantly slower than the *2-stage* technique (around 6 s versus 0.6 s, respectively). The reason is that during huge loop closure, m impacts the entire calculation of the *Rectangular* method (Lemma 2), while in the *2-stage* technique only the second stage is affected (Lemma 1 or Lemma 4). Lemma 2 is more computationally demanding than Lemma 1 or Lemma 4, thus producing such a big runtime difference during a loop closure.

On the other hand, the *Backsubstitution* method does not depend on m or $|^I X|$; instead its complexity mainly depends on the state dimension n and the sparsity level of a matrix R_k . The n , the overall dimension of all state variables, is not affected by loop closures. While, in general, the matrix R_k (a factorization of information matrix Λ_k) becomes denser during the loop closures, an appropriate

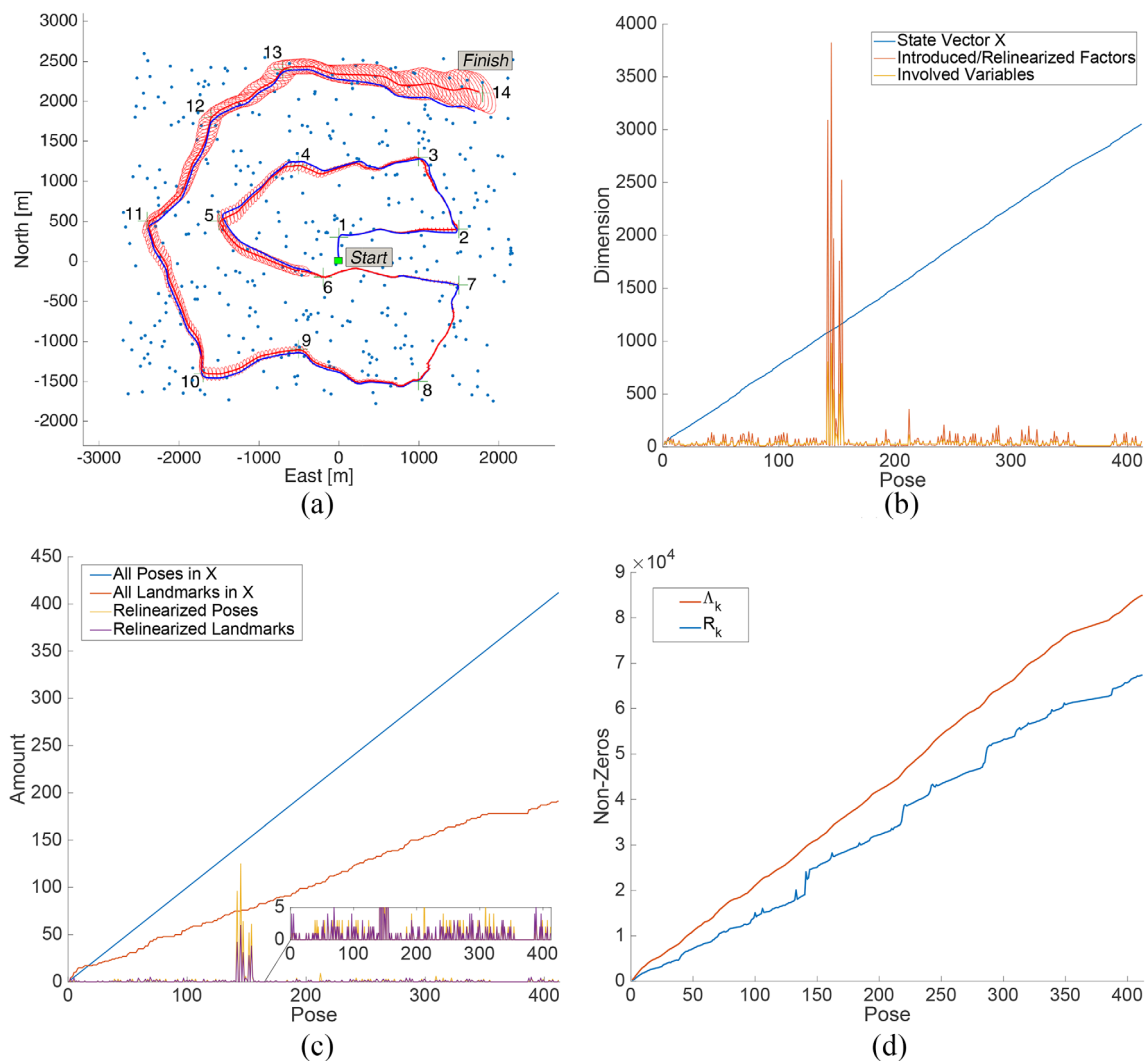


Fig. 6. Robot follows a predefined path by navigating through given way-points. (a) Robot trajectory. Blue dots are mapped landmarks, red line with small ellipses is estimated trajectory with pose covariances, blue line is the real trajectory, pluses with numbers beside them are robot's predefined goals. Green mark is robot's start position. (b) Dimensions at each time-step of state vector X , of overall introduced and relinearized factors and of *involved* variables in these factors. (c) Number of overall poses/landmarks inside state vector X and number of relinearized poses/landmarks. (d) Number of non-zero entries inside the information matrix Λ_k and the square-root information matrix R_k .

variable reordering (of the entire matrix) can mitigate this effect. In our simulations we used SYMAMD ordering (symmetric approximate minimum degree permutation) (Amestoy et al., 1996) to reorder an entire Λ_k before producing R_k . As can be seen in Figure 6d (blue line), the resulting sparsity of R_k grows smoothly with time, with only a minor increase during the loop-closure event (around pose 150). Thus, we can see no peaks in calculation time plot of *Backsubstitution* approach around this time (see Figures 7c–7d, purple line). In practice, when implementing our incremental approach on a real robot, to handle this loop-closure shortcoming we can check whether a big loop closure happens (m or $|^l X|$ are bigger than current state dimension n) and use *Backsubstitution* as a fallback.

Comparing *Recursive* versus *Backsubstitution* we can see that the former is considerably slower. The first was implemented by us in C++ code, while the second is based on highly optimized Matlab implementation of backsubstitution. Apparently, our current C++ implementation of *Recursive* method is not properly optimized. We foresee that it can be done in much better way so that both *Recursive* versus *Backsubstitution* techniques will have very similar runtime complexity.

6.1.1. Comparison with Ila et al. (2015). Here we compare our approach with that from Ila et al. (2015); that approach was briefly described in Section 5.1.5. As mentioned

