Towards Efficient Inference Update through Planning via JIP - Joint Inference and Belief Space Planning

Elad I. Farhi and Vadim Indelman

***Abstract*— Inference and decision making under uncertainty are essential in numerous robotics problems. In recent years, the similarities between inference and control triggered much work, from developing unified computational frameworks to pondering about the duality between the two. In spite of the aforementioned efforts, inference and control, as well as inference and belief space planning (BSP) are still treated as two separate processes. In this paper we propose a novel approach that utilizes the similarities between inference and BSP and make the *key observation* that inference can be efficiently updated using the precursory planning stage, thus paving the way towards a joint inference and BSP paradigm. We develop four different methods that implement our novel approach under simplifying assumptions and validate them in the context of autonomous navigation in unknown environment. Results indicate that not only our methods improve running time by at least two orders of magnitude, compared to iSAM2 paradigm, they also found to be less sensitive to state dimensionality and loop closures.**

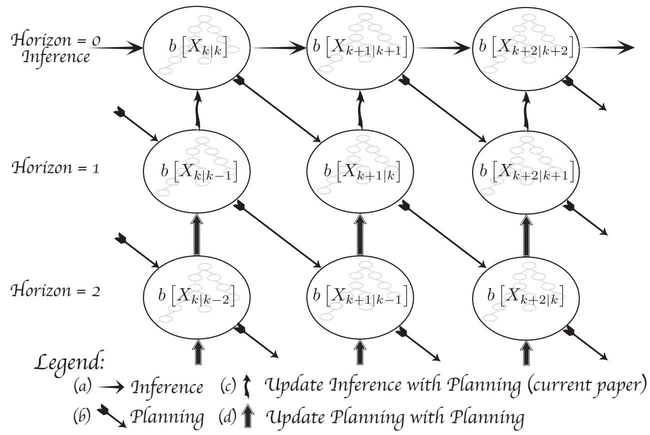

Fig. 1: Visualization of *JIP*, a novel approach to address both inference and belief space planning under a single process. Here $b[X_{k+1|k}]$ stands for the belief of the joint state in time instance $k+1$ while current time is k and each row stands for a different planning horizon. The relations between different beliefs in the graph are denoted by different arrows. (a) Inference (b) Planning step (c) Updating Inference with precursory planning (d) Update planning with precursory planning.

I. Introduction

Computationally efficient inference and decision making under uncertainty are required in numerous problem domains in robotics, with applications such as autonomous navigation and SLAM, search and rescue scenarios, object manipulation and robot-assisted surgery. Until recently, the two processes, inference and decision making, have been treated separately: the inference stage maintains a belief over variables of interest (e.g. robot poses) given available information, while decision making under uncertainty and the related belief space planning problem are entrusted with determining the best next action(s) to realize a certain objective.

The inference problem has been addressed by the research community extensively over the past decades. In particular, focus was given to inference over high-dimensional state spaces, with SLAM being a representative problem, and to computational efficiency to facilitate online operation, as required in numerous robotics systems. Over the years, the solution paradigm for the inference problem has evolved. From EKF based methods (e.g. [4], [8]), through information form recursive (e.g. [23]) and smoothing methods (e.g. [6], [7])), and in recent years up to incremental smoothing approaches, such as iSAM [13] and iSAM2 [12].

Given a belief from the inference stage, decision making under uncertainty and Belief Space Planning (BSP) approaches reason about belief evolution for different candidate actions while taking into account different sources of uncertainty. The corresponding problem is an instantiation of a Partially Observable Markov Decision Process (POMDP) problem, which is known to be computationally intractable [1]. Over the years, numerous approaches have been developed that trade-off suboptimal performance with reduced computational complexity, see e.g. [9], [17], [19]. While the majority of these approaches, including [2], [20], [21], [26], assumed some sources of absolute information (GPS, known landmarks) are available or considered the environment to be known, recent research relaxed these assumptions, accounting for the uncertainties in the mapped environment thus far as part of the decision making process [10], [14] at the price of increased state dimensionality.

Interestingly, be the decision making approach as it may, it has to solve numerous inference problems in order to determine the (sub)optimal actions. Obvious similarities between inference and decision making, triggered much work in recent years. For example, Kobilarov et al. [15] and Ta et al. [22] developed Differential Dynamic Programming (DDP) and Factor Graph (FG) based unified computational frameworks, respectively, for inference and decision making. Toussaint and Storkey [25] provided an approximate solution to the Markov Decision Process (MDP) problem using inference optimization methods, and Todorov [24] investigated

Elad I. Farhi is with the Technion Autonomous Systems Program (TASP), Technion - Israel Institute of Technology, Haifa 32000, Israel, eladf@campus.technion.ac.il. Vadim Indelman is with the Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel, vadim.indelman@technion.ac.il.

the duality between optimal control and inference for the MDP case. While both Toussaint and Storkey [25] and Todorov [24] refer to the MDP case, in this paper we address the more general case where the state is partially observable and thus has to be inferred probabilistically within a POMDP framework. Despite the aforementioned research efforts, however, inference and BSP are still handled as two separate processes.

Our *key observation* is that similarities between inference and decision making paradigms can be utilized in order to save valuable computation time. In this paper, we provide an approach to save precious computation time in the inference update stage by reusing calculations from precursory planning. Our approach is rooted in the new concept of *JIP*, a Joint Inference and belief space Planning approach, which strives to handle both inference and decision making as a single process, thus sharing and updating similar calculations becomes trivial. In contrast to the notion of joint inference and control, which considers an MDP setting, the JIP paradigm considers a partially observable setting (POMDP).

Figure 1 provides a graphical representation of *JIP*. The joint inference and belief space planning approach incorporates both inference and decision making stages into a single process. Each node in the graph represents a belief, i.e. $b[X_{k+1|k}]$ denotes the joint belief of state X at a future time instant $k+1$ given that the current time is k. The right facing arrows i.e. (a) in Figure 1 denote inference at sequential time instances. The diagonal arrows i.e. (b) in Figure 1 represent optimal controls that lead up to the appropriate beliefs. The upward facing arrows denote either future belief update i.e. (d) in Figure 1 or inference update i.e. (c) in Figure 1, while both by using precursory future belief.

