

The International Journal of Robotics Research

<http://ijr.sagepub.com/>

FIRM: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements

Ali-akbar Agha-mohammadi, Suman Chakravorty and Nancy M Amato
The International Journal of Robotics Research published online 15 November 2013
DOI: 10.1177/0278364913501564

The online version of this article can be found at:
<http://ijr.sagepub.com/content/early/2013/11/14/0278364913501564>

Published by:



<http://www.sagepublications.com>

On behalf of:



Multimedia Archives

Additional services and information for *The International Journal of Robotics Research* can be found at:

Email Alerts: <http://ijr.sagepub.com/cgi/alerts>

Subscriptions: <http://ijr.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

>> [OnlineFirst Version of Record](#) - Nov 15, 2013

[What is This?](#)

FIRM: Sampling-based feedback motion-planning under motion uncertainty and imperfect measurements

The International Journal of
Robotics Research
0(0) 1–37
© The Author(s) 2013
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/0278364913501564
ijr.sagepub.com



Ali-akbar Agha-mohammadi¹, Suman Chakravorty² and Nancy M Amato¹

Abstract

In this paper we present feedback-based information roadmap (FIRM), a multi-query approach for planning under uncertainty which is a belief-space variant of probabilistic roadmap methods. The crucial feature of FIRM is that the costs associated with the edges are independent of each other, and in this sense it is the first method that generates a graph in belief space that preserves the optimal substructure property. From a practical point of view, FIRM is a robust and reliable planning framework. It is robust since the solution is a feedback and there is no need for expensive replanning. It is reliable because accurate collision probabilities can be computed along the edges. In addition, FIRM is a scalable framework, where the complexity of planning with FIRM is a constant multiplier of the complexity of planning with PRM. In this paper, FIRM is introduced as an abstract framework. As a concrete instantiation of FIRM, we adopt stationary linear quadratic Gaussian (SLQG) controllers as belief stabilizers and introduce the so-called SLQG-FIRM. In SLQG-FIRM we focus on kinematic systems and then extend to dynamical systems by sampling in the equilibrium space. We investigate the performance of SLQG-FIRM in different scenarios.

Keywords

Planning, control, uncertainty, information, belief space

1. Introduction

Decision-making under uncertainty is a crucial ability for most robotic systems. In the presence of uncertainty in a robot's motion and uncertainty in its sensory readings, the true robot state is not available for decision-making purposes. In such cases, a state estimation module can provide a probability distribution over all possible states, referred to as *information state* or *belief*. Therefore, decision-making under motion and sensing uncertainties needs to be performed in the information space (belief space). In its most general form, this decision-making can be formulated as a partially observable Markov decision process (POMDP) problem (Astrom, 1965; Smallwood and Sondik, 1973; Kaelbling et al., 1998). However, only a very small class of problems formulated using POMDP can be solved exactly due to its computational complexity (Papadimitriou and Tsitsiklis, 1987; Madani et al., 1999). In particular, planning (i.e. solving POMDPs) over continuous state, control, and observation spaces is a big challenge.

On the other hand, in the absence of uncertainty, sampling-based path-planning algorithms including graph-based methods such as probabilistic roadmap methods (PRMs) (Kavraki et al., 1996) and their variants (see e.g. Amato et al., 1998), and tree-based methods such as

rapidly-exploring randomized trees (RRTs) (Lavalle and Kuffner, 2001), expansive space trees (Hsu, 2000) and their variants (e.g. Karaman and Frazzoli, 2011) have shown great success in solving robot motion planning problems. Nevertheless, direct transformation of the roadmap-based methods to planning under uncertainty (in belief space) is a challenge for two main reasons. The first issue is ensuring that the roadmap nodes are reachable. The second challenge is that the incurred costs on different edges of the roadmap depend on each other, which violates a basic assumption in roadmap-based methods that each roadmap edge represent an independent planning problem.

In this paper, we generalize the PRM framework to obtain the feedback-based information roadmap (FIRM) framework that takes into account both motion and sensing uncertainties. FIRM is constructed as a roadmap (graph)

¹Department of Computer Science and Engineering, Texas A&M University, USA

²Department of Aerospace Engineering, Texas A&M University, USA

Corresponding author:

Ali-akbar Agha-mohammadi, Department of Computer Science and Engineering, Texas A&M University, College Station, TX 77840, USA.

Email: aliagha@tamu.edu

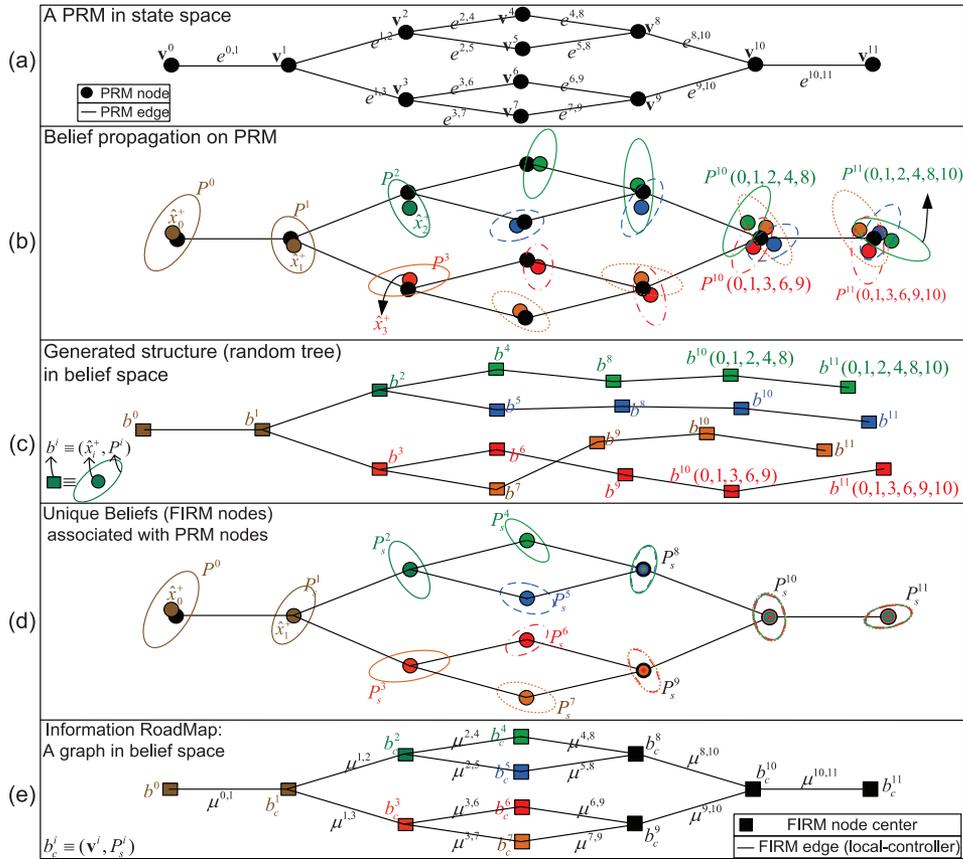


Fig. 1. (a) A simple PRM in state space. (b) Assuming Gaussian belief space, belief snapshots along different paths starting from v^0 and ending at v^{11} are shown. As can be seen, the obtained belief depends on the path traveled by the robot. For example, $P^{11}(0, 1, 3, 6, 9, 10)$ denotes the estimation covariance at node v^{11} , when the robot has traversed a path through nodes $(0, 1, 3, 6, 9, 10)$ prior to node 11. (c) Corresponding belief paths in the belief space. Belief at each node depends on the initial belief, actions taken (edges), and obtained observations (random). Therefore, the generated structure in the belief space is not a graph but a random tree. (d) Unique beliefs assigned to each PRM node. Using stabilizers, regardless of the action and observation history, the belief at each node stops at these predefined beliefs. (e) The FIRM corresponding to the given PRM; b_c^i denotes graph nodes in the belief space and μ^{ij} denotes local planners (graph edges).