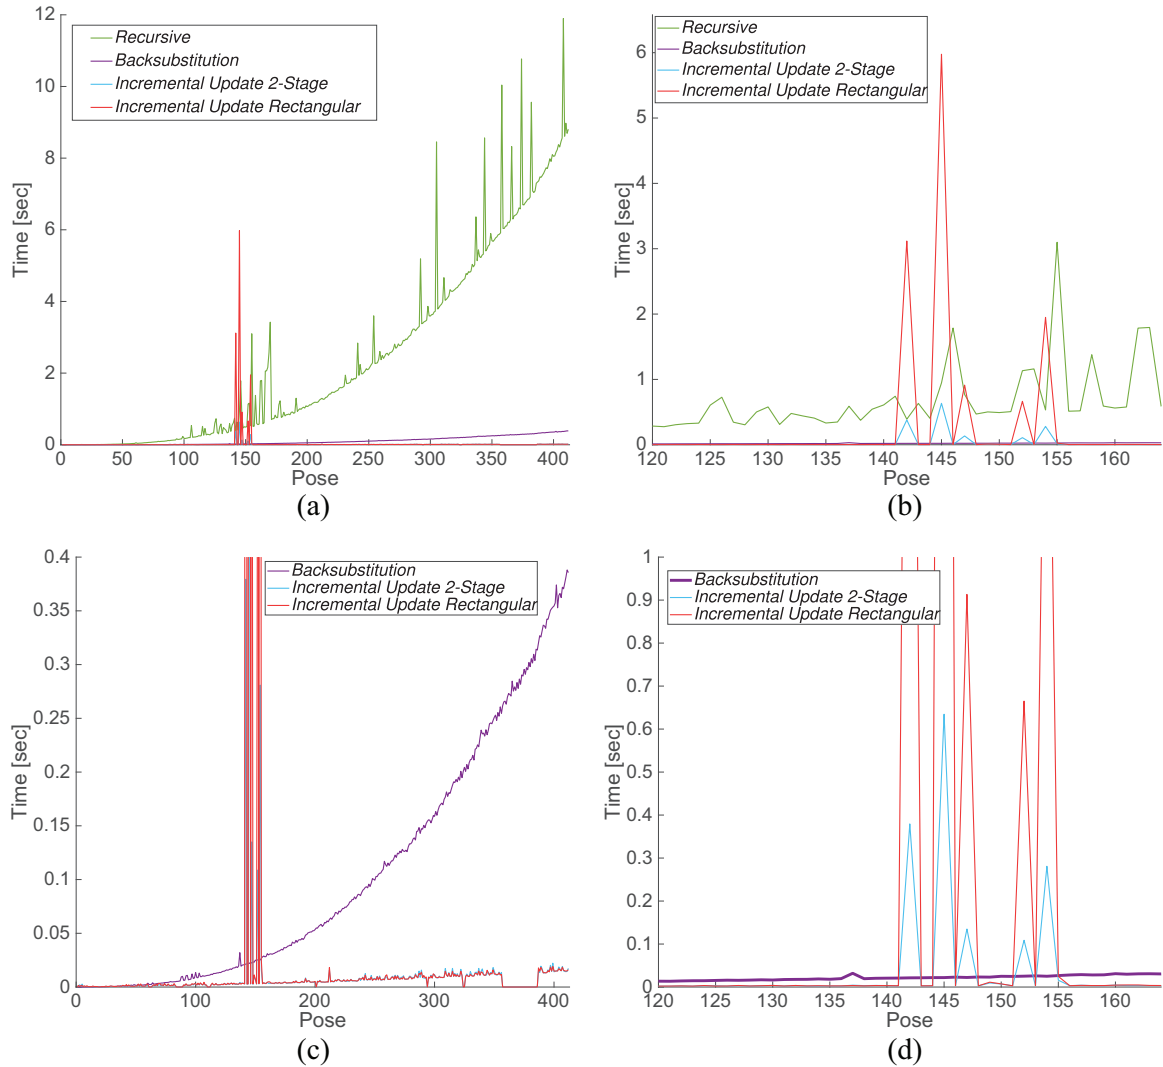


Fig. 7. Robot follows a predefined path by navigating through given way-points. (a) Running time of marginal covariance recovery, i.e., evaluating the marginal covariance matrix for each pose and each landmark. (b) Running time from (a) with zoom-in on loop closure around pose 150. (c) Running time from (a) with zoom-in on three fastest approaches. (d) Running time from (c) with zoom-in on loop closure around pose 150.

above, their method is limited and cannot be applied after relinearization of state vector, thus we execute this method only at time steps where no variable was relinearized. In Figure 6c we can see that there is only a sparse number of such time steps. Nonetheless, from Figure 8 we can see that it is enough to understand the overall trend. Specifically, as we expected, both our and Ila15icra’s methods have very similar runtime complexity, since the runtime of both techniques depends on m and $|^lX|$. Furthermore, note that in some edge cases our approach is much more efficient (see time-steps 357–386 in Figure 8). When during the time-step only new poses and only new odometry factors are added to the inference system, that is when $m = |^lX|$, our approach updates covariances incrementally via Lemma 3. This lemma is computationally very efficient since it does not update covariances of old variables which are unchanged in such a case. In contrast, the method in Ila




Fig. 8. Running time of marginal covariance recovery during robot navigation, i.e., evaluating the marginal covariance matrix for each pose and each landmark. The figure compares the proposed approach and that from Ila et al. (2015). The latter is only shown for time-steps when no state relinearization occurred.

et al. (2015) does not exploit the properties of this specific setting $m = |^lX|$.

In addition, the incremental recovery of *conditional* covariance entries is essential to efficiently solve the information-theoretic BSP problem which is considered in Section 6.2; yet, this key aspect was also not supported in Ila et al. (2015).

6.2. BSP

Thus far, we have performed simulation of a passive SLAM problem, where the robot follows a predefined trajectory. As can be seen in Figure 6a, by the end of the trajectory the covariance of robot position (red ellipse) is considerably big. Such uncertainty in robot localization may fail the navigation task and is undesirable in general. In this section, we focus on an active SLAM scenario, where the robot autonomously decides whether to follow the navigation path or to perform a loop closure and reduce state uncertainty. At each time step, the robot autonomously decides its next action according to a specified objective function that is discussed in the following.

We compare the performance of the proposed BSP approach, that we denote as *rAMDL-Tree*, with our previous method *rAMDL* (Kopitkov and Indelman, 2017), which was shown to be superior in runtime complexity to other state-of-the-art information-based BSP methods. Note that both *rAMDL* and *rAMDL-Tree*, as well as other relevant state-of-the-art alternatives, make *identical* decisions, i.e., calculate the same optimal actions. Thus, the only difference is the runtime complexity, the reduction of which is the main motivation behind the work presented herein.

In our simulation, at each time step we sample a set of trajectories to the current goal g , and also to (clusters of) already mapped landmarks for uncertainty reduction via loop closures (see Figure 9a). The overall number of candidate actions (number of trajectories) is around 200 (see Figure 9b). In our simulations different trajectories were sampled by using a probabilistic roadmap (PRM) algorithm (Kavraki et al., 1996), see Figures 9c–9d, although the planning method presented herein does not depend on the sampling procedure and other sampling techniques can be applied.

We consider the objective function

$$J(a) = \alpha_1 d(x_{k+L}, g) + \alpha_2 c(a) + \alpha_3 J_{inf}(a) \quad (62)$$

where $d(x_{k+L}, g)$ is the distance between the current goal g and candidate's last pose x_{k+L} for a given action a , $c(a)$ is the control cost, and $J_{inf}(a)$ is an information-theoretic term. As was mentioned previously, both first and second terms can be calculated very quickly and do not require belief propagation. Thus, in the following, we ignore these terms and discuss runtime only for the term $J_{inf}(a)$.

In our first simulation, $J_{inf}(a)$ calculates the posterior entropy of the robot's last pose x_{k+L} within the candidate trajectory. We evaluated this term for all action candidates independently through our previous method, *rAMDL*, and through the approach proposed herein, *rAMDL-Tree*,

which, using the FGP action tree, accounts for candidate actions' mutual parts and evaluates them only once. In Figures 10a–10b we can see that *rAMDL-Tree* is twice as fast as *rAMDL* and succeeds to evaluate more than 200 actions in less than 100 ms. In addition, we can see that the only algorithmic part that depends on the state dimension n , i.e., the one-time calculation of prior covariances at G_- during the incremental covariance update termed in the figure as *GTSAM-PriorMarginals*, takes a small portion of the overall runtime (green lines in Figures 10a–10b); most of the time is consumed by propagation of covariance entries within the FGP action tree and IG calculation for each edge in this tree. Note that the marginal/conditional covariances required by *rAMDL-Tree* are propagated from the root of FGP action tree to its leaves based on our incremental covariance update technique (see Section 5.1). Specifically, as described in Section 5.2.3, covariance propagation is performed in two phases. In the first, each tree node in bottom-to-top order determines what covariance entries are required from its belief. In the second phase, the required covariances are calculated in top-to-bottom order using incremental covariance update lemmas (see also Figure 5).

In addition, we have performed a similar simulation considering this time $J_{inf}(a)$ calculating the IG of the landmarks that were mapped until now. The results are shown in Figure 11. Comparing the time performance between the first and second scenarios we can see that “focused-landmarks” requires more time; while *rAMDL-Tree* takes around 50–100 ms in the first, it requires 100–250 ms in the second. This is due to the fact that calculation of the *focused* IG contains a one-time computation that depends on the dimension of *focused* variables set X_+^F . In the “focused-landmarks” scenario this is the dimension of all landmarks mapped thus far, and it increases with time as more landmarks are observed and introduced into the state vector. Also in this scenario we see a similar trend where *rAMDL-Tree* performs twice as fast as *rAMDL* (100–250 ms versus 200–500 ms), while determining the same optimal actions.

7. Conclusions

We have developed computationally efficient approaches that address incremental covariance recovery and BSP over high-dimensional state spaces. Our incremental covariance update technique allows us to efficiently update specific covariance entries (both marginal and conditional) after any change in the inference problem, including introduction of new state variables, addition of new measurement factors, and relinearization of the entire state vector or only a subset of the state variables. It can be applied whenever an efficient method is required to track covariance entries within the estimation system (e.g., in SLAM for data association or safety), and is also an indispensable part of our BSP approach.