We provide a novel paradigm for saving computation time by updating each inference stage with its precursory planning stage, corresponds to the upward facing arrow (c) in Figure 1. We present four methods that utilize calculation from the planning stage in inference, along with a thorough analysis and comparison of those methods to the regular batch inference and the efficient iSAM paradigms.

To summarize, our contributions in this paper are as follows: (a) We present JIP, a novel approach to address both Inference and Planning as a single process, which enables to share information between the two in a single paradigm. (b) We introduce a novel approach for saving computation time during the inference stage by reusing calculations made during the precursory planning stage. (c) We provide four methods, that utilize our main contribution. (d) We evaluate our methods and compare them to the state of the art in simulation considering the problem of autonomous navigation in unknown environments.

II. Notations and Problem Formulation

The joined Inference and belief space Planning problem consists out of two main successive and recursive stages, the Inference stage and the Planning stage. The Inference stage produces a state estimate using all information up to present time, while the Planning stage produces the next control action using the former inference stage information and a Cost/Reward function. Both the inference and belief space planning stages will be reviewed in the following sections.

Let x_t denote the robot's state in time instance t and $\mathcal{L}$ represent the world state if the latter is uncertain or unknown. For example, for SLAM problem, it could represent landmark locations. The joint state, up to and including time t, is defined as $X_t = \{x_0, ..., x_t, \mathcal{L}\} \in \mathbb{R}^n$. Throughout this paper we shall use the notation $t|k$ to refer to some time instant t while considering information up to time k - as will be shown in the sequel, this notation will allow to refer to *sequential* inference and planning phases in a unified manner.

Let $z_{t|k}$ and $u_{t|k}$ denote, respectively, the obtained measurements and the applied control action at time t, while the current time is k. For example, $z_{k+1|k}$ represents measurements from a future time instant $k+1$ while $z_{k-1|k}$ represents measurements from a past time instant $k-1$, with the present time being k in both cases. For a stochastic state sequence, the conditional probability density function (pdf) over the joint state can be formulated. Representing the measurements and controls up to time t, given current time is k, as

$$z_{1:t|k} \doteq \{z_{1|k}, ..., z_{t|k}\} \ , \ u_{0:t-1|k} \doteq \{u_{0|k}, ..., u_{t-1|k}\}, \quad (1)$$

the joint pdf, the *belief*, is given by

$$b[X_{t|k}] \doteq \mathbb{P}(X_t|z_{1:t|k}, u_{0:t-1|k}). \quad (2)$$

Eq. (2) represents the posterior at current time k for $t = k$, while for $t > k$ it represents the posterior from the planning phase for a specific sequence of future actions and observations. Using Bayes rule, Eq. (2) can be rewritten as

$$\mathbb{P}(X_t|z_{1:t|k}, u_{0:t-1|k}) \propto \mathbb{P}(x_0) \cdot \prod_{i=1}^{t} \left[\mathbb{P}(x_i|x_{i-1}, u_{i-1|k}) \prod_{j \in \mathcal{M}_{i|k}} \mathbb{P}(z^j_{i|k}|x_i, l_j) \right], \quad (3)$$

where $\mathbb{P}(x_0)$ is the prior on x_0, and $\mathbb{P}(x_i|x_{i-1}, u_{i-1|k})$ and $\mathbb{P}(z^j_{i|k}|x_i, l_j)$ denote, respectively, the motion model and measurement likelihood terms. Data association given information up to time k, i.e. landmark indices observed at each time i, is denoted by the set $\mathcal{M}_{i|k}$, while a particular observation of landmark j is represented by $z^j_{i|k} \in z_{i|k}$. The motion and observation models from Eq. (3) are assumed to be with additive Gaussian noise

$$x_{i+1} = f(x_i, u_i) + w_i \quad , \quad w_i \sim \mathcal{N}(0, \Sigma_w) \quad (4)$$

$$z^j_i = h(x_i, l_j) + v_i \quad , \quad v_i \sim \mathcal{N}(0, \Sigma_v), \quad (5)$$

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

A. Inference

For the inference problem $t \leq k$, i.e time instances that are equal or smaller than current time. The maximum a posteriori

(MAP) estimate of the joint state X_k is given by

$$X^{\star}_{k|k} = \underset{X_k}{\arg\max}\ \mathbb{P}(X_k|z_{1:k|k}, u_{0:k-1|k}) \tag{6}$$

The MAP estimate from Eq. (6) is referred to as the *inference solution*, in which, all controls and observations until time instance k are known. Introducing Eqs. (3-5) into Eq. (6) and taking the negative logarithm yields the following non-linear least squares problem (NLS)

$$X^{\star}_{k|k} = \underset{X_k}{\arg\min} \|x_0 - x_0^{\star}\|^2_{\Sigma_0} + \sum_{i=1}^{k}\left[\|x_i - f(x_{i-1}, u_{i-1|k})\|^2_{\Sigma_w} + \sum_{j\in\mathcal{M}_i|k}\|z^j_{i|k} - h(x_i, l_j)\|^2_{\Sigma_v}\right] \tag{7}$$

where $\|a\|^2_{\Sigma} \doteq a^T\Sigma^{-1}a$ is the Mahalanobis norm. Linearizing each of the terms in Eq. (7) and performing standard algebraic manipulations (see, e.g., [10]) yields

$$\Delta X^{\star}_{k|k} = \underset{\Delta X_k}{\arg\min} \|A_{k|k}\Delta X_k - b_{k|k}\|^2 \tag{8}$$

where $A_{k|k} \in \mathbb{R}^{m\times n}$ is the Jacobian matrix and $b_{k|k} \in \mathbb{R}^m$ is the right hand side (RHS) vector. In a more elaborated representation

$$A_{k|k} = \begin{bmatrix} \Sigma_0^{-\frac{1}{2}} \\ \mathcal{F}_{1:k|k} \\ \mathcal{H}_{1:k|k} \end{bmatrix} \quad , \quad b_{k|k} = \begin{bmatrix} 0 \\ \breve{b}^{\mathcal{F}}_{1:k|k} \\ \breve{b}^{\mathcal{H}}_{1:k|k} \end{bmatrix} \tag{9}$$