in the belief space, where graph nodes are beliefs (rigorously speaking, small subsets of the belief space) and edges are local controllers in belief space. FIRM is an abstract generic framework that relies on the existence of an appropriate belief node sampler and connector (local controller). We also construct a stationary linear quadratic Gaussian controller-based (SLQG-based) instantiation of this generic framework, called SLQG-FIRM, where we provide a specific node sampler and connector. In SLQG-FIRM we first focus on the kinematic systems and then extend it to dynamical systems by restricting sampling space to the equilibrium space. SLQG-FIRM is the first method that generalizes the PRM to the belief space such that the incurred costs on different edges of the roadmap are independent of each other, while providing a straightforward approach to sample reachable belief nodes. This property is a direct consequence of utilizing feedback controllers in the construction of FIRM. Based on this property, the FIRM framework breaks the curse of history in POMDPs

(Pineau et al., 2003), and provides the optimal *feedback policy* over the roadmap instead of returning a single nominal path.

Figure 1 illustrates the problem of edge dependence in the direct transformation of PRM to stochastic domains. It also shows the approach of FIRM in generating a graph in belief space with independent edges. Figure 1(a) depicts a simple PRM in the state space with twelve nodes $\mathcal{V} = \{v^0, \dots, v^{11}\}$. Figure 1(b) shows the belief evolution on the underlying PRM. Assuming the belief is Gaussian in this example, we represent a point in belief space using a mean \hat{x}^+ and a covariance P , in other words, a belief b is characterized by the pair $b \equiv (\hat{x}^+, P)$. In Figure 1(b), mean values are shown by small filled circles, and covariance matrices are shown by their corresponding 3σ ellipses centered at the mean. We drive the system from v^0 toward the node v^{11} . The initial belief at node v^0 is $b^0 \equiv (\hat{x}_0^+, P^0)$. The belief propagation from left to right starting from b_0 is shown in Figure 1(b).

Although there exists a single edge $e^{(10,11)}$ between nodes v^{10} and v^{11} in PRM (see Figure 1(a)), the belief evolution along $e^{(10,11)}$ is not unique (see Figures 1(b)–(c)) since it depends on (i) the initial belief, (ii) obtained observations (observation history), and (iii) the path taken (action history) that has led to v^{10} . Figure 1(c) shows the corresponding belief propagation in the belief space, where each rectangle encodes a mean and covariance. As seen in Figure 1(c), the belief paths do not form a graph; rather, they form a random tree in belief space. Hence, in practice, where observations are random, not only does the number of possible beliefs grow exponentially, but the belief also evolves randomly.

Therefore, to predict edge costs, full knowledge of the belief at the start of the edge is required. This in turn requires full knowledge of the history of observations and actions leading up to the start of the edge.

Even if future observations were assumed to be deterministic for the purpose of planning, the generated structure would still be a tree that grows exponentially in the size of the underlying PRM graph.

In FIRM, we use local feedback planners to drive the belief process toward the predefined unique beliefs associated with PRM nodes (see Figure 1(d)). As a result, the evolution of belief after a FIRM node is reached is independent of the evolution of belief before that node is reached. This breaks the curse of history, allowing us to construct a PRM-like roadmap in the belief space with independent edge costs. Therefore, in contrast to the main body of the literature in motion planning under uncertainty, FIRM can be re-used for future queries and does not need to reconstruct the roadmap every time a new query is submitted.

From an algorithmic perspective, this edge independence is an example of the optimal substructure property. A problem has an optimal substructure only if the optimal solution can be obtained from a combination of optimal solutions to its subproblems (Cormen et al., 2001). To solve a problem using dynamic programming (DP) or its successive approximation schemes, such as Dijkstra's algorithm, the optimal substructure assumption has to hold (Snedovich, 2006), that is, the cost of any subpath has to be independent of what precedes it and what succeeds it. As mentioned, the direct transformation of sampling-based methods to belief space breaks this assumption, while FIRM preserves it. Furthermore, edge independence allows the challenging task of computing collision probabilities to be done offline, for each edge separately, without performing costly computations repeatedly and without any simplifying assumption.

The current paper draws on earlier work published in conference papers (Agha-mohammadi et al., 2011, 2012b, 2013b). Compared with Agha-mohammadi et al. (2011), in this paper, we construct the FIRM framework more rigorously by detailing the procedure of transforming the POMDP problem to the belief semi-Markov decision process (SMDP) problem, and then to the FIRM Markov

decision process (FIRM MDP) problem, where the policy on the graph and overall hybrid policy generated by FIRM are distinguished clearly. Also, in this paper we provide a clearer distinction between the abstract FIRM framework and its instantiations, and we provide more rigorous explanation and proofs on SLQG-FIRM. Further, we append the proofs of the probabilistic completeness of FIRM to this paper, which completes the work in Agha-mohammadi et al. (2012b). We also present new unpublished results on the performance of SLQG-FIRM in more difficult environments, and demonstrate its real-time planning capabilities. Further, we provide a complexity analysis of the method and compare it to state-of-the-art methods.

The outline is as follows. In the next section, we review the most relevant related work. Section 3 provides an overview of the method and its contributions. In Section 4 we describe the general problem of feedback motion-planning under uncertainty, present notation, and formulate the POMDP problem. In Section 5, we present the SLQG-based instantiation of the abstract FIRM framework by providing concrete belief samplers and connectors (local planners). In Section 6, assuming the existence of belief samplers and connectors, we introduce the abstract FIRM framework and detail the process of transforming POMDP to a FIRM MDP. In Section 7, aiming at evaluating the quality of the FIRM solution, we extend the concepts of success and probabilistic completeness to the stochastic setting and prove the probabilistic completeness of the FIRM framework. Experimental results are presented in Section 8. In Section 9, we discuss limitations of the framework, future work, and open issues. Section 10 concludes the paper.

2. Related work

In this section we review the related work and place our work into context. First, we review the related work on planning algorithms under uncertainty, and then we consider the work concerning probabilistic completeness.