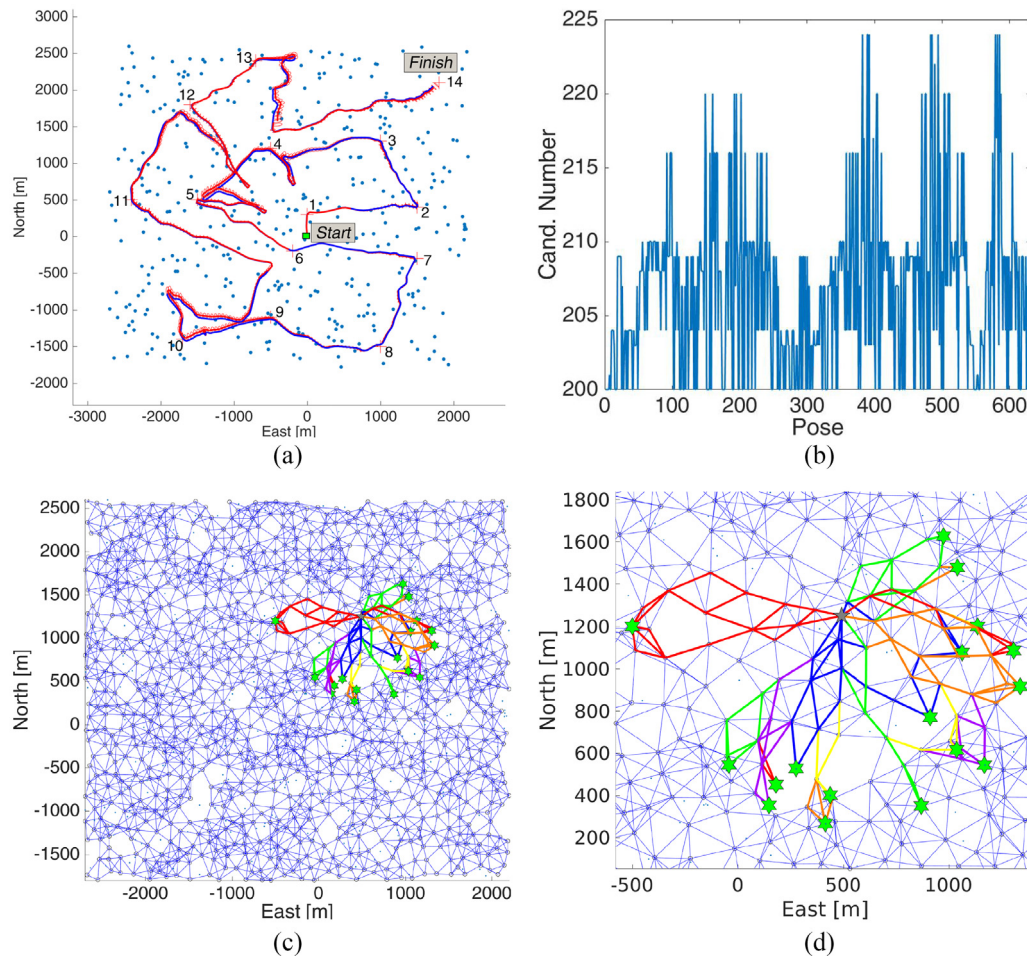


Fig. 9. *Focused* BSP scenario with *focused* robot's last pose. (a) Final robot trajectory. Blue dots are mapped landmarks, red line with small ellipses is estimated trajectory with pose covariances, blue line is the real trajectory, red pluses with numbers beside them are robot's predefined goals. Green mark is robot's start position. (b) Number of action candidates at each time. (c) Candidate trajectories at time step 69. (d) Enlarged view of (c). The pentagram represents current robot pose and hexagrams are current goal g and various landmark clusters. As can be seen, there are a number of candidate trajectories to each hexagram and many candidate trajectories share path sections. Our approach exploits the shared trajectory parts for faster planning.

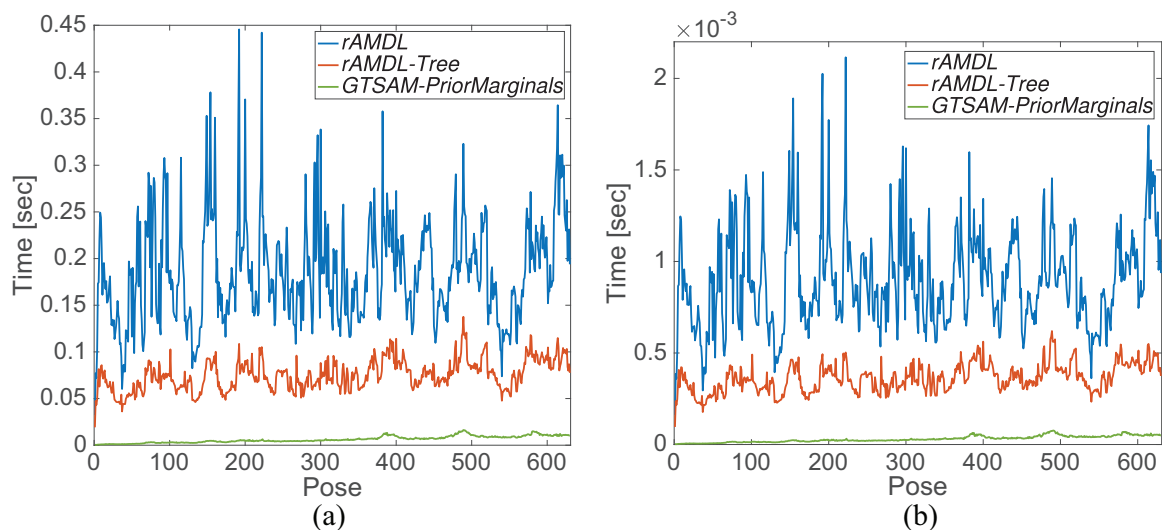


Fig. 10. *Focused* BSP scenario with *focused* robot's last pose. (a) Running time of planning, i.e., evaluating the impact of all candidate actions, each representing a possible trajectory. (b) Running time from (a) normalized by the number of candidates.

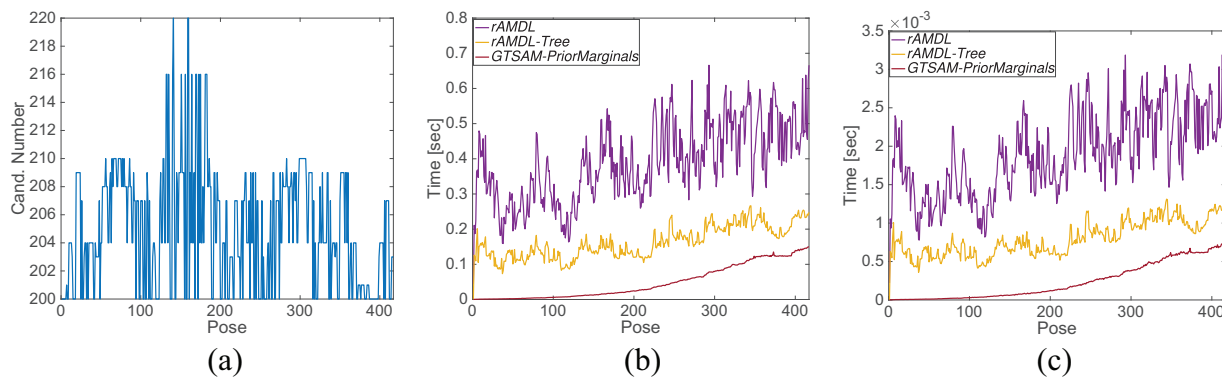


Fig. 11. *Focused* BSP scenario with *focused* mapped until now landmarks. (a) Number of action candidates at each time. (b) Running time of planning, i.e., evaluating the impact of all candidate actions, each representing possible trajectory. (c) Running time from (b) normalized by the number of candidates.

Furthermore, considering the BSP problem our key observation is that in many robotics applications, candidate actions have mutual parts where each part can be evaluated only once, independently of the number of candidate actions that share it. For this purpose, we presented a novel approach to model future posterior beliefs of different candidate actions within a single graphical model which we called an FGP action tree. This tree model allows shared parts of different actions to be evaluated only once by representing belief propagation of posterior factor graphs from the current factor graph sequentially. The FGP action tree has a consecutive hierarchic form, with intermediate vertices that represent beliefs after applying only part of a candidate action. Further, we use the aforementioned incremental covariance recovery technique to efficiently calculate covariances at intermediate and final beliefs within the FGP action tree, doing so independently of state dimension n . The calculated covariance entries let us reason about probabilistic properties of the beliefs and actions within the FGP action tree. Specifically, this allowed us to efficiently calculate the information impact of all candidate actions by reusing calculation from candidates' mutual parts. Overall, our method involves two passes over the FGP action tree, bottom-to-top query of required covariance entries and top-to-bottom propagation of these covariance entries. We evaluated the proposed approach in simulation considering the problem of autonomous navigation in unknown environments, and showed it reduces runtime by half compared with our previous approach (Kopitkov and Indelman, 2017), *rAMD L*.