where $\mathcal{F}_{1:k|k}$, $\mathcal{H}_{1:k|k}$, $\breve{b}^{\mathcal{F}}_{1:k|k}$ and $\breve{b}^{\mathcal{H}}_{1:k|k}$ denote the Jacobian matrices and RHS vectors of all motion and observation terms accordingly, for time instances $1:k$ when current time is k. These Jacobians, with the RHS can be referred to by

$$\mathcal{A}_{1:k|k} = \begin{bmatrix} \mathcal{F}_{1:k|k} \\ \mathcal{H}_{1:k|k} \end{bmatrix} \quad , \quad \breve{b}_{1:k|k} = \begin{bmatrix} \breve{b}^{\mathcal{F}}_{1:k|k} \\ \breve{b}^{\mathcal{H}}_{1:k|k} \end{bmatrix} \tag{10}$$

While there are few methods to solve Eq. (8), we choose QR factorization as presented, e.g., in [13] The QR factorization of the Jacobian matrix $A_{k|k}$ is given by the Orthonormal rotation matrix $Q_{k|k}$ and the upper triangular matrix $R_{k|k}$

$$A_{k|k} = Q_{k|k}R_{k|k} \tag{11}$$

Eq. (11) is introduced into Eq. (8) thus producing

$$R_{k|k}\Delta X_k = d_{k|k}, \tag{12}$$

where $R_{k|k}$ is un upper triangular matrix and $d_{k|k}$ is the corresponding RHS vector, given by the original RHS vector and the orthonormal rotation matrix $Q_{k|k}$

$$d_{k|k} \doteq Q^T_{k|k}b_{k|k} \tag{13}$$

We can now solve Eq. (12) for ΔX_k via back substitution, updating the linearization point, and repeat the process until convergence. One can substantially reduce running time by exploiting sparsity and updating the QR factorization from the previous step with new information instead of calculating a factorization from scratch, see e.g. iSAM2 algorithm [12]. The information matrix is thus given by

$$\Lambda_{k|k} = A^T_{k|k}A_{k|k} \tag{14}$$

To summarize this section, the belief $b[X_{k|k}]$ can be represented as the Gaussian

$$b[X_{k|k}] \doteq \mathbb{P}(X_k|z_{1:k|k}, u_{0:k-1|k}) = \mathcal{N}(X^{\star}_{k|k}, \Lambda^{-1}_{k|k}). \tag{15}$$

B. Planning in the Belief Space

Finite horizon belief space planning for L look ahead steps involves inference over the beliefs $b[X_{k+L|k}]$

$$b[X_{k+l|k}] = \mathbb{P}(X_{k+l}|z_{1:k+l|k}, u_{0:k+l-1|k}) \tag{16}$$

with $l \in [k+1, k+L]$, and where we use the same notation as in Eq. (2) to denote the current time is k. The belief (16) can be written recursively as a function of the belief $b[X_{k|k}]$ from the inference phase as

$$b[X_{k+l|k}] = b[X_{k|k}] \cdot \prod_{i=k+1}^{k+l}\left[\mathbb{P}(x_i|x_{i-1}, u_{i-1|k}) \prod_{j\in\mathcal{M}_{i|k}} \mathbb{P}(z^j_{i|k}|x_i, l_j)\right], \tag{17}$$

for the considered action sequence $u_{k:k+l-1|k}$ at planning time k, and observations $z_{k+1:k+l|k}$ that are expected to be obtained upon execution of these actions. The set $\mathcal{M}_{i|k}$ denotes landmark indices that are expected to be observed at a future time instant i.

One can now define an objective function

$$J(u_{k-1:k+L-1|k}) \doteq \underset{z_{k+1:k+L|k}}{\mathbb{E}} \sum_{i=k+1}^{k+L} c_i\left(b[X_{i|k}], u_{i-1|k}\right), \tag{18}$$

with immediate costs (or rewards) c_i and where the expectation considers all the possible realizations of the future observations $z_{k+1:k+L|k}$. Conceptually, one could also reason whether these observations will actually be obtained, e.g. by considering also different realizations of $\mathcal{M}_{i|k}$. Note that for information-theoretic costs (e.g. entropy) and Gaussian distributions considered herein, it can be shown that the expectation operator can be omitted (see e.g. [10]), while another alternative is to simulate future observations via sampling (see e.g. [18]), if such a simulator is available.

The optimal control can now be defined as

$$u^{\star}_{k:k+L-1|k} = \underset{u_{k:k+L-1|k}}{\arg\min}\ J(u_{k:k+L-1|k}). \tag{19}$$

Evaluating the objective function (18) for a candidate action sequence involves calculating belief evolution for different look ahead steps. An interesting insight, that we will exploit in the sequel, is that the underlying equations are similar to those seen in Section II-A.

In particular, evaluating the belief at the lth look ahead step, $b[X_{k+l|k}]$, involves MAP inference

$$X^{\star}_{k+L|k} = \underset{X_{k+L}}{\arg\min} \|X_k - X^{\star}_{k|k}\|^2_{\Lambda^{-1}_{k|k}} + \sum_{i=k+1}^{k+L}\left[\|x_i - f(x_{i-1}, u_{i-1|k})\|^2_{\Sigma_w} + \sum_{j\in\mathcal{M}_i|k}\|z^j_{i|k} - h(x_i, l_j)\|^2_{\Sigma_v}\right] \tag{20}$$

Following a similar procedure as in Section II-A, i.e. linearizing and re-writing in a matrix form, we get

$$\Delta X^{\star}_{k+L|k} = \underset{\Delta X_{k+L}}{\arg\min} \|A_{k+L|k}\Delta X_{k+L} - b_{k+L|k}\|^2. \quad (21)$$

The Jacobian matrix $A_{k+L|k}$ and RHS vector $b_{k+L|k}$ are defined as

$$A_{k+L|k} \doteq \begin{bmatrix} A_{k|k} \\ \mathcal{A}_{k+1:k+L|k} \end{bmatrix}, \; b_{k+L|k} \doteq \begin{bmatrix} b_{k|k} \\ \breve{b}_{k+1:k+L|k} \end{bmatrix}, \quad (22)$$

where $A_{k|k}$ and $b_{k|k}$ are taken from inference, see Eq. (8), and $\mathcal{A}_{k+1:k+L|k}$ and $\breve{b}_{k+1:k+L|k}$ correspond to the new terms obtained at the first L look ahead steps (e.g. see Eq. (10)). Note that $\mathcal{A}_{k+1:k+L|k}$ is not a function of the (unknown) measurements $z_{k+1:k+L|k}$.