2.1. Planning algorithms

Uncertainty in robotic systems usually stems from three sources: (i) motion uncertainty, which results from the noise that affects system dynamics; (ii) sensing uncertainty, caused by noisy sensory measurements, which is also referred to as imperfect state information; and (iii) uncertainty in the environment map, such as uncertain obstacle locations or uncertain locations of features (information sources) in the environment.

Methods such as those in Missiuro and Roy (2006), Guibas et al. (2008), and Nakhaei and Lamiroux (2008) deal with map uncertainty. However, we do not scrutinize these methods, since we assume there is no uncertainty in the environment map. Methods such as those in Alterovitz et al. (2007), Melchior and Simmons (2007),

and Chakravorty and Kumar (2009, 2011) exploit sampling-based motion-planning ideas to deal with motion uncertainty. However, methods that are most related to FIRM consider both motion and sensing uncertainties in planning, where the ultimate goal is to solve a POMDP problem, in other words, to find the best policy that generates optimal actions as a function of belief. However, due to the intractability of the POMDP solution, the practical results using these methods are usually limited to problems with small sets of discrete states (Kaelbling et al., 1998). Point-based POMDP solvers such as those in Porta et al. (2006), Kurniawati et al. (2008), Bai et al. (2010), and Ong et al. (2010) have increased the size of problems that can be solved by POMDPs. However, they do not handle continuous state, control, and observation spaces. For the Gaussian belief case, Van den Berg et al. (2011, 2012) handle continuous spaces locally around a given trajectory in belief space. Platt et al. (2011) generalize the local approaches to non-Gaussian beliefs.

In continuous state, control, and observation space, the main body of methods does not follow the POMDP framework due to its extreme complexity. Instead, these methods return a nominal path as the solution of the planning problem, which is fixed regardless of the process and sensor noise in the execution phase. Censi et al. (2008) propose a planning algorithm based on graph search and constraint propagation on a grid-based representation of the space. Platt et al. (2010) plan in continuous space by finding the best nominal path using nonlinear optimization methods. In the linear quadratic Gaussian motion planning (LQG-MP) method (Van den Berg et al., 2010), among the finite number of RRT paths, the best path is found by simulating the performance of LQG on all RRT paths. Bry and Roy (2011) propose a tree-based approach, in which the underlying nominal trajectory is optimized using RRT*. Vitus and Tomlin (2011) also propose an approach to optimize the underlying trajectory by formulating the problem as a chance-constrained optimal control problem. In Van den Berg et al. (2011), the authors also extend the LQG-MP to roadmaps. Prentice and Roy (2009) and Huynh and Roy (2009) also utilize roadmap-based methods based on the PRM approach, where the best path is found through a breadth-first search on the belief roadmap (BRM). However, in all these roadmap-based methods, the optimal substructure assumption is violated, in other words, the costs of different edges on the graph depend on each other. The point-based POMDP planner in Kurniawati et al. (2012) takes into account motion, observation, and map uncertainties, and advances the previous point-based methods by introducing guided cluster sampling. It starts with a roadmap in the configuration space, and grows a single-query tree in the belief space, rooted in the initial belief. Since these methods return a nominal path instead of a feedback policy, the path needs to be recomputed (in other words, replanning has to be performed) in the case of large

deviations or when starting from a new point. However, unless the planning domain is small (e.g. in Platt et al., 2010), replanning using these methods is computationally very expensive. The reason for this is that the constructed planning tree depends on the starting belief, and therefore all computations needed to construct the tree (including predicting future costs) have to be reproduced from the new starting belief. BRM ameliorates this expensive computation using covariance factorization techniques, but it still does not satisfy the optimal substructure assumption. Thus, for a new query from a new initial point, BRM needs to perform the search algorithm again. In the presence of obstacles, recomputing the collision probabilities is also needed, which makes replanning even more expensive. In other words, these methods are single-query, in the sense that the edge costs are computed for a given query.

Since these methods are single-query, online replanning can be done only if the planning domain is small (e.g. in Platt et al., 2010) or if the planning horizon is short, such as receding-horizon-control-based (RHC-based) approaches (e.g. Chakravorty and Erwin, 2011). The method proposed in Toit and Burdick (2010) is an RHC-based method, where the nominal path is updated dynamically over an N -step horizon. The PUMA framework proposed in He et al. (2011) is also an RHC-based framework, where instead of a single action, a sequence of actions (macro-action) is selected at every decision stage. However, these methods entail repeatedly solving open loop optimal control problems at every time step, which is computationally very expensive as the previous computations cannot be reused for the queries from the new initial point. In FIRM, however, a feedback policy (that is, a mapping from belief space to actions) is computed offline. Thus in replanning from a new initial point, the computations need not be reproduced. Thus for a fixed goal, the algorithm is robust to changes in the start point of the query. It is also robust to changes in the goal point, because graph feedback can be computed (see Equation (32)) online, which results in a multi-query roadmap in the belief space.

In the methods that account for sensing uncertainty, the state has to be estimated based on measurements. To handle unknown future measurements in the planning stage, the methods in Censi et al. (2008), Huynh and Roy (2009), Prentice and Roy (2009), Platt et al. (2010), and Toit and Burdick (2010) consider the maximum likelihood (ML) observation sequence to predict the estimation performance. In contrast, FIRM takes all possible future observations into account in the planning stage. The methods in Van den Berg et al. (2010, 2011) also consider all possible future observations.

In the presence of obstacles, due to the dependency of collision events on each other in different time steps, it is a burdensome task to include the collision probabilities in planning. Thus, the methods in Censi et al. (2008), Van den Berg et al. (2010, 2011), and Toit and Burdick (2010)

design some safety measures to account for obstacles in planning. A problem with some of these collision probability measures is that they are built on the assumption that the collision probabilities at different stages along the path are independent of each other, which is not true in general and may lead to very conservative plans (see Figure 2). As a result, different methods (e.g. Patil et al., 2012) aim at providing more accurate and faster ways of computing collision probabilities. In FIRM, however, collision probabilities can be computed and seamlessly incorporated into the planning stage without making any simplifying assumptions.

2.2. Probabilistic completeness

Due to the success of sampling-based methods in many practical planning problems, researchers have investigated the theoretical basis for this success. However, almost all of these investigations have been performed for algorithms that are designed for planning in the absence of uncertainty. The literature in this direction falls into two categories: path-isolation-based methods and space-covering-based methods.

Path-isolation-based analysis: In this approach, one path is chosen, and it is tiled with some sets such as ϵ -balls (Kavraki et al., 1998) or sets with arbitrary shapes but strictly positive measures (Ladd and Kavraki, 2004). Then the success probability is analyzed by investigating the probability of sampling in each of the sets that tile the given path in the obstacle-free space. The methods in Švestka and Overmars (1997), Kavraki et al. (1998), Bohlin (2002), and Ladd and Kavraki (2004) are among those that perform path-isolation-based analysis of the planning algorithm.