There are several avenues for future research to take the proposed concept of reusing calculations between different candidate actions further. In this work, we considered a specific realization of the FGP action tree, using the structure of candidate trajectories in an autonomous navigation scenario. However, given posterior factor graphs for different candidate actions, multiple FGP action trees can be constructed. A key question that will be addressed as part of future research is how to construct an FGP action tree so that most of actions' similarity would be exploited. Another

direction for future research is to make BSP runtime complexity totally independent of state dimension. The process of information evaluation via FGP action tree consists of only a single, one-time, calculation that depends on state dimension n , i.e., recovering the prior marginal (or conditional) covariance entries of variables *involved* in candidate actions. When there are many candidates (above 500 in our simulations), this one-time part is insignificant with respect to the overall process time and can be ignored. However, for a smaller number of candidate actions this part takes considerable time (about 50% of time in our simulation in Figure 11) and additional research efforts are required in order to reduce its time complexity. Anticipating what state variables would be involved in the near-future candidate actions and incrementally tracking specific covariance entries of these involved variables during the inference process through our incremental covariance recovery method may be an appropriate solution and will be investigated in the future.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the Israel Science Foundation.

References

- Agha-Mohammadi A-A, Chakravorty S and Amato NM (2014) FIRM: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements. *The International Journal of Robotics Research* 33(2): 268–304.
- Amestoy P, Davis T and Duff I (1996) An approximate minimum degree ordering algorithm. *SIAM Journal on Matrix Analysis and Applications* 17(4): 886–905.
- Bai Z, Fahey G and Golub G (1996) Some large-scale matrix computation problems. *Journal of Computational and Applied Mathematics* 74(1): 71–89.
- Chaves SM and Eustice RM (2016) Efficient planning with the Bayes tree for active SLAM. In: *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 4664–4671.

- Dellaert F (2012) Factor graphs and GTSAM: A hands-on introduction. Technical Report GT-RIM-CP&R-2012-002, Georgia Institute of Technology.
- Dellaert F and Kaess M (2006) Square Root SAM: Simultaneous localization and mapping via square root information smoothing. *The International Journal of Robotics Research* 25(12): 1181–1203.
- Dunn E and Frahm J-M (2009) Next best view planning for active model improvement. In: *British Machine Vision Conference (BMVC)*.
- Golub G and Plemmons R (1980) Large-scale geodetic least-squares adjustment by dissection and orthogonal decomposition. *Linear Algebra and Its Applications* 34: 3–28.
- He R, Brunskill E and Roy N (2011) Efficient planning under uncertainty with macro-actions. *Journal of Artificial Intelligence Research* 40: 523–570.
- Ila V, Polok L, Solomy M, Smrz P and Zemcik P (2015) Fast covariance recovery in incremental nonlinear least square solvers. In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 4636–4643.
- Ila V, Polok L, Solomy M and Svoboda P (2017) SLAM++ - a highly efficient and temporally scalable incremental SLAM framework. *The International Journal of Robotics Research* 36(2): 210–230.
- Indelman V, Carlone L and Dellaert F (2015) Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments. *The International Journal of Robotics Research* 34(7): 849–882.
- Kaelbling LP, Littman ML and Cassandra AR (1998) Planning and acting in partially observable stochastic domains. *Artificial intelligence* 101(1): 99–134.
- Kaess M and Dellaert F (2009) Covariance recovery from a square root information matrix for data association. *Robotics and Autonomous Systems* 57(12): 1198–1210.
- Kaess M, Johannsson H, Roberts R, Ila V, Leonard J and Dellaert F (2012) iSAM2: Incremental smoothing and mapping using the Bayes tree. *The International Journal of Robotics Research* 31: 217–236.
- Kaess M, Ranganathan A and Dellaert F (2008) iSAM: Incremental smoothing and mapping. *IEEE Transactions on Robotics* 24(6): 1365–1378.
- Kavraki L, Svestka P, Latombe J-C and Overmars M (1996) Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation* 12(4): 566–580.
- Kim A and Eustice RM (2014) Active visual SLAM for robotic area coverage: Theory and experiment. *The International Journal of Robotics Research* 34(4–5): 457–475.
- Kopitkov D and Indelman V (2017) No belief propagation required: Belief space planning in high-dimensional state spaces via factor graphs, matrix determinant lemma and re-use of calculation. *The International Journal of Robotics Research* 36(10): 1088–1130.
- Kschischang F, Frey B and Loeliger H-A (2001) Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory* 47(2): 498–519.
- Levine D and How JP (2013) Sensor selection in high-dimensional Gaussian trees with nuisances. In: *Advances in Neural Information Processing Systems (NIPS)*, pp. 2211–2219.
- Patil S, Kahn G, Laskey M, Schulman J, Goldberg K and 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 (WAFR)*, pp. 515–533.
- Pineau J, Gordon GJ and Thrun S (2006) Anytime point-based approximations for large POMDPs. *Journal of Artificial Intelligence Research* 27: 335–380.
- Platt R, Tedrake R, Kaelbling L and Lozano-Pérez T (2010) Belief space planning assuming maximum likelihood observations. In: *Robotics: Science and Systems (RSS)*, Zaragoza, Spain, pp. 587–593.
- Prentice S and Roy N (2009) The belief roadmap: Efficient planning in belief space by factoring the covariance. *The International Journal of Robotics Research* 28(11–12): 1448–1465.
- Stachniss C, Grisetti G and Burgard W (2005) Information gain-based exploration using Rao-Blackwellized particle filters. In: *Robotics: Science and Systems (RSS)*, pp. 65–72.
- Van Den Berg J, Patil S and Alterovitz R (2012) Motion planning under uncertainty using iterative local optimization in belief space. *The International Journal of Robotics Research* 31(11): 1263–1278.
- Wenhardt S, Deutsch B, Angelopoulou E and Niemann H (2007) Active visual object reconstruction using d-, e-, and t-optimal next best views. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Appendix A. Proof of Lemma 1: Non-augmented case

The variables set W in this case is $\{Y_{old}, {}^I X\} = \{Y, {}^I X\}$. Define the prior marginal covariance matrices: $\Sigma_-^I \equiv \Sigma_-^{M, {}^I X}$, $\Sigma_-^Y \equiv \Sigma_-^{M, Y}$. Also denote the prior cross-covariance between Y and ${}^I X$ as Σ_-^C . Then, $\Sigma_-^{M, W}$ will have the following form:

$$\Sigma_-^{M, W} = \begin{pmatrix} \Sigma_-^Y & \Sigma_-^C \\ (\Sigma_-^C)^T & \Sigma_-^I \end{pmatrix} \quad (63)$$

In addition, let us separate prior (old) state variables X_- into *involved* ${}^I X$ (in new factors F_{new}) and *not involved* ${}^{-I} X$ variables. Similarly, let us partition the Jacobian matrix A into

$$A = \begin{pmatrix} {}^{-I} A & {}^I A \end{pmatrix} = \begin{pmatrix} 0 & {}^I A \end{pmatrix} \quad (64)$$

where ${}^{-I} A$ contains noise-weighted Jacobians with respect to X^{-I} , and ${}^I A$ with respect to ${}^I X$. From its definition we can conclude that ${}^{-I} A$ contains only zeros.

Next, using the Woodbury matrix identity and information update equation $\Lambda_+ = \Lambda_- + A^T \cdot A$, the posterior covariance matrix is

$$\begin{aligned} \Sigma_+ &= (\Lambda_+)^{-1} = (\Lambda_- + A^T \cdot A)^{-1} = \Sigma_- - \Sigma_- \cdot A^T \\ &\cdot [I_m + A \cdot \Sigma_- \cdot A^T]^{-1} \cdot A \cdot \Sigma_- = \Sigma_- - \Sigma_- \cdot A^T \\ &\cdot [I_m + {}^I A \cdot \Sigma_-^I \cdot ({}^I A)^T]^{-1} \cdot A \cdot \Sigma_- \end{aligned} \quad (65)$$

where $A \cdot \Sigma_- \cdot A^T = {}^I A \cdot \Sigma_-^I \cdot ({}^I A)^T$ because of A 's sparsity structure.