Performing QR factorization, yields

$$A_{k+L|k} = Q_{k+L|k}R_{k+L|k}, \quad (23)$$

from which the information matrix, required in the information-theoretic term in cost, can be calculated.

C. *Problem Statement*

Our goal is to save computation time in the inference stage by reusing calculations made during the precursory planning stage. In other words, instead of performing inference over $b[X_{k+1|k+1}]$ at time $k+1$ by updating the belief $b[X_{k+1|k}]$ from the previous time with newly acquired information (actions, measurements), we investigate whether inference over $b[X_{k+1|k+1}]$ can be performed by exploiting $b[X_{k+1|k}]$, which was already calculated as part of the planning phase at time k. This concept of re-using planning for inference is denoted in Figure 1 by the upward facing arrow (c), which connects the appropriate first horizon of the precursory planning phase to the current inference phase. In terms of the graphical representation of the joint inference-planning system in Figure 1, we wish to "shortcut" the regular inference procedure using pre-calculated planning step and through it to save precious calculation time.

III. Approach

While different approaches exist for calculating the optimal actions (19), all such approaches necessarily calculate belief evolution, i.e. Eqs. (21)-(22), for the optimal action. This optimal action is then executed, and new observations are obtained at the next time instant $k+1$, followed by a new inference stage.

Our *key observation* is that the involved calculations in the inference stage (Section II-A) and belief evolution, as part of a precursory planning stage (Section II-B), are similar and therefore can be re-used. In terms of Figure 1, this observation refers to the beliefs from the first row (inference) and the second row (first look ahead step in precursory planning stage). We now discuss the actual difference between the corresponding beliefs, considering without loosing generality $b[X_{k+1|k+1}]$ and $b[X_{k+1|k}]$ from the inference and planning stages, respectively, where $b[X_{k+1|k}]$ was evaluated, as part of the planning stage, for the determined optimal action.

Both $b[X_{k+1|k+1}]$ and $b[X_{k+1|k}]$ are Gaussian beliefs. Thus, they are represented by the first two moments, see e.g. Eq. (15), that are calculated via MAP inference as a function of the square root information matrix R and RHS vector d. Performing inference over $b[X_{k+1|k+1}]$ involves calculating $R_{k+1|k+1}$ and $d_{k+1|k+1}$, as described in Section II-A. On the other hand, the terms $R_{k+1|k}$ and $d_{k+1|k}$ describing $b[X_{k+1|k}]$ are already available from the precursory planning stage.

Rather than updating $b[X_{k+1|k+1}]$ from $b[X_{k|k}]$ and the new information from time $k+1$, we propose to update $b[X_{k+1|k+1}]$ from $b[X_{k+1|k}]$, where similar calculations have already been made, so that a valuable computation time would be spared.

In this paper we make the following simplifying assumption that will be relaxed in future research: we assume that the considered data association in the first look ahead step of the planning stage at time k is identical to the actual data association in the succeeding inference stage at time $k+1$. We would denote this as the consistent data association assumption. Recalling the definition of $\mathcal{M}_{i|k}$ (see e.g. Eq. (3)), this assumption is equivalent to writing

$$\mathcal{M}_{k+1|k} \equiv \mathcal{M}_{k+1|k+1}. \quad (24)$$

In other words, landmarks considered to be observed at a future time $k+1$, will indeed be observed at that time. Note this does *not* necessarily imply that actual measurements and robot poses will be as considered within the planning stage.

We now observe that the motion models in both $b[X_{k+1|k+1}]$ and $b[X_{k+1|k}]$ are evaluated considering the *same* control (i.e. the optimal control $u^{\star}_k$). Moreover, the robot pose x_{k+1} is initialized to the *same* value in both cases as $f(x_k, u^{\star}_k)$, see e.g. [10], and thus the linearization point of all probabilistic terms in inference and planning is *identical*. This, together with the above assumption (i.e. Eq. (24) holds) allows us to write $A_{k+1|k} = A_{k+1|k+1}$, and hence

$$R_{k+1|k+1} \equiv R_{k+1|k}, \quad (25)$$

for the first iteration in the inference stage at time $k+1$.

Hence in order to solve $b[X_{k+1|k+1}]$ we are left to find the RHS vector $d_{k+1|k+1}$, while $R_{k+1|k+1}$ can be *entirely re-used*. In the sequel we present four methods that can be used for this purpose, and examine computational aspects of each. First, however, we discuss a graphical models perspective for the above concept, and relate it to the iSAM2 algorithm [12].

A. JIP - *Joint Inference and Belief Space Planning*

Figure 1 provides a 2D graphical representation of *JIP*. Our approach interleaves both inference and decision making stages into a single process.

Each node in the graph represents a belief, i.e. $b[X_{k+1|k}]$ denotes the joint belief of state X at a future time instant $k+1$ given that the current time is k. The right facing arrows i.e. (a) in Figure 1 denote inference at sequential time instances.

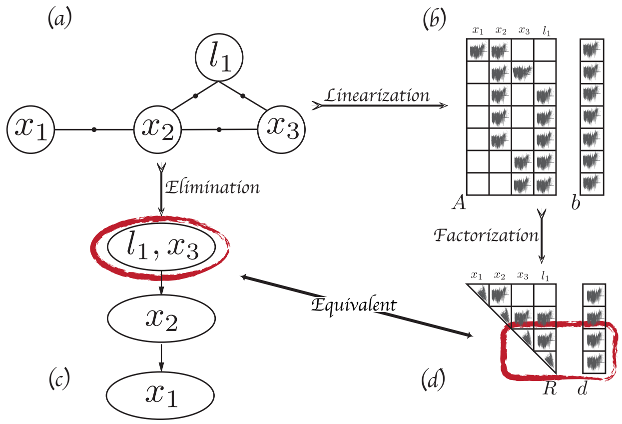

Fig. 2: The relations between different problem representations. (a) Factor graph (b) Jacobian matrix A with RHS vector b (c) Bayes Tree (d) Factorized Jacobian matrix R with equivalent RHS vector d.