Space-covering-based analysis: In space-covering-based analysis, an adequate number of sampled points needed to find a successful path is expressed in terms of a parameter ϵ , which is a property of the environment. A space is ϵ -good if every point in the state space can be connected to at least an ϵ fraction of the space using local planners. The methods in Kavraki et al. (1995) and Hsu (2000) are among these.

These methods were developed for the situation where the desired result from the planning algorithm is a path. However, in the presence of uncertainty, the concept of ‘successful path’ is no longer meaningful, because on a given path, different policies may result in different success probabilities, where some are interpreted as successful and some are not. Thus, since the planning algorithm returns a policy instead of a path, success has to be defined for a policy. This paper extends these concepts to probabilistic spaces, that is, to sampling-based methods concerning planning under uncertainty. In Section 7, we define and formulate the concepts of successful policy and probabilistic completeness under uncertainty (PCUU).

3. Method overview and contributions

The highlights and contributions of this paper can be divided into theoretical and practical parts. The theoretical highlights can be summarized as follows:

- *Abstract frameworks:* We introduce the abstract information roadmap (IRM) framework as a graph in the belief space, where the graph nodes are beliefs (rigorously speaking, small subsets of the belief space) and edges are local controllers. The abstract FIRM framework is defined as an IRM where local controllers are feedback controllers. These abstract frameworks rely on the existence of an appropriate belief node sampler and connector (local controller) and are general enough to capture any form of belief. Discussing the concept of belief reachability under feedback controllers, we detail the reduction of a POMDP to a tractable MDP on the FIRM graph, referred to as the FIRM MDP.
- *SLQG-FIRM:* To instantiate a FIRM, we need concrete belief samplers and connectors. A concrete example of these components based on SLQG controllers is given in Section 5. Basically, it is shown that under an SLQG controller the belief can be driven into the ϵ -neighborhood of the sampled Gaussian beliefs in finite time, and thus node reachability is achieved. In this fashion, SLQG-FIRM addresses the hard task of sampling in reachable belief space that is required in belief-space planning (Spaan and Vlassis, 2005; Pineau et al., 2006; Kurniawati et al., 2010). The focus of SLQG-FIRM is on kinematic systems. However, we also extend it to dynamical systems by restricting the nodes to the equilibrium space.
- *Graph (multi-query roadmap) in belief space:* FIRM is the first framework that generates a *graph* in the belief space with independent edges. In other words, it is a multi-query roadmap, which distinguishes it from other methods in the belief space.
- *Breaking the curse of history:* A fundamental contribution of FIRM is that the optimal action at a given node does not depend on the traversed nodes, actions, and observations prior to this node, in other words, it is independent of the history of the information process (see Figure 1). This is a direct consequence of inducing reachable belief nodes using feedback controllers, which breaks the curse of history in POMDPs. In addition, the sampling-based nature of the method, borrowed from PRM, allows us to ameliorate the curse of dimensionality.
- *Probabilistic completeness:* Finally, we generalize the conventional concept of ‘probabilistic completeness’ (which is defined for motion-planning methods in deterministic environments) to the concept of PCUU (which is defined for the planners in the presence of uncertainty). According to this definition, we prove that FIRM is a probabilistically complete algorithm. Moreover, we perform an analysis on the absorption

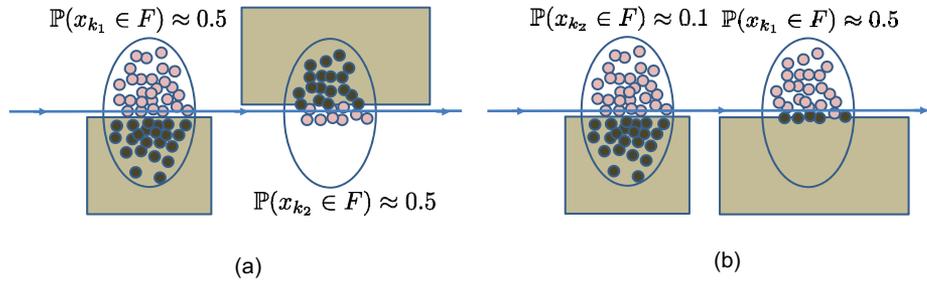


Fig. 2. The dependence of collision events on each other at different time steps. x_k is the robot state and F is the obstacle set (rectangles). $\mathbb{P}(x_k \in F)$ is the collision probability at the k -th time step. Drawn ellipses are 3σ ellipses of Gaussian distributions obtained by Kalman filtering. Also, the samples in Monte Carlo simulation are shown by small circles. The dark ones have collided with obstacles and do not get propagated, and the light ones are the safe samples. Although the overall collision probability in (a) is much more than the collision probability in (b), simplified safety measures based on the ellipse–obstacle intersection area lead to the same safety measure in (a) and (b), and are unable to capture this dependency.

probability of the local planners in the belief space, which provides useful general tools that can be used in analyzing planning methods under uncertainty.

More importantly, FIRM offers a set of practical contributions, which we believe provides an important step toward utilizing POMDPs as a practical tool for robot motion planning under uncertainty. The main practical highlights can be summarized as follows:

- *Efficient planning:* The construction of FIRM is offline and thus online planning (and replanning) is feasible and almost instantaneous.
- *Robustness:* The optimal feedback policy, instead of a nominal path, is computed offline. It is obtained by solving the DP problem associated with the FIRM MDP on the belief graph. As a result, no replanning is needed even in the case of large deviations (or just local real-time replanning is sufficient), and the feedback over the belief space can take care of deviations. Therefore, the method is robust to large deviations. It is also less sensitive to linearization errors, since if the system goes out of the linearization region of a controller, it falls into the valid linearization region of some other controller (assuming a sufficient number of FIRM nodes) that can take the belief and drive it to the goal.
- *Reliability (incorporating obstacles in planning):* In the FIRM framework, collision probabilities can be computed, which leads to more accurate plans, as opposed to simplified collision measures that may lead to conservative plans (see Figure 2). The obstacles add a *failure node* to the FIRM graph into which the robot can be absorbed. Further, due to the offline construction of FIRM, the heavy computational burden of estimating collision probabilities can be done offline.
- *Scalability:* Belief-space planners usually have an exponential planning complexity either in the number of nodes (if they are sampling-based methods) or in the

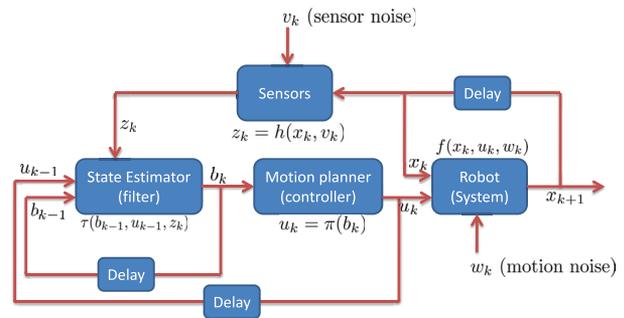


Fig. 3. Block diagram corresponding to the problem of planning under motion and sensing uncertainty.

size of grid (if they rely on discretizing the environment). However, the complexity of the FIRM construction is a constant multiplier of the complexity of the PRM construction. Moreover, the complexity of planning (or replanning) with FIRM is a constant, which is independent of the size of the underlying graph.