Then Σ_+ can be calculated as

$$\Sigma_+ = \Sigma_- - \Sigma_- \cdot A^T \cdot C^{-1} \cdot A \cdot \Sigma_- \quad (66)$$

$$C = I_m + {}^I A \cdot \Sigma_-^I \cdot ({}^I A)^T \quad (67)$$

Further, $\Sigma_+^{M,Y}$ can be calculated by retrieving from Σ_+ rows and columns that belong to variables Y :

$$\begin{aligned} \Sigma_+^{M,Y} &= \Sigma_-^{M,Y} - \Sigma_-^{(Y,:)} \cdot A^T \cdot C^{-1} \cdot A \cdot \Sigma_-^{(:,Y)} = \Sigma_-^{M,Y} \\ &\quad - [A \cdot \Sigma_-^{(:,Y)}]^T \cdot C^{-1} \cdot [A \cdot \Sigma_-^{(:,Y)}] \\ &= \Sigma_-^{M,Y} - [{}^I A \cdot (\Sigma_-^C)^T]^T \cdot C^{-1} \cdot [{}^I A \cdot (\Sigma_-^C)^T] = \Sigma_-^{M,Y} \\ &\quad - [\Sigma_-^C \cdot ({}^I A)^T]^T \cdot C^{-1} \cdot [\Sigma_-^C \cdot ({}^I A)^T]^T \end{aligned} \quad (68)$$

where using Matlab syntax we have $\Sigma_-^{(Y,:)} \doteq \Sigma_-(Y, :)$ and $\Sigma_-^{(:,Y)} \doteq \Sigma_-(:, Y)$. ■

Note that the columns inside information matrices do not have to be ordered in any particular way, and that the above-provided proof is correct for any ordering whatsoever.

Appendix B. Proof of Lemma 2: Rectangular case

In this case we can partition variables set Y into two subsets $Y_{old} \doteq X_- \cap Y$ and $Y_{new} \doteq X_{new} \cap Y$ or, in other words, into old and new state variables. The posterior marginal covariance matrix $\Sigma_+^{M,Y}$ will have then the following form:

$$\Sigma_+^{M,Y} = \begin{pmatrix} \Sigma_+^{M,Y_{old}} & \Sigma_+^{(Y_{old}, Y_{new})} \\ (\Sigma_+^{(Y_{old}, Y_{new})})^T & \Sigma_+^{M,Y_{new}} \end{pmatrix} \quad (69)$$

and we are looking for an efficient way to calculate matrices $\Sigma_+^{M,Y_{old}}$, $\Sigma_+^{M,Y_{new}}$ and $\Sigma_+^{(Y_{old}, Y_{new})}$.

The variables set W in this case is $\{Y_{old}, {}^I X\}$. Define next the prior marginal covariance matrices: $\Sigma_-^I \equiv \Sigma_-^{M, {}^I X}$, $\Sigma_-^{Y_{old}} \equiv \Sigma_-^{M, Y_{old}}$. Also denote the prior cross-covariance between Y_{old} and ${}^I X$ as Σ_-^C . Then, $\Sigma_-^{M,W}$ will have the following form:

$$\Sigma_-^{M,W} = \begin{pmatrix} \Sigma_-^{Y_{old}} & \Sigma_-^C \\ (\Sigma_-^C)^T & \Sigma_-^I \end{pmatrix} \quad (70)$$

In addition, let us separate prior (old) state variables X_- into *involved* ${}^I X$ (in new factors F_{new}) and *not involved* ${}^{-I} X$. The posterior state vector is then $X_+ = \{{}^I X, {}^{-I} X, X_{new}\}$. Similarly, let us partition the Jacobian matrix A into

$$A = (A_{old} \ A_{new}), \quad A_{old} = ({}^{-I} A \quad {}^I A) = (0 \quad {}^I A) \quad (71)$$

where A_{old} contains noise-weighted Jacobians with respect to old variables X_- , A_{new} with respect to new variables X_{new} , ${}^{-I} A$ with respect to ${}^{-I} X$, and ${}^I A$ with respect to ${}^I X$.

From its definition we can conclude that ${}^{-I} A$ contains only zeros.

Following the information update equation $\Lambda_+ = \Lambda_+^{Aug} + A^T \cdot A$ (see also Figure 3), the posterior information matrix can be partitioned using separation $X_+ = \{X_-, X_{new}\}$ as

$$\Lambda_- = \begin{pmatrix} \Lambda_- + A_{old}^T \cdot A_{old} & A_{old}^T \cdot A_{new} \\ A_{new}^T \cdot A_{old} & A_{new}^T \cdot A_{new} \end{pmatrix} \quad (72)$$

Now, let us partition the posterior covariance matrix Σ_+ in a similar way:

$$\Sigma_+ = \begin{pmatrix} \Sigma_+^{old} & \Sigma_+^{cross} \\ (\Sigma_+^{cross})^T & \Sigma_+^{new} \end{pmatrix} \quad (73)$$

Given the setup until now, we derive each of the matrices $\Sigma_+^{M,Y_{old}}$, $\Sigma_+^{M,Y_{new}}$, and $\Sigma_+^{(Y_{old}, Y_{new})}$ from (69) using parts from $\Sigma_-^{M,W}$ defined in (70).

$\Sigma_+^{M,Y_{new}}$ By using block-wise matrix inversion (which is based on the notion of Schur complements), Σ_+^{new} is equal to

$$\begin{aligned} \Sigma_+^{new} &= (A_{new}^T \cdot A_{new} - A_{new}^T \cdot A_{old} \cdot (\Lambda_- + A_{old}^T \cdot A_{old})^{-1} \\ &\quad \cdot A_{old}^T \cdot A_{new})^{-1} = (A_{new}^T \cdot (I_m - A_{old} \\ &\quad \cdot (\Lambda_- + A_{old}^T \cdot A_{old})^{-1} \cdot A_{old}^T) \cdot A_{new})^{-1} \end{aligned} \quad (74)$$

Now, let us define matrix C as follows:

$$C \triangleq I_m + A_{old} \cdot \Sigma_- \cdot A_{old}^T = I_m + {}^I A \cdot \Sigma_-^I \cdot ({}^I A)^T \quad (75)$$

Through the Woodbury matrix identity it can be easily shown that C 's inverse is

$$C^{-1} = I - A_{old} \cdot (\Lambda_- + A_{old}^T \cdot A_{old})^{-1} \cdot A_{old}^T \quad (76)$$

Therefore, Σ_+^{new} is equal to

$$\Sigma_+^{new} = (A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1} \quad (77)$$

and $\Sigma_+^{M,Y_{new}}$ can be calculated in the following way (note that such calculation's complexity is independent of state dimension):

$$\Sigma_+^{M,Y_{new}} = P^{(Y_{new}, :)}, \quad P \triangleq [(A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1}]^{(:, Y_{new})} \quad (78)$$

where in brackets we are using Matlab syntax to index relevant rows/columns. Note that $[(A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1}]^{(:, Y_{new})}$ can be calculated without calculation of full $(A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1}$, by using backslash operator in Matlab:

$$P = [A_{new}^T \cdot C^{-1} \cdot A_{new}] \setminus I^{(:, Y_{new})} \quad (79)$$

where $I^{(:, Y_{new})}$ are particular columns of the identity matrix. where

$\Sigma_+^{M, Y_{old}}$

Using the block-wise matrix inversion again we know that Σ_+^{old} from (73) is equal to

$$\begin{aligned}\Sigma_+^{old} &= (\Lambda_- + A_{old}^T \cdot A_{old} - A_{old}^T \cdot A_{new} \cdot (A_{new}^T \cdot A_{new})^{-1} \\ &\quad \cdot A_{new}^T \cdot A_{old})^{-1} \\ &= (\Lambda_- + A_{old}^T \cdot (I_m - A_{new} \cdot (A_{new}^T \cdot A_{new})^{-1} \cdot A_{new}^T) \\ &\quad \cdot A_{old})^{-1} = (\Lambda_k + A_{old}^T \cdot K \cdot A_{old})^{-1}\end{aligned}\quad (80)$$

with

$$\begin{aligned}K &\triangleq I_m - A_{new} \cdot (A_{new}^T \cdot A_{new})^{-1} \cdot A_{new}^T \\ &= I_m - A_{new} \cdot F \cdot A_{new}^T, \quad F \triangleq (A_{new}^T \cdot A_{new})^{-1}\end{aligned}\quad (81)$$

where K is a singular, symmetric, idempotent projection matrix, with properties $K = K^2$ and $K = K^T$.