From each node there is a possible planning action distinguished by different possible controls. The optimal planning action is denoted as a diagonal arrow in the 2D representation i.e. (b) in Figure 1. While the full 3D representation of *JIP* would contain more branches, going into and outside of the 2D plane, denoting all possible actions and the according beliefs. For simplicity reasons we would continue to refer only the 2D representation.

Each row in Figure 1 represents a different planning horizon step, while the first row denotes inference, each of the others denote future beliefs. Each column in Figure 1 contains beliefs which reason about the same time instance, but the deeper we go down the column, the older the information this reasoning is based upon, e.g $b[X_{k+1|k+1}]$, $b[X_{k+1|k}]$ and $b[X_{k+1|k-1}]$ all denote the joint belief of state X at time instant k+1 but for different current time instances. The upward facing arrows denote update using precursory belief, e.g. (d) in Figure 1 denote updating the current time the belief is based upon, when the upward facing arrows from the second row to the first i.e. (c) in Figure 1 denote updating the precursory first planning horizon step with the new obtained measurements. Following methods provides an implementation for (c) in Figure 1. While the implementation for (d) in Figure 1 is preserved for future work.

B. Graphical Models Perspective

The inference problem is naturally represented using graphical models such as factor graph [16] and Bayes Tree (BT) [11]. Figure 2 presents the connections between those different representations that are exploited by incremental smoothing approaches such as iSAM2 [12]. The factor graph ((a) in Figure 2) encodes the joint pdf (3), or equivalently for Gaussian distributions, the original non-linear least squares problem (7). On the other hand, BT represents a factorization of the joint pdf in terms of conditionals for a given variable elimination order. For Gaussian distributions, BT efficiently represents the square root information matrix R and the RHS vector d, where each clique over a subset of variables encodes non-zero entries in appropriate sub-blocks and entries of R and d, see illustration in (c) and (d) in Figure 2. When new information is received, the iSAM2 algorithm efficiently updates only the impacted parts in BT, an operation that corresponds to updating an existing factorization rather than calculating a new one from scratch.

Our approach can be also framed within the Bayes tree graphical model. The simplifying assumption regarding consistent data association between the inference stage and its precursory planning stage leads to Eq. (25). Hence, Bayes trees from both stages will have the same clique formation and entries that correspond to the R matrix, and will only differ by the RHS vector d. Thus, to get to the succeeding BT (inference) from the already-available BT (from planning), only appropriate values of the RHS vector d, that are stored within the impacted cliques, need to be updated. This observation suggests that the same iSAM2 machinery could be used to identify and appropriately update these entries within the d vector, further reducing running time. We leave further investigation of this direction to future research, and consider in this paper updating the entire d vector.

C. The Orthogonal Transformation Matrix Method - `OTM`

In the `OTM` method, we obtain $d_{k+1|k+1}$ following the definition as written in Eq. (13). Recall that at time $k+1$ in the inference stage, the posterior should be updated with new terms that correspond, for example, to motion model and obtained measurements. The RHS vector's augmentation, that corresponds to these new terms is denoted by $\breve{b}_{k+1|k+1}$, see Eq. (10). Given $R_{k|k}$ and $d_{k|k}$ from the inference stage at time k, the augmented system at time $k+1$ is

$$A_{k+1|k+1}\Delta X_{k+1} \doteq \begin{bmatrix} R_{k|k} \\ \mathcal{A}_{k+1|k+1} \end{bmatrix} \Delta X_{k+1} = \begin{bmatrix} d_{k|k} \\ \breve{b}_{k+1|k+1} \end{bmatrix}$$

which after factorization of $A_{k+1|k+1}$ (see Eqs. (11)-(13)) becomes

$$R_{k+1|k+1}\Delta X_{k+1} = d_{k+1|k+1}, \tag{26}$$

where

$$d_{k+1|k+1} = Q^T_{k+1|k} \begin{bmatrix} d_{k|k} \\ \breve{b}_{k+1|k+1} \end{bmatrix}. \tag{27}$$

The above calculation of $d_{k+1|k+1}$ requires $Q_{k+1|k+1}$. Since $A_{k+1|k} \equiv A_{k+1|k+1}$ (see Section III), we get $Q_{k+1|k+1} = Q_{k+1|k}$. However, $Q_{k+1|k}$, is already available from the precursory planning stage, see Eq. (23), and thus calculating $d_{k+1|k+1}$ via Eq. (27) does *not* involve QR factorization in practice.

D. The Downdate Update Method - `DU`

In the `DU` method we propose to re-use the $d_{k+1|k}$ vector from the planning stage to calculate $d_{k+1|k+1}$.

While not actually required within the planning stage, $d_{k+1|k}$ could be calculated at that stage from $b_{k+1|k}$

and $Q_{k+1|k}$, see Eqs. (22)-(23). However, $b_{k+1|k}$ (but not $A_{k+1|k}$) is a function of the unknown future observations $z_{k+1|k}$, which would seem to complicate things. Our solution to this issue is as follows: We assume *some* value for the observations $z_{k+1|k}$ and then calculate $d_{k+1|k}$ within the planning stage. As in inference at time $k+1$, the actual measurements $z_{k+1|k+1}$ will be different, we remove the contribution of $z_{k+1|k}$ to $d_{k+1|k}$ via information downdating [3], and then appropriately incorporate $z_{k+1|k+1}$ to get $d_{k+1|k+1}$.