4. Problem formulation

The main sources of uncertainty in motion planning are the lack of exact knowledge of the robot's motion model, the robot's sensing model, and the environment model, which are referred to as motion uncertainty, sensing uncertainty, and map uncertainty, respectively. In this paper, we focus on motion and sensing uncertainty, but some of the concepts are extensible to problems with map uncertainty. The MDP problem and the POMDP problem are the most general formulations, respectively, for planning problems under motion uncertainty and for planning problems under both motion and sensing uncertainty.

While in the deterministic setting, we seek an optimal path as the solution of the motion-planning problem, in the stochastic setting, we seek an optimal feedback (mapping) π as the solution of the motion-planning problem. In the case of an MDP, π is a mapping from the state space to the

control space, while in the case of POMDP, π is a mapping from the belief space to the control space (see Figure 3). In the rest of paper, we focus on POMDPs.

4.1. Preliminaries and notation

As mentioned, the POMDP formulation is the most general formulation for the planning problem under process (motion) uncertainty and imperfect state information (sensing uncertainty). POMDPs are introduced in Astrom (1965), Smallwood and Sondik (1973), and Kaelbling et al. (1998). In the following, we first explain different elements of the POMDP problem, and then present a form of the POMDP formulation which is known as the belief MDP problem (Thrun et al., 2005; Bertsekas, 2007).

State, control, and observation: Let $x_k \in \mathbb{X}$, $u_k \in \mathbb{U}$, and $z_k \in \mathbb{Z}$ denote the system state, control, and observation at time step k , respectively, where $\mathbb{X} \subseteq \mathbb{R}^{d_x}$, $\mathbb{U} \subseteq \mathbb{R}^{d_u}$, and $\mathbb{Z} \subseteq \mathbb{R}^{d_z}$ are the state, control, and observation spaces. Scalars d_x , d_u , and d_z are the state, control, and observation dimensions and \mathbb{R}^d denotes the d -dimensional Euclidean space.

Basically, system state x encodes all information needed for decision-making at a specific time instant. It is worth noting that the state space in our problem is continuous. Control space \mathbb{U} , which contains all possible control inputs (or actions), can also be continuous, and $u_{0:k} := \{u_0, u_1, \dots, u_k\}$ denotes the control history up to step k . Similarly, the observation space \mathbb{Z} that contains all possible observations (sensor measurements) can also be continuous, and $z_{0:k} := \{z_0, z_1, \dots, z_k\}$ is the observation history up to step k .

State evolution model: The process model (or the motion model) $x_{k+1} = f(x_k, u_k, w_k)$ describes how the system state evolves as a function of the applied control u_k and the process (motion) noise w_k , which is distributed according to the (known) probability density function (pdf) $p(w_k|x_k, u_k)$. An equivalent representation of this evolution model is through the transition pdf $p(x'|x, u) : \mathbb{X} \times \mathbb{U} \times \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$, which encodes the probability density of the transition from state x to state x' under the control u .

Observation (sensor) model: Although x_k is sufficient information to make the decision (generate control u_k), in partially observable systems, the system state is *unknown* and the only available data for decision-making is the imperfect measurements of the state made by the sensors. The observation model $z_k = h(x_k, v_k)$ encodes the relation between system state x_k and its measurements z_k , where v_k is the observation noise at time step k , which is distributed according to the (known) pdf $p(v_k|x_k)$. An equivalent representation of this observation model is through the likelihood pdf $p(z|x) : \mathbb{X} \times \mathbb{Z} \rightarrow \mathbb{R}_{\geq 0}$.

Information state (belief): In partially observable environments, the available data for decision-making in time step k is the history of observations we have made, $z_{0:k}$, and the history of actions we have taken, $u_{0:k-1}$. Let us denote

this data history by $\mathcal{H}_k = \{z_{0:k}, u_{0:k-1}\}$. This data history can be compressed to a conditional probability distribution over all possible states, that is, $b_k = p(x_k|z_{0:k}; u_{0:k-1})$. The pdf $b_k : \mathbb{X} \times \mathbb{Z}^k \times \mathbb{U}^{k-1} \rightarrow \mathbb{R}_{\geq 0}$ is called the information state or belief at the k th step. \mathbb{B} denotes the belief space of the problem, containing all possible beliefs $b \in \mathbb{B}$.

Belief evolution model (filter model): In recursive state estimation techniques, belief can be computed recursively. The belief evolution model (or belief dynamics) introduced by this recursion is shown by function $\tau : \mathbb{B} \times \mathbb{U} \times \mathbb{Z} \rightarrow \mathbb{B}$, which computes the next belief based on the last action and current observation $b_{k+1} = \tau(b_k, u_k, z_{k+1})$. This belief evolution model can be derived using Bayes' rule and the law of total probability (Thrun et al., 2005; Bertsekas, 2007) as follows:

$$b_{k+1} = p(z_{k+1}|\mathcal{H}_k, u_k)^{-1} p(z_{k+1}|x_{k+1}) \int_{\mathbb{X}} p(x_{k+1}|x_k, u_k) b_k dx_k =: \tau(b_k, u_k, z_{k+1}) \quad (1)$$

An equivalent representation of the belief evolution model is through the transition pdf $p(b'|b, u) : \mathbb{B} \times \mathbb{U} \times \mathbb{B} \rightarrow \mathbb{R}_{\geq 0}$ that encodes the probability density of the transition from belief b to belief b' under the control u .

Policy: In a partially observable system, the planner π (also called the policy or feedback controller) has to be a function that returns an action u_k given the available data \mathcal{H}_k . However, it can be shown that the compression of data \mathcal{H}_k to belief b_k preserves all the information needed for decision-making (Kumar and Varaiya, 1986). Therefore, a policy $\pi(\cdot)$ has to be a function that returns an action u_k given the belief b_k , in other words, $\pi(\cdot) : \mathbb{B} \rightarrow \mathbb{U}$:

$$u_k = \pi(b_k), \quad \forall b_k \in \mathbb{B} \quad (2)$$

The space of all possible $\pi(\cdot)$ is denoted by Π .