Further, Σ_+^{old} can be now rewritten as

$$\begin{aligned}\Sigma_+^{old} &= (\Lambda_- + A_{old}^T \cdot K^T \cdot K \cdot A_{old})^{-1} \\ &= \Lambda_-^{-1} - \Lambda_-^{-1} \cdot A_{old}^T \cdot K^T \\ &\quad \cdot (I_m + K \cdot A_{old} \cdot \Lambda_-^{-1} \cdot A_{old}^T \cdot K^T)^{-1} \\ &\quad \cdot K \cdot A_{old} \cdot \Lambda_-^{-1} = \Sigma_- - \Sigma_- \cdot A_{old}^T \cdot K^T \\ &\quad \cdot (I_m + K \cdot A_{old} \cdot \Sigma_- \cdot A_{old}^T \cdot K^T)^{-1} \\ &\quad \cdot K \cdot A_{old} \cdot \Sigma_- = \Sigma_- - \Sigma_- \cdot A_{old}^T \\ &\quad \cdot K^T \cdot G^{-1} \cdot K \cdot A_{old} \cdot \Sigma_-\end{aligned}\quad (82)$$

with

$$\begin{aligned}G &\triangleq I_m + K \cdot A_{old} \cdot \Sigma_- \cdot A_{old}^T \cdot K^T = I_m + K \cdot I A \\ &\quad \cdot \Sigma_-^I \cdot (I A)^T \cdot K^T = I_m + K_1 \cdot \Sigma_-^I \cdot K_1^T\end{aligned}\quad (83)$$

$$K_1 \triangleq K \cdot I A \quad (84)$$

where K_1 are non-zero columns from A_{old} projected outside of vector space that is spanned by columns in A_{new} . In other words, K_1 contains information from A_{old} that is not contained within A_{new} .

Then $\Sigma_+^{M, Y_{old}}$ can be calculated by retrieving from Σ_+^{old} rows and columns that belong to variables Y_{old} :

$$\begin{aligned}\Sigma_+^{M, Y_{old}} &= \Sigma_-^{M, Y_{old}} - \Sigma_-^{(Y_{old}, :)} \cdot A_{old}^T \cdot K^T \cdot G^{-1} \\ &\quad \cdot K \cdot A_{old} \cdot \Sigma_-^{(:, Y_{old})} = \Sigma_-^{M, Y_{old}} - [K \cdot A_{old} \cdot \Sigma_-^{(:, Y_{old})}]^T \cdot G^{-1} \\ &\quad \cdot [K \cdot A_{old} \cdot \Sigma_-^{(:, Y_{old})}] \\ &= \Sigma_-^{M, Y_{old}} - [K \cdot I A \cdot \Sigma_-^{(X^I, Y_{old})}]^T \cdot G^{-1} \cdot [K \cdot I A \cdot \Sigma_-^{(X^I, Y_{old})}] \\ &= \Sigma_-^{M, Y_{old}} - [K_1 \cdot (\Sigma_-^C)^T]^T \cdot G^{-1} \cdot [K_1 \cdot (\Sigma_-^C)^T] \\ &= \Sigma_-^{M, Y_{old}} - B \cdot G^{-1} \cdot B^T\end{aligned}\quad (85)$$

$$B \triangleq \Sigma_-^C \cdot K_1^T = \Sigma_-^C \cdot (I A)^T \cdot K \quad (86)$$

$\Sigma_+^{(Y_{old}, Y_{new})}$: **Method 1**

Using the block-wise matrix inversion again we know that Σ_+^{cross} from (73) is equal to

$$\begin{aligned}\Sigma_+^{cross} &= -(\Lambda_- + A_{old}^T \cdot A_{old})^{-1} \cdot A_{old}^T \cdot A_{new} \\ &\quad \cdot (A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1} \\ &= -(\Sigma_- - \Sigma_- \cdot A_{old}^T \cdot C^{-1} \cdot A_{old} \cdot \Sigma_-) \\ &\quad \cdot A_{old}^T \cdot A_{new} \cdot (A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1} \\ &= -\Sigma_- \cdot A_{old}^T \cdot A_{new} \cdot (A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1} \\ &\quad + \Sigma_- \cdot A_{old}^T \cdot C^{-1} \cdot A_{old} \cdot \Sigma_- \cdot A_{old}^T \cdot A_{new} \\ &\quad \cdot (A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1} \\ &= \Sigma_- \cdot A_{old}^T \cdot [-I_m + C^{-1} \cdot A_{old} \cdot \Sigma_- \cdot A_{old}^T] \\ &\quad \cdot A_{new} \cdot (A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1} \\ &= \Sigma_- \cdot A_{old}^T \cdot [C^{-1} \cdot I A \cdot \Sigma_-^I \cdot (I A)^T - I_m] \cdot A_{new} \\ &\quad \cdot (A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1}\end{aligned}\quad (87)$$

where matrix C is defined in (75)

Then $\Sigma_+^{(Y_{old}, Y_{new})}$ can be calculated by retrieving from Σ_+^{cross} the entries that correspond to Y_{old} rows and Y_{new} columns:

$$\begin{aligned}\Sigma_+^{(Y_{old}, Y_{new})} &= \Sigma_-^{(Y_{old}, :)} \cdot A_{old}^T \cdot [C^{-1} \cdot I A \cdot \Sigma_-^I \cdot (I A)^T - I_m] \\ &\quad \cdot A_{new} \cdot [(A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1}]^{(:, Y_{new})} \\ &= \Sigma_-^C \cdot (I A)^T \cdot [C^{-1} \cdot I A \cdot \Sigma_-^I \cdot (I A)^T - I_m] \\ &\quad \cdot A_{new} \cdot P\end{aligned}\quad (88)$$

where matrix P is defined in (78).

$\Sigma_+^{(Y_{old}, Y_{new})}$: **Method 2**

Using another form of block-wise matrix inversion, Σ_+^{cross} from (73) is equal to

$$\begin{aligned}\Sigma_+^{cross} &= -[\Sigma_- - \Sigma_- \cdot A_{old}^T \cdot K^T \cdot G^{-1} \cdot K \cdot A_{old} \cdot \Sigma_-] \\ &\quad \cdot A_{old}^T \cdot A_{new} \cdot (A_{new}^T \cdot A_{new})^{-1} \\ &= [\Sigma_- \cdot A_{old}^T \cdot K^T \cdot G^{-1} \cdot K \cdot A_{old} \cdot \Sigma_- \cdot A_{old}^T - \Sigma_- \cdot A_{old}^T] \\ &\quad \cdot A_{new} \cdot (A_{new}^T \cdot A_{new})^{-1} \\ &= [\Sigma_- \cdot A_{old}^T \cdot K^T \cdot G^{-1} \cdot K \cdot A_{old} \cdot \Sigma_- \cdot A_{old}^T - \Sigma_- \cdot A_{old}^T] \\ &\quad \cdot A_{new} \cdot F \\ &= [\Sigma_- \cdot A_{old}^T \cdot K^T \cdot G^{-1} \cdot K \cdot I A \cdot \Sigma_-^I \cdot (I A)^T - \Sigma_- \cdot A_{old}^T] \\ &\quad \cdot A_{new} \cdot F\end{aligned}\quad (89)$$

where matrix F is defined in (81).

Then $\Sigma_+^{(Y_{old}, Y_{new})}$ can be calculated by retrieving from Σ_+^{cross} the entries that correspond to Y_{old} rows and Y_{new} columns:

$$\begin{aligned}\Sigma_+^{(Y_{old}, Y_{new})} &= [\Sigma_-^{(Y_{old}, :)} \cdot A_{old}^T \cdot K^T \cdot G^{-1} \cdot K^I A \cdot \Sigma_-^I \\ &\cdot ({}^I A)^T - \Sigma_-^{(Y_{old}, :)} \cdot A_{old}^T] \cdot A_{new} \cdot F^{(:, Y_{new})} \\ &= [\Sigma_-^C \cdot ({}^I A)^T \cdot K^T \cdot G^{-1} \cdot K^I A \cdot \Sigma_-^I \cdot ({}^I A)^T - \Sigma_-^C \cdot ({}^I A)^T] \\ &\cdot A_{new} \cdot F^{(:, Y_{new})} \\ &= \Sigma_-^C \cdot [({}^I A)^T \cdot K^T \cdot G^{-1} \cdot K^I A \cdot \Sigma_-^I - I_k] \cdot ({}^I A)^T \cdot A_{new} \\ &\cdot F^{(:, Y_{new})} \\ &= \Sigma_-^C \cdot [K_1^T \cdot G^{-1} \cdot K_1 \cdot \Sigma_-^I - I_k] \cdot ({}^I A)^T \cdot A_{new} \cdot F^{(:, Y_{new})}\end{aligned}\quad (90)$$

where matrix K_1 is defined in (84) and the identity matrix I_k has dimension $|{}^I X|$. ■

Note that the columns inside information matrices do not have to be ordered in any particular way, and that the above-provided proof is correct for any ordering whatsoever.