More specifically, downdating the measurements $z_{k+1|k}$ from $d_{k+1|k}$ is done via (see e.g. [3])

$$d^{aug}_{k+1|k} = R^{aug^{-T}}_{k+1|k}\left(R^T_{k+1|k}d_{k+1|k} - \mathcal{A}^T_{k+1|k}\breve{b}_{k+1|k}\right), \quad (28)$$

where $\breve{b}_{k+1|k}$ is a function of $z_{k+1|k}$, see Eqs. (21)-(22), and where $R^{aug}_{k+1|k}$ is the downdated $R_{k+1|k}$ matrix which is given by

$$R^{aug^T}_{k+1|k}R^{aug}_{k+1|k} = A^T_{k+1|k}A_{k+1|k} - \mathcal{A}^T_{k+1|k}\mathcal{A}_{k+1|k}. \quad (29)$$

Interestingly, the above calculations are not really required: Since we already have $d_{k|k}$ from the previous inference stage, we can attain the downdated $d^{aug}_{k+1|k}$ vector more efficiently by augmenting $d_{k|k}$ with zero padding.

$$d^{aug}_{k+1|k} = \begin{bmatrix} d_{k|k} \\ 0 \end{bmatrix} \quad (30)$$

Similarly, $R^{aug}_{k+1|k}$ can be calculated as

$$R^{aug}_{k+1|k} = \begin{bmatrix} R_{k|k} & 0 \\ 0 & 0 \end{bmatrix}, \quad (31)$$

where $R_{k|k}$ is zero padded to match dimensions of $R_{k+1|k}$.

Now, all which is left to get $d_{k+1|k+1}$, is to incorporate the new measurements $z_{k+1|k+1}$ (encoded in $\breve{b}_{k+1|k+1}$). Following [3], this can be done via

$$d_{k+1|k+1} = R^{-T}_{k+1|k+1}\left(R^{aug^T}_{k+1|k}d^{aug}_{k+1|k} + \mathcal{A}^T_{k+1|k+1}\breve{b}_{k+1|k+1}\right),$$

where $\mathcal{A}_{k+1|k+1} \equiv \mathcal{A}_{k+1|k}$ due to the assumption regarding data association from Section III.

E. The OTM - Only Observations Method - `OTM-OO`

The `OTM-OO` method aspire to utilize more information from the planning stage. Since the motion models from inference and the precursory planning first step are identical, i.e. same function $f(.,.)$, see Eqs. (7) and (21), and as in both cases the *same* control is considered (the determined optimal action), there is no reason to change the motion model data from the RHS vector $d_{k+1|k}$. In order to enable the aforementioned, the planning stage as described in Section II-B has to be broken down into two stages, in which the motion and observation models are updated separately. The breakdown would impose no effect over the computation time or accuracy of the planning stage solution. So following Section III-C, we attain from planning the RHS vector already with the motion model ($d^{\mathcal{F}}_{k+1|k}$), augment it with the new measurements and rotate it with the corresponding rotation matrix obtained from the planning stage.

$$d_{k+1|k+1} = Q^{\mathcal{F}^T}_{k+1|k}\begin{bmatrix} d^{\mathcal{F}}_{k+1|k} \\ \breve{b}^{\mathcal{H}}_{k+1|k+1} \end{bmatrix} \quad (32)$$

While the rotation matrix $Q^{\mathcal{F}}_{k+1|k}$ is given from the precursory planning stage where

$$Q^{\mathcal{F}}_{k+1|k}R^{\mathcal{F}}_{k+1|k} = \begin{bmatrix} R_{k|k} \\ \mathcal{F}_{k+1|k} \end{bmatrix} = A^{\mathcal{F}}_{k+1|k} \quad (33)$$

F. The DU - Only Observations Method - `DU-OO`

The `DU-OO` method is a variant of the second method, where, similarly to Section III-E, we utilize the fact that there is no reason to change the motion model data from the RHS vector $d_{k+1|k}$. Hence we would downdate all data with the exception of the motion model, and then update accordingly. As opposed to the second method (Section III-D), now we do need to downdate using [3]

$$d^{\mathcal{F}}_{k+1|k} = R^{\mathcal{F}^{-T}}_{k+1|k}\left(R^T_{k+1|k}d_{k+1|k} - \mathcal{H}^T_{k+1|k}\breve{b}^{\mathcal{H}}_{k+1|k}\right), \quad (34)$$

where $d^{\mathcal{F}}_{k+1|k}$ is the RHS vector, downdated from all new measurements with the exception of the motion model and $R^{\mathcal{F}}_{k+1|k}$ is the equivalent downdated $R_{k+1|k}$ matrix which is given by

$$R^{\mathcal{F}^T}_{k+1|k}R^{\mathcal{F}}_{k+1|k} = A^T_{k+1|k}A_{k+1|k} - \mathcal{H}^T_{k+1|k}\mathcal{H}_{k+1|k}, \quad (35)$$

where $\mathcal{H}_{k+1|k}$ denotes the portion of the planning stage Jacobian, of the new factors with the exception of the motion model. Now, all which is left, is to update $d^{\mathcal{F}}_{k+1|k}$ with the new measurements from the inference stage.

$$d_{k+1|k+1} = R^{-T}_{k+1|k+1}\left(R^{\mathcal{F}^T}_{k+1|k}d^{\mathcal{F}}_{k+1|k} + \mathcal{H}^T_{k+1|k+1}\breve{b}^{\mathcal{H}}_{k+1|k+1}\right)$$

IV. Results

All four methods were implemented in *MATLAB*. In order to test our methods' running time we compared them to the batch inference update paradigm (here on denoted as `STD`) and to iSAM2 efficient inference update paradigm [12] (here on denoted as `iSAM`). The comparison was made in two different simulations. The presented running time is a result of an average between 10^3 repetitions per each method. The `iSAM` approach uses the `GTSAM C++` implementation with the supplied *MATLAB* wrapper [5]. Considering the general rule of thumb, that *MATLAB* implementation is at least one order of magnitude slower, the comparison to `iSAM` as a reference is extremely conservative. Both simulations were executed on the same Linux machine, with Xeon E3-1241v3 3.5 GHz processor with 32 GB of memory.