Cost-to-go: To choose an optimal policy, we need to have a cost function, which is a task-dependent quantity. But let us in general denote the one-step cost of taking action u at belief b by $c(b, u) : \mathbb{B} \times \mathbb{U} \rightarrow \mathbb{R}_{\geq 0}$. Then, we can define the cost-to-go function $J^\pi(\cdot) : \mathbb{B} \rightarrow \mathbb{R}_{\geq 0}$ from a belief b_0 under the policy π as

$$J^\pi(b_0) := \sum_{k=0}^{\infty} \mathbb{E}[c(b_k, \pi(b_k))] \quad (3)$$

s.t. $b_{k+1} = \tau(b_k, \pi(b_k), z_{k+1}), \quad z_k \sim p(z_k|x_k)$

where $\mathbb{E}[\cdot]$ is the expectation operator. Consider a goal region $B^{goal} \subset \mathbb{B}$ such that, for all u , we have $c(b \in B^{goal}, u) = 0$; in other words, the goal region is cost-absorbing. Then, the above cost-to-go would be finite for a policy that can drive the state to the goal region in finite time.

4.2. POMDP problem

Given the motion model f , observation model h , and cost-to-go J^π , the POMDP problem seeks the best policy that

minimizes the cost-to-go function from every belief in the belief space. Formally, if we denote the *optimal cost-to-go* function by $J(\cdot)$, we can define *optimal policy* $\pi(\cdot): \mathbb{B} \rightarrow \mathbb{U}$, which is the solution of POMDP as follows:

$$J(\cdot) := \min_{\Pi} J^{\pi}(\cdot) \quad (4)$$

$$\pi = \arg \min_{\Pi} J^{\pi}(\cdot) \quad (5)$$

This formulation of the POMDP problem is also known as the *belief-MDP* problem (Thrun et al., 2005; Bertsekas, 2007), because it is an MDP over the belief space.

Dynamic programming: It is well known that the optimal cost-to-go is obtained by solving the following stationary DP equation on the belief space \mathbb{B} (Thrun et al., 2005; Bertsekas, 2007). Subsequently, the solution of the POMDP (i.e. π) can be computed as a function that returns the argument of this minimization, that is, returns the optimal action at every belief:

$$J(b) = \min_u \{c(b, u) + \int_{\mathbb{B}} p(b'|b, u) J(b') db'\}, \forall b \in \mathbb{B} \quad (6a)$$

$$\pi(b) = \arg \min_u \{c(b, u) + \int_{\mathbb{B}} p(b'|b, u) J(b') db'\}, \forall b \in \mathbb{B} \quad (6b)$$

However, it is well known that this DP equation is exceedingly difficult to solve since it is defined over the entire belief space and suffers from the curse of history (Pineau et al., 2003) and the curse of dimensionality.

Constrained POMDP problems: The presence of constraints makes this problem even more difficult. We denote the constraint set (or the failure set) in the state and control space by $F \subset \mathbb{X} \times \mathbb{U}$, which needs to be avoided by the system, in other words, $(x_k, u_k) \notin F$, for all k .

4.3. Problem description

We aim at constructing a sampling-based solution to the belief MDP problem. The main goals of this paper are as follows.

SLQG-based FIRM: First, we construct a roadmap in belief space utilizing SLQG controllers as belief stabilizers. We perform this construction for a certain class of systems, and show that the belief reachability condition is guaranteed. In designing SLQG-FIRM, we first focus on kinematic systems (satisfying $x = f(x, 0, 0)$). Then, using the notion of equilibrium space and restricting the sampling to this space, we apply the method to dynamical systems as well.

General FIRM framework: After studying the concrete SLQG-FIRM example, we consider the more general case, where, for a general system, assuming that there exists a controller under which belief reachability is guaranteed, we (i) construct a graph in the belief space encoding the failure probabilities on its edges, (ii) reduce the intractable belief MDP in Equation (4) into a tractable MDP problem on this graph, and (iii) compute a feedback solution on this graph.

5. SLQG-FIRM

In this section, we discuss a particular instance of the FIRM framework in which belief reachability is accomplished by SLQG controllers. In Section 6, we propose the general FIRM framework.

We start this section by restricting our attention to the class of systems that SLQG-FIRM can handle. Then, we present a brief review of LQG controllers and address how we can define nodes in the belief space to satisfy reachability using SLQG controllers. Next we explain the procedure of constructing local controllers (i.e. FIRM edges) and the SLQG-based FIRM graph. Finally, we compute transition probabilities and costs associated with each graph edge and compute the graph feedback.

5.1. Preliminaries on SLQG

In this section, we assume the noise is Gaussian, and we start by defining the notation needed in dealing with Gaussian beliefs.

Gaussian belief space: We denote the random estimation vector by x^+ , whose distribution is $b_k = p(x_k^+ | z_{0:k}, u_{0:k-1})$, and denote the mean and covariance of x^+ by $\hat{x}^+ = \mathbb{E}[x^+]$ and $P = \mathbb{E}[(x^+ - \hat{x}^+)(x^+ - \hat{x}^+)^T]$, respectively. Denoting the Gaussian belief space by $\mathbb{G}\mathbb{B}$, every function $b(\cdot) \in \mathbb{G}\mathbb{B}$ can be characterized by a mean-covariance pair (\hat{x}^+, P) . Abusing the notation, we also show this pair by $b \equiv (\hat{x}^+, P) \in \mathbb{R}^n \times \mathbb{S}_+^n$, where the mean vector belongs to the n -dimensional Euclidean space \mathbb{R}^n and the covariance matrix belongs to the space of all positive semi-definite $n \times n$ matrices \mathbb{S}_+^n .

LQG controllers: An LQG controller is composed of a Kalman filter (KF) as the state estimator and a linear quadratic regulator (LQR) controller (see Figure 3). Thus, the belief dynamic $b_{k+1} = \tau(b_k, u_k, z_{k+1})$ is known and comes from the Kalman filtering equations, and the controller $u_k = \mu(b_k)$ that acts on the belief comes from the LQR equations. Considering a quadratic cost for state error and control error, LQG is an optimal controller for linear systems with Gaussian noise (Bertsekas, 2007). However, it is also often used for stabilization of nonlinear systems around a given trajectory or around a given point.

Stationary and time-varying LQG: Time-varying LQG is designed to track a given trajectory, in which at every time step a different feedback policy is utilized. SLQG is a time-invariant policy, in which LQG is designed around a given point, say \mathbf{v} , to steer the state of the system to \mathbf{v} (Bertsekas, 2007). In Appendices B and C we review these controllers in detail.

Equilibrium space: Let us denote a configuration of a robotic system (Lozano-Perez, 1983) by q . Kinematic models are specified in terms of the configuration variable q , while dynamical models are specified by the state $x = (q, \dot{q})$, where \dot{q} denotes the corresponding velocities. In SLQG-FIRM, we sample the underlying PRM nodes (stabilizer parameters) from the configuration space. Thus, for

dynamical systems, we impose the condition $\dot{q} = 0$ on the samples, in other words, we sample from the equilibrium space of the system, which is denoted by \mathbf{X} in this paper.