Appendix C. Proof of Lemma 3: Squared case

The *squared* case is a special instance of the *rectangular* case and, thus, we will use here the same setup as for the *rectangular* case. In other words, we use the partitioning that was defined in (69), (70), and (71).

In the *squared* case, we have that $m = |X_{new}|$ from which we can conclude that matrix A_{new} from (71) is a squared matrix. Then, matrix K from (26) is equal to a zero matrix:

$$\begin{aligned}K &= I_m - A_{new} \cdot (A_{new}^T \cdot A_{new})^{-1} \cdot A_{new}^T \\ &= I_m - A_{new} \cdot A_{new}^{-1} \cdot (A_{new}^T)^{-1} \cdot A_{new}^T = 0\end{aligned}\quad (91)$$

Further, matrices K_1 and B from (27) and (28) contain only zeros, and $\Sigma_+^{M, Y_{old}}$ is equal to

$$\Sigma_+^{M, Y_{old}} = \Sigma_-^{M, Y_{old}} - B \cdot G^{-1} \cdot B^T = \Sigma_-^{M, Y_{old}}\quad (92)$$

Next, Σ_+^{new} from (77) can be calculated as

$$\begin{aligned}\Sigma_+^{new} &= (A_{new}^T \cdot C^{-1} \cdot A_{new})^{-1} = A_{new}^{-1} \cdot C \cdot (A_{new}^T)^{-1} \\ &= A_{new}^{-1} \cdot C \cdot (A_{new}^{-1})^T\end{aligned}\quad (93)$$

and $\Sigma_+^{M, Y_{new}}$ is equal to

$$\begin{aligned}\Sigma_+^{M, Y_{new}} &= [A_{new}^{-1}]^{(Y_{new}, :)} \cdot C \cdot [(A_{new}^{-1})^T]^{(:, Y_{new})} \\ &= [A_{new}^{-1}]^{(Y_{new}, :)} \cdot C \cdot [A_{new}^{-1}]^{(Y_{new}, :)}{}^T = A_{iv} \cdot C \cdot A_{iv}^T\end{aligned}\quad (94)$$

where

$$A_{iv} \triangleq [A_{new}^{-1}]^{(Y_{new}, :)}\quad (95)$$

and can be efficiently calculated through the Matlab backslash operator:

$$A_{iv} = A_{new} \setminus I^{(:, Y_{new})}\quad (96)$$

Next, we can reduce (31) to

$$\begin{aligned}\Sigma_+^{(Y_{old}, Y_{new})} &= \Sigma_-^C \cdot [K_1^T \cdot G^{-1} \cdot K_1 \cdot \Sigma_-^I - I_k] \\ &\cdot ({}^I A)^T \cdot A_{new} \cdot F^{(:, Y_{new})} \\ &= -\Sigma_-^C \cdot ({}^I A)^T \cdot A_{new} \cdot F^{(:, Y_{new})} \\ &= -\Sigma_-^C \cdot ({}^I A)^T \cdot A_{new} \cdot F^{(:, Y_{new})} \\ &= -\Sigma_-^C \cdot ({}^I A)^T \cdot A_{new} \cdot A_{new}^{-1} \cdot (A_{new}^T)^{-1}]^{(:, Y_{new})} \\ &= -\Sigma_-^C \cdot ({}^I A)^T \cdot (A_{new}^T)^{-1}]^{(:, Y_{new})} \\ &= -\Sigma_-^C \cdot ({}^I A)^T \cdot [(A_{new}^T)^{-1}]^{(:, Y_{new})} = -\Sigma_-^C \cdot ({}^I A)^T \\ &\cdot [(A_{new}^{-1})^{(Y_{new}, :)}]{}^T = -\Sigma_-^C \cdot ({}^I A)^T \cdot (A_{iv})^T\end{aligned}\quad (97)$$

Note that the columns inside information matrices do not have to be ordered in any particular way, and that the above-provided proof is correct for any ordering whatsoever. ■

Appendix D. Proof of Lemma 4: relinearization case

As we saw in (40), the information update here has the following form:

$$\Lambda_+ = \Lambda_- + B^T \cdot B, \quad B \triangleq \begin{pmatrix} iA_- \\ A_+ \end{pmatrix}\quad (98)$$

First, denote by ${}^I A_-$ the non-zero columns of A_- and by ${}^I A_+$ the non-zero columns of A_+ (note that indices of such columns are the same in both A_- and A_+). Next, apply Lemma 1 as follows:

$$\begin{aligned}\Sigma_+^{M, Y} &= \Sigma_-^Y - V \cdot C^{-1} \cdot V^T, \quad V \triangleq \Sigma_-^C \cdot ({}^I B)^T, \\ C &\triangleq I_m + {}^I B \cdot \Sigma_-^I \cdot ({}^I B)^T, \quad {}^I B \triangleq \begin{pmatrix} i{}^I A_- \\ {}^I A_+ \end{pmatrix}\end{aligned}\quad (99)$$

where $\Sigma_-^I \equiv \Sigma_-^{M, {}^I X}$ is the prior marginal covariance of variables ${}^I X$ involved in the relinearized factors F_R ; $\Sigma_-^Y \equiv \Sigma_-^{M, Y}$ is the prior marginal covariance of variables Y ; Σ_-^C is the cross-covariance between Y and X^I .

Next we can see that

$$\begin{aligned}C &= I_m + \begin{pmatrix} i{}^I A_- \\ {}^I A_+ \end{pmatrix} \cdot \Sigma_-^I \cdot \begin{pmatrix} i \cdot ({}^I A_-)^T & ({}^I A_+)^T \end{pmatrix} \\ &= \begin{pmatrix} I - {}^I A_- \cdot \Sigma_-^I \cdot ({}^I A_-)^T & i{}^I A_- \cdot \Sigma_-^I \cdot ({}^I A_+)^T \\ i{}^I A_+ \cdot \Sigma_-^I \cdot ({}^I A_-)^T & I + {}^I A_+ \cdot \Sigma_-^I \cdot ({}^I A_+)^T \end{pmatrix} \\ &\triangleq \begin{pmatrix} C_{11} & i \cdot C_{12} \\ i \cdot C_{12}^T & C_{22} \end{pmatrix}\end{aligned}\quad (100)$$

where I is the identity matrix of appropriate dimension. Note that while C_{11} can sometimes not be a positive-definite (PD) matrix, C_{22} is always PD and therefore is invertible.

Next, we use a block-wise inversion in order to calculate C^{-1} :

$$C^{-1} \triangleq \begin{pmatrix} C_{11}^{inv} & C_{12}^{inv} \\ (C_{12}^{inv})^T & C_{22}^{inv} \end{pmatrix}, \quad (101)$$

$$C_{11}^{inv} \triangleq (C_{11} + C_{12} \cdot C_{22}^{-1} \cdot C_{12}^T)^{-1}, \quad (102)$$

$$C_{12}^{inv} \triangleq -i \cdot C_{11}^{inv} \cdot C_{12} \cdot C_{22}^{-1}, \quad (103)$$

$$C_{22}^{inv} \triangleq C_{22}^{-1} - C_{22}^{-1} \cdot C_{12}^T \cdot C_{11}^{inv} \cdot C_{12} \cdot C_{22}^{-1} \quad (104)$$