Each method is required to provide with $R_{k+1|k+1}$ and $d_{k+1|k+1}$ while only data that was already calculated during the precursory planning stage $k+1|k$ is given as an input. All other calculations are considered a part of the method hence timed along with it. The data association assumption (see Eq. (25)) is kept in both simulations.

We will first review the reference paradigms formulation, then present and discuss the results of each simulation

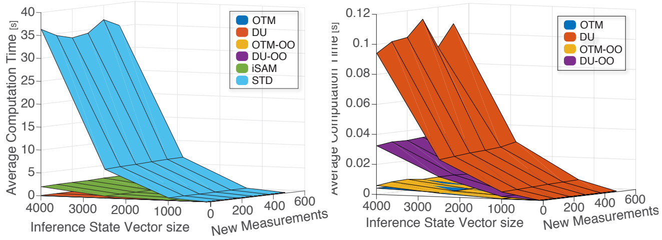

(a) **(b)**

Fig. 3: Method comparison through basic analysis simulation, checking sensitivity to new added measurements and the size of the inference state vector: (a) All the tested methods i.e. `STD`, `iSAM` and our four methods (b) Our four methods, i.e. `OTM`, `UD`, `OTM-OO` and `UD-OO`.

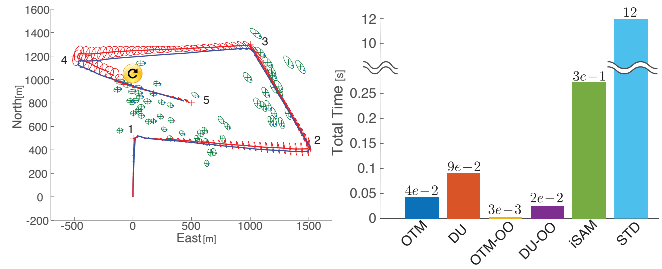

(a) **(b)**

Fig. 4: Second simulation layout and results: (a) The Synthetic Environment, where landmarks are marked in green, targets are numbered and marked with red crosses, the ground truth is denoted by a blue line, the trajectory is denoted by a red line while the covariance is visualized by red ellipse (b) Total average running time of inference update for each method.

In the batch inference update paradigm (`STD`), we follow Eqs. (9)-(13) for time instance $k+1|k+1$ to receive both the R matrix and RHS vector d.

Before covering `iSAM` paradigm there are few issues that need to be addressed. First for consistency reasons, since iSAM2 is formulated in terms of the FG and BT graphical models (see Section III-B), the underlying calculations were interpreted to the Jacobian matrix and its QR factorization. Secondly, again for consistency reasons, since other methods already consider the Jacobian is available, i.e. post linearization e.g. Eq. (21), in timing iSAM2 we do not include the linearization step so the comparison would be valid.

In iSAM2 efficient inference update paradigm (`iSAM`), we first augment the inference factorized Jacobian matrix and the RHS vector with new factors

$$A_{k+1|k+1} = \begin{bmatrix} R_{k|k} \\ \mathcal{A}_{k+1|k+1} \end{bmatrix}, \; b_{k+1|k+1} = \begin{bmatrix} d_{k|k} \\ \breve{b}_{k+1|k+1} \end{bmatrix} \quad (36)$$

Then we use Givens Rotations to incrementally update only relevant entries in Eq. (36), thus obtaining the factorized Jacobian and RHS.

A. First Simulation - Basic Analysis

The first simulation preforms a single horizon BSP calculation, followed by an inference step and a single inference update. The simulation provides a basic analysis of running time for each method, denoted by the z axis, for a fully dense information matrix and with no loop closures. The simulation analyzes the sensitivity of each method to the initial state vector size, denoted by the y axis, and to the number of new factors, denoted by x axis. Since we perform a single horizon step with a single inference update no re-linearization in necessary, hence iSAM comparison is valid.

Figure 3a presents average timing results for all methods. After inspecting the results, we found that for all methods, running time is a non linear, positive gradient function of the inference state vector size and a linear function of the number of new measurements. While the running time dependency over the number of new measurements grows with the inference state vector size. For all inspected parameters our methods scores the lowest running time with a difference of up to *three orders of magnitude* comparing to `iSAM`.

Figure 3b presents running time of our suggested methods. Interestingly, the OTM methodology proves to be more time efficient than the DU methodology, while for both the OO addition improves running time, thus scoring all methods from the fastest to the slowest with a time difference of *four orders of magnitude* between the opposites:

`OTM-OO` ⇒ `OTM` ⇒ `DU-OO` ⇒ `DU`⇒ `iSAM`⇒ `STD`

B. Second Simulation - BSP in Unknown Environment

The second simulation preforms continuous BSP in an unknown synthetic environment. In contrast to Section IV-A, since the synthetic environment replicates a real world scenario, the information matrix is of course sparse (e.g. Fig. 2). A robot was given five targets (see Figure 4a) while all landmarks were a-priori unknown, and was required to visit all targets whilst not crossing a covariance value threshold. The largest loop closure in the trajectory of the robot, and the first in a series of large loop closures, is denoted by a yellow ↻ sign across all relevant graphs. The robot preforms continuous BSP with a finite horizon of five look ahead steps (see [10]). During the inference update stage each of the aforementioned methods were timed performing a single inference update step.

Similarly to Section IV-A, as can be seen in Figure 4b, the suggested MATLAB implemented methods are up to *two orders of magnitude* faster than `iSAM` used in a MATLAB C++ wrapper. Interestingly, the use of sparse information matrices changed the methods timing hierarchy. While `OTM-OO` still has the best timing results (3×10^{-3} sec), *two orders of magnitude faster than* `iSAM`, `OTM` and `DU-OO` switched places. So the timing hierarchy from fastest to slowest is:

`OTM-OO`⇒ `DU-OO` ⇒ `OTM` ⇒ `DU`⇒ `iSAM`⇒ `STD`

After demonstrating our methods drastically reduce running time, we continue on to showing that in few aspects they are also less sensitive. From carefully inspecting Figure 5, alongside the trajectory of the robot in Figure 4a the following observations can be made. Our methods, in particularly