Remark 1. *FIRM can be generalized to cases that do not need to sample in equilibrium space. For example, in systems such as fixed-wing aircraft, the system cannot reach the zero velocity $\dot{q} = 0$. In such cases, SLQG is not a suitable choice and one needs to design more appropriate controllers, such as periodic controllers as detailed in Agha-mohammadi et al. (2012c, 2013a). In such a case, we sample periodic maneuvers as FIRM nodes. In other words, we go from periodic trajectory to periodic trajectory instead of going from point to point (Agha-mohammadi et al., 2012c, 2013a).*

5.2. Belief stabilizers

In SLQG-FIRM nodes, we use SLQG controllers as belief stabilizers, that is, as a tool to reach (stabilize to) a predefined belief (FIRM node). To explain how SLQG works as a belief stabilizer, consider a fixed point $\mathbf{v} \in \mathbf{X}$ in the state space and consider the following linear (linearized) system about \mathbf{v} :

$$x_{k+1} = \mathbf{A}x_k + \mathbf{B}u_k + \mathbf{G}w_k, \quad w_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}) \quad (7a)$$

$$z_k = \mathbf{H}x_k + v_k, \quad v_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}) \quad (7b)$$

SLQG controller: The goal of the SLQG controller designed about \mathbf{v} is to keep the state as close as possible to the desired point \mathbf{v} and also keep the energy consumed at a reasonable level. More rigorously, SLQG minimizes the following quadratic cost:

$$J = \mathbb{E} \left\{ \sum_{k \geq 0} (x_k - \mathbf{v})^T \mathbf{W}_x (x_k - \mathbf{v}) + u_k^T \mathbf{W}_u u_k \right\} \quad (8)$$

where \mathbf{W}_x and \mathbf{W}_u are positive-definite weight matrices that are defined by the user. In Appendix C, the SLQG controller minimizing the above cost is discussed in detail. However, in brief, the belief propagation and control generation is carried out as follows:

$$\begin{aligned} b_{k+1} &\equiv \begin{bmatrix} \hat{x}_{k+1}^+ \\ P_{k+1}^+ \end{bmatrix} = \\ &\begin{bmatrix} \mathbf{A}\hat{x}_k^+ + \mathbf{B}u_k + \mathbf{K}_{k+1}(z_{k+1} - \mathbf{H}(\mathbf{A}\hat{x}_k^+ + \mathbf{B}u_k)) \\ (\mathbf{I} - \mathbf{K}_{k+1}\mathbf{H})(\mathbf{A}P_k^+ \mathbf{A}^T + \mathbf{G}\mathbf{Q}\mathbf{G}^T) \end{bmatrix} \\ &\equiv \tau(b_k, u_k, z_{k+1}) \end{aligned} \quad (9)$$

where \mathbf{K}_k is called the Kalman gain at the k th time step and is computed as follows:

$$\mathbf{K}_{k+1} = (\mathbf{A}P_k^+ \mathbf{A}^T + \mathbf{G}\mathbf{Q}\mathbf{G}^T) \mathbf{H}^T (\mathbf{H}(\mathbf{A}P_k^+ \mathbf{A}^T + \mathbf{G}\mathbf{Q}\mathbf{G}^T) \mathbf{H}^T + \mathbf{M}\mathbf{R}\mathbf{M}^T)^{-1} \quad (10)$$

The control signal is generated using a stationary feedback gain \mathbf{L}_s :

$$u_k = -\mathbf{L}_s(\hat{x}_k^+ - \mathbf{v}) =: \mu(b_k), \quad \mathbf{L}_s = (\mathbf{B}_s^T \mathbf{S}_s \mathbf{B}_s + \mathbf{W}_u)^{-1} \mathbf{B}_s^T \mathbf{S}_s \mathbf{A}_s \quad (11)$$

where \mathbf{S}_s is the solution of the following discrete algebraic Riccati equation (DARE):

$$\mathbf{S}_s = \mathbf{W}_x + \mathbf{A}_s^T \mathbf{S}_s \mathbf{A}_s - \mathbf{A}_s^T \mathbf{S}_s \mathbf{B}_s (\mathbf{B}_s^T \mathbf{S}_s \mathbf{B}_s + \mathbf{W}_u)^{-1} \mathbf{B}_s^T \mathbf{S}_s \mathbf{A}_s \quad (12)$$

Controllable and observable pairs: Consider an $n \times n$ matrix \mathbf{A} . A pair of matrices (\mathbf{A}, \mathbf{B}) is called a controllable pair if the controllability matrix $\mathfrak{C} = [\mathbf{B}, \mathbf{A}\mathbf{B}, \mathbf{A}^2\mathbf{B}, \dots, \mathbf{A}^{n-1}\mathbf{B}]$ has rank n (Bertsekas, 2007). A pair of matrices (\mathbf{A}, \mathbf{H}) is called observable if the pair $(\mathbf{A}^T, \mathbf{H}^T)$ is controllable (Bertsekas, 2007).

Controllable and observable systems: Let us also define the matrices $\check{\mathbf{Q}}$ and $\check{\mathbf{W}}_x$ such that $\mathbf{G}\mathbf{Q}\mathbf{G}^T = \check{\mathbf{Q}}\check{\mathbf{Q}}^T$, $\mathbf{W}_x = \check{\mathbf{W}}_x^T \check{\mathbf{W}}_x$. We next consider a class of linear systems and quadratic cost weights that satisfy the following property.

Property 1. *Pairs (\mathbf{A}, \mathbf{B}) and $(\mathbf{A}, \check{\mathbf{Q}})$ are controllable pairs, and pairs (\mathbf{A}, \mathbf{H}) and $(\mathbf{A}, \check{\mathbf{W}})$ are observable pairs.*

In the following, we present three lemmas, through which we can construct reachable SLQG-FIRM nodes for the systems that satisfy Property 1. However, approaches such as periodic LQG (PLQG)-based FIRM (Agha-mohammadi et al., 2012c) or dynamic feedback linearization (DFL)-based FIRM (Agha-mohammadi et al., 2012a) extend this class of systems by excluding the controllability part in Property 1, and thus consider a broader class of systems.

Lemma 1. *Consider the SLQG controller designed to drive the state of the system in Equation (7) to a point $\mathbf{v} \in \mathbf{X}$. Given that Property 1 is satisfied, in the absence of a stopping region, the belief b_k under SLQG controller converges to a unique stationary belief b_s , in distribution (i.d.). In other words, the distribution over belief converges to a unique distribution. That is,*

$$b_k \xrightarrow{i.d.} b_s \sim \mathcal{N}(b_s, \mathbf{C}) \quad (13)$$

Note that b_k is a random belief that converges to another random belief b_s . In the Gaussian setting, the distribution over the random belief b_s is $\mathcal{N}(b_s, \mathbf{C})$, where $b_s = \mathbb{E}[b_s] \equiv (\mathbf{v}, P_s)$. The stationary estimation covariance matrix P_s is characterized in Lemma 2, and the covariance \mathbf{C} is characterized in Appendix C.