Using the above notation we can see that

$$\begin{aligned} ({}^l B)^T \cdot C^{-1} \cdot {}^l B &= \begin{pmatrix} i \cdot ({}^l A_-)^T & ({}^l A_+)^T \end{pmatrix} \\ &\cdot \begin{pmatrix} C_{11}^{inv} & C_{12}^{inv} \\ (C_{12}^{inv})^T & C_{22}^{inv} \end{pmatrix} \cdot \begin{pmatrix} i \cdot {}^l A_- \\ {}^l A_+ \end{pmatrix} \\ &= ({}^l A_+)^T \cdot C_{22}^{-1} \cdot {}^l A_+ - ({}^l A_-)^T \cdot C_{11}^{inv} \cdot {}^l A_- + ({}^l A_-)^T \\ &\cdot C_{11}^{inv} \cdot C_{12} \cdot C_{22}^{-1} \cdot {}^l A_+ \\ &+ ({}^l A_+)^T \cdot C_{22}^{-1} \cdot C_{12}^T \cdot C_{11}^{inv} \cdot {}^l A_- - ({}^l A_+)^T \cdot C_{22}^{-1} \cdot C_{12}^T \\ &\cdot C_{11}^{inv} \cdot C_{12} \cdot C_{22}^{-1} \cdot {}^l A_+ \\ &= ({}^l A_+)^T \cdot C_{22}^{-1} \cdot {}^l A_+ - [({}^l A_-)^T - ({}^l A_+)^T \cdot C_{22}^{-1} \cdot C_{12}^T] \\ &\cdot C_{11}^{inv} \cdot [({}^l A_-)^T - ({}^l A_+)^T \cdot C_{22}^{-1} \cdot C_{12}^T]^T \\ &= ({}^l A_+)^T \cdot C_{22}^{-1} \cdot {}^l A_+ - [({}^l A_-)^T - ({}^l A_+)^T \cdot C_{22}^{-1} \cdot C_{12}^T] \\ &\cdot [C_{11} + C_{12} \cdot C_{22}^{-1} \cdot C_{12}^T]^{-1} \\ &\cdot [({}^l A_-)^T - ({}^l A_+)^T \cdot C_{22}^{-1} \cdot C_{12}^T]^T \\ &= ({}^l A_+)^T \cdot [I + {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_+)^T]^{-1} \cdot {}^l A_+ \\ &- [({}^l A_-)^T - ({}^l A_+)^T \cdot [I + {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_+)^T]^{-1}] \cdot {}^l A_+ \\ &\cdot \Sigma_-^l \cdot ({}^l A_-)^T \cdot J^{-1} \cdot [({}^l A_-)^T - ({}^l A_+)^T \\ &\cdot [I + {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_+)^T]^{-1}] \cdot {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_-)^T]^T \end{aligned} \quad (105)$$

$$\begin{aligned} J &\triangleq C_{11} + C_{12} \cdot C_{22}^{-1} \cdot C_{12}^T = I - {}^l A_- \cdot \Sigma_-^l \cdot ({}^l A_-)^T + {}^l A_- \\ &\cdot \Sigma_-^l \cdot ({}^l A_+)^T \cdot [I + {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_+)^T]^{-1} \cdot {}^l A_+ \\ &\cdot \Sigma_-^l \cdot ({}^l A_-)^T \end{aligned} \quad (106)$$

Next, introduce new notation

$$M_2 \triangleq ({}^l A_+)^T / chol[I + {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_+)^T], \quad (107)$$

$$G \triangleq M_2^T \cdot \Sigma_-^l \cdot ({}^l A_-)^T \quad (108)$$

where $chol(\cdot)$ represents Cholesky decomposition, which returns an upper triangular matrix, and “/” is the backslash operator from Matlab syntax ($A/B = A \cdot B^{-1}$).

It can be clearly seen that

$$M_2 \cdot M_2^T = ({}^l A_+)^T \cdot [I + {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_+)^T]^{-1} \cdot {}^l A_+ \quad (109)$$

$$\begin{aligned} G^T \cdot G &= {}^l A_- \cdot \Sigma_-^l \cdot ({}^l A_+)^T \cdot [I + {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_+)^T]^{-1} \cdot {}^l A_+ \\ &\cdot \Sigma_-^l \cdot ({}^l A_-)^T, \end{aligned} \quad (110)$$

$$M_2 \cdot G = ({}^l A_+)^T \cdot [I + {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_+)^T]^{-1} \cdot {}^l A_+ \cdot \Sigma_-^l \cdot ({}^l A_-)^T \quad (111)$$

Using matrices M_2 and G , we can rewrite an expression for $({}^l B)^T \cdot C^{-1} \cdot {}^l B$ as

$$\begin{aligned} ({}^l B)^T \cdot C^{-1} \cdot {}^l B &= M_2 \cdot M_2^T - [({}^l A_-)^T - M_2 \cdot G] \\ &\cdot J^{-1} \cdot [({}^l A_-)^T - M_2 \cdot G]^T, \end{aligned} \quad (112)$$

$$J = I - {}^l A_- \cdot \Sigma_-^l \cdot ({}^l A_-)^T + G^T \cdot G \quad (113)$$

Next, let us define another matrix:

$$M_1 \triangleq [({}^l A_-)^T - M_2 \cdot G] / chol[I - {}^l A_- \cdot \Sigma_-^l \cdot ({}^l A_-)^T + G^T \cdot G] \quad (114)$$

with $M_1 \cdot M_1^T$ being equal to

$$\begin{aligned} M_1 \cdot M_1^T &= [({}^l A_-)^T - M_2 \cdot G] \\ &\cdot [I - {}^l A_- \cdot \Sigma_-^l \cdot ({}^l A_-)^T + G^T \cdot G]^{-1} \\ &\cdot [({}^l A_-)^T - M_2 \cdot G]^T = [({}^l A_-)^T - M_2 \cdot G] \\ &\cdot J^{-1} \cdot [({}^l A_-)^T - M_2 \cdot G]^T \end{aligned} \quad (115)$$

Thus, we will have

$$\begin{aligned} ({}^l B)^T \cdot C^{-1} \cdot {}^l B &= M_2 \cdot M_2^T - M_1 \cdot M_1^T = M \cdot M^T, \\ M &\triangleq (iM_1 \quad M_2) \end{aligned} \quad (116)$$

By combining the above equation with (99) we can conclude that

$$\begin{aligned} \Sigma_+^{M,Y} &= \Sigma_-^Y - \Sigma_-^C \cdot M \cdot M^T \cdot (\Sigma_-^C)^T = \Sigma_-^Y - U \cdot U^T, \\ U &\triangleq \Sigma_-^C \cdot M \end{aligned} \quad (117)$$

Note that the columns inside information matrices do not have to be ordered in any particular way, and that the above-provided proof is correct for any ordering whatsoever. ■

Appendix E. Sum of information gains

Consider action a with increment $I(a) = \{F_{new}, X_{new}\}$. Further, consider specific partitioning of a into sub-actions $a = \{a'_1, \dots, a'_k\}$ where each sub-action a'_i has increment $I_i(a'_i) = \{F_{i,new}, X_{i,new}\}$. The factor sets $F_{i,new}$ are disjoint, so are the new variable sets $X_{i,new}$. In addition, for proper action partitioning we will have $\bigcup_{i=1}^k F_{i,new} = F_{new}$ and $\bigcup_{i=1}^k X_{i,new} = X_{new}$.

Next, we will prove that the IG of a is equal to sum of IGs of sub-actions $\{a'_i\}_{i=1}^k$ in the *unfocused* scenario. A similar proof can also be shown for *focused* BSP.

The *unfocused* IG of action a by definition is

$$J_{IG}(a) = \mathcal{H}(b[X_-]) - \mathcal{H}(b[X_+]) \quad (118)$$

where $b[X_-]$ is a prior belief before applying action a , and $b[X_+]$ is a posterior belief after applying it.

In addition, denote the posterior belief of each sub-action a'_i as $b_i[X_{i,+}]$. When applying sub-actions consecutively in sequence, belief propagation will have the following form:

$$\begin{aligned} b[X_-] &\Rightarrow b_1[X_{1,+}] \Rightarrow b_2[X_{2,+}] \Rightarrow \dots \Rightarrow b_{k-1}[X_{k-1,+}] \\ &\Rightarrow b[X_+] \end{aligned} \quad (119)$$

Then, the IG of each sub-action is equal to

$$J_{IG}(a'_1) = \mathcal{H}(b[X_-]) - \mathcal{H}(b_1[X_{1,+}])$$

$$J_{IG}(a'_2) = \mathcal{H}(b_1[X_{1,+}]) - \mathcal{H}(b_2[X_{2,+}])$$

...

$$J_{IG}(a'_{k-1}) = \mathcal{H}(b_{k-2}[X_{k-2,+}]) - \mathcal{H}(b_{k-1}[X_{k-1,+}])$$

$$J_{IG}(a'_k) = \mathcal{H}(b_{k-1}[X_{k-1,+}]) - \mathcal{H}(b[X_+])$$

and sum of these IGs is equal to

$$\sum_{i=1}^k J_{IG}(a'_i) = \mathcal{H}(b[X_-]) - \mathcal{H}(b[X_+]) = J_{IG}(a) \quad (120)$$

■