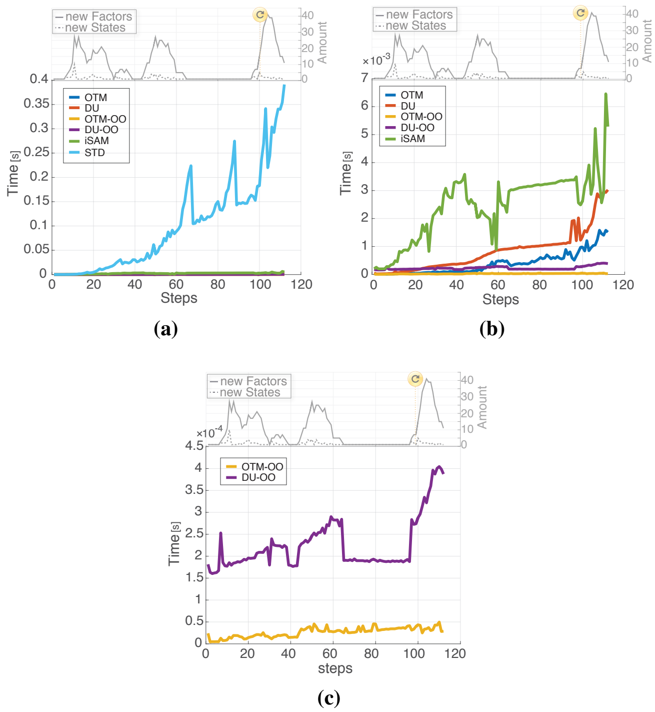

Fig. 5: Second simulation timing results for the scenario presented in Figure 4a. Upper part of each graph provides indication on new factors and new states per computation step while the lower presents the methods timing results: (a) All six methods (b) OTM, DU, OTM-OO, DU-OOand iSAM methods (c) OTM-OO and DU-OO methods.

OTM-OO seems to be impervious to large loop closures, variant state vector sizes, various new measurements or even the combination of the aforementioned.

V. Conclusions

In this paper we presented a novel approach that utilizes decision making calculations for inference and paves the way towards *JIP*-joint inference and BSP paradigm. Under the consistent data association assumption, we provided four different methods to update inference using information from precursory planning stage. We compared our methods to iSAM with continuous BSP in unknown environment, for both dense and sparse information matrix cases. We successfully showed that updating inference with precursory planning is more efficient by at least two orders of magnitude when compared to iSAM.

Future work will include relaxing the consistent data association assumption and utilizing iSAM methodologies for efficient selective updates.

References

[1] D. Bernstein, R. Givan, N. Immerman, and S. Zilberstein. The complexity of decentralized control of markov decision processes. *Mathematics of operations research*, 27(4):819–840, 2002.

[2] A. Bry and N. Roy. Rapidly-exploring random belief trees for motion planning under uncertainty. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 723–730, 2011.

[3] A. Cunningham, V. Indelman, and F. Dellaert. DDF-SAM 2.0: Consistent distributed smoothing and mapping. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Karlsruhe, Germany, May 2013.

[4] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *IEEE Trans. Pattern Anal. Machine Intell.*, 29(6):1052–1067, 2007.

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

[6] F. Dellaert and M. Kaess. Square Root SAM: Simultaneous localization and mapping via square root information smoothing. *Intl. J. of Robotics Research*, 25(12):1181–1203, Dec 2006.

[7] R.M. Eustice, H. Singh, and J.J. Leonard. Exactly sparse delayed-state filters for view-based SLAM. *IEEE Trans. Robotics*, 22(6):1100–1114, Dec 2006.

[8] Simon S Haykin et al. *Kalman filtering and neural networks*. Wiley Online Library, 2001.

[9] G. A. Hollinger and G. S. Sukhatme. Sampling-based robotic information gathering algorithms. *Intl. J. of Robotics Research*, pages 1271–1287, 2014.

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

[11] M. Kaess, V. Ila, R. Roberts, and F. Dellaert. The Bayes tree: An algorithmic foundation for probabilistic robot mapping. In *Intl. Workshop on the Algorithmic Foundations of Robotics*, Dec 2010.

[12] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert. iSAM2: Incremental smoothing and mapping using the Bayes tree. *Intl. J. of Robotics Research*, 31:217–236, Feb 2012.

[13] M. Kaess, A. Ranganathan, and F. Dellaert. iSAM: Incremental smoothing and mapping. *IEEE Trans. Robotics*, 24(6):1365–1378, Dec 2008.

[14] A. Kim and R. M. Eustice. Active visual slam for robotic area coverage: Theory and experiment. *Intl. J. of Robotics Research*, 2014.

[15] Marin Kobilarov, Duy-Nguyen Ta, and Frank Dellaert. Differential dynamic programming for optimal estimation. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 863–869. IEEE, 2015.

[16] F.R. Kschischang, B.J. Frey, and H-A. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Trans. Inform. Theory*, 47(2), February 2001.

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

[18] S. Pathak, A. Thomas, A. Feniger, and V. Indelman. Da-bsp: Towards data association aware belief space planning for robust active perception. In *Eur. Conf. on AI (ECAI)*, September 2016.

[19] J. Pineau, G. J. Gordon, and S. Thrun. Anytime point-based approximations for large pomdps. *J. of Artificial Intelligence Research*, 27:335–380, 2006.

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

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

[22] Duy-Nguyen Ta, Marin Kobilarov, and Frank Dellaert. A factor graph approach to estimation and model predictive control on unmanned aerial vehicles. In *International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 181–188. IEEE, 2014.

[23] S. Thrun, Y. Liu, D. Koller, A.Y. Ng, Z. Ghahramani, and H. Durrant-Whyte. Simultaneous localization and mapping with sparse extended information filters. *Intl. J. of Robotics Research*, 23(7-8):693–716, 2004.

[24] Emanuel Todorov. General duality between optimal control and estimation. In *IEEE Conference on Decision and Control*, pages 4286–4292. IEEE, 2008.

[25] Marc Toussaint and Amos Storkey. Probabilistic inference for solving discrete and continuous state markov decision processes. In *Intl. Conf. on Machine Learning (ICML)*, pages 945–952. ACM, 2006.

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