Proof. In Appendix C, we review the SLQG and prove Lemma 1. \square

Lemma 2. *Given Property 1, the following DARE has a unique symmetric positive-definite solution (Bertsekas, 2007), denoted by P_s^- :*

$$P_s^- = \mathbf{G}\mathbf{Q}\mathbf{G}^T + \mathbf{A}(P_s^- - P_s^- \mathbf{H}^T (\mathbf{H}P_s^- \mathbf{H}^T + \mathbf{R})^{-1} \mathbf{H}P_s^-) \mathbf{A}^T \quad (14)$$

Moreover, the stationary covariance matrix P_s introduced in Lemma 1 is computed as:

$$P_s = P_s^- - P_s^- \mathbf{H}^T (\mathbf{H}P_s^- \mathbf{H}^T + \mathbf{R})^{-1} \mathbf{H}P_s^- \quad (15)$$

Proof. See Appendix C or Bertsekas (2007). \square

Now we state the main result, through which we can construct *reachable* FIRM nodes under SLQG-based belief stabilizers:

Lemma 3. *Consider the SLQG controller designed to drive the state of the system in Equation (7) to a point $\mathbf{v} \in \mathbf{X}$. Suppose matrix \mathbf{H} is full rank and Property 1 is satisfied. Then, any set $B \subset \mathbb{B}$ whose interior contains $b_c \equiv (\mathbf{v}, P_s)$ is reachable under the designed SLQG controller starting from any Gaussian distribution. Moreover, the estimation covariance P_k converges to the unique deterministic stationary covariance P_s .*

Proof. See Appendix D. \square

Therefore, based on Lemma 3, SLQG can accomplish the belief reachability for appropriately chosen region B . In the next subsection we explicitly characterize region B .

5.3. Designing SLQG-FIRM nodes

Underlying PRM: As mentioned, to construct a FIRM we first construct an underlying PRM (Kavraki et al., 1996). In the SLQG-FIRM, nodes of the underlying PRM, denoted by $\{\mathbf{v}^j\}_{j=1}^N$, are sampled from the obstacle-free space. Considering linear systems or nonlinear systems that are locally well approximated by linearization, we linearize the system about every PRM node. Let us denote the linear (linearized) system about \mathbf{v}^j as follows:

$$x_{k+1} = \mathbf{A}^j x_k + \mathbf{B}^j u_k + \mathbf{G}^j w_k, \quad w_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}^j) \quad (16a)$$

$$z_k = \mathbf{H}^j x_k + v_k, \quad v_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}^j) \quad (16b)$$

where w_k and v_k are motion and measurement noise, respectively, drawn from zero-mean Gaussian distributions with covariances \mathbf{Q}^j and \mathbf{R}^j .

FIRM nodes: To design the j th FIRM node B^j , we first design the SLQG controller μ_s^j (see Equations (9) and (11)) corresponding to the system in Equation (16). The controller μ_s^j is called the j th node controller or the j th belief stabilizer. Given Property 1, based on Lemma 1, the limiting random belief $b_s^j \equiv (\widehat{x}_s^j, P_s^j)$ exists, and \widehat{x}_s^j and P_s^j are the stationary estimation mean and covariance, respectively. Note that under SLQG, \widehat{x}_s^j is a random variable and P_s^j is a deterministic matrix. Moreover, in Lemma 1, it is shown that $b_c^j = \mathbb{E}[b_s^j] \equiv (\mathbf{v}^j, P_s^j)$, where P_s^j is shown to be unique and computed in Lemma 2.

Thus, we can characterize the j th node center:

$$b_c^j \equiv (\mathbf{v}^j, P_s^j) \quad (17)$$

As a result, considering B^j as a ball with an arbitrary radius $\epsilon > 0$ centered at b_c^j , the pair (B^j, μ_s^j) is a *proper pair*, based on Lemma 3; in other words, B^j is reachable under μ_s^j . Thus, one can define the j th FIRM node as $B^j = \{b : \|b - b_c^j\|_b < \delta\}$, where $\|\cdot\|_b$ denotes a suitable norm in belief space and

δ defines the FIRM node size. A typical example of such a FIRM node in Gaussian belief space can be defined by considering mean and covariance separately:

$$B^j = \{b \equiv (x, P) : \|x - \mathbf{v}^j\| < \delta_1, \|P - P_s^j\|_m < \delta_2\} \quad (18)$$

where δ_1 and δ_2 are suitably small thresholds that determine the size of FIRM node B^j . Moreover, $\|\cdot\|$ is a suitable vector norm and $\|\cdot\|_m$ is a suitable matrix norm. We denote the set of all SLQG-FIRM nodes as $\mathbb{V} = \{B^j\}$.

5.4. Designing SLQG-FIRM edges

A FIRM edge is actually a local planner (local feedback controller). In SLQG-based FIRM, the local controller representing the (i, j) th edge is denoted by μ^{ij} . The role of μ^{ij} is to drive the belief from the node B^i to the node B^j . Based on Lemma 3, for a linear system, if we choose $\mu^{ij} = \mu_s^j$, as was done in Agha-mohammadi et al. (2011), the node B^j is reachable under μ^{ij} . However, to better cope with nonlinearities, we construct the local controller μ^{ij} by preceding the node controller with a time-varying LQG controller $\bar{\mu}_k^{ij}$, which is called an *edge controller* here. Time-varying LQG controllers are described in detail in Appendix B.

PRM edge: To design edge controllers, first the underlying PRM edges, denoted by $\mathcal{E} = \{\mathbf{e}^{ij}\}$, have to be constructed. For kinematics-based models there are many different methods in the PRM literature to construct such edges. For dynamical models, there are fewer choices. A few examples are in Van den Berg and Overmars (2007) and Agha-mohammadi et al. (2012c).

Edge controllers: An edge controller $\bar{\mu}_k^{ij}$ in SLQG-FIRM is built by linearizing the system along the (i, j) th PRM edge \mathbf{e}^{ij} and designing a time-varying LQG controller to track it (see Appendix B). The edge controller has two major roles. First, it tries to track the PRM edge and thus exploits the available information on the PRM edges, such as some clearance from the obstacles. Second, in the case where the neighboring PRM nodes are not close to each other, it takes the belief into the valid linearization region of the j th belief stabilizer, where it hands the system over to the belief stabilizer, and the belief stabilizer in turn takes the system to the j th FIRM node.

Local controllers: Thus, overall, the (i, j) th local controller μ^{ij} is the concatenation of the (i, j) th edge controller $\bar{\mu}_k^{ij}$ and j th node controller (belief stabilizer) μ_s^j . We denote the set of all SLQG-FIRM edges by $\mathbb{M} = \{\mu^{ij}\}$ and the set of all SLQG-FIRM edges originating from B^i by $\mathbb{M}(i)$.

SLQG-FIRM: Formally, we define SLQG-FIRM as a graph with the set of nodes $\mathbb{V} = \{B^j\}$ and the set of edges (or local controllers) $\mathbb{M} = \{\mu^{ij}\}$. The set of controllers originating from B^i is denoted by $\mathbb{M}(i) \subset \mathbb{M}$.

5.5. Transition probabilities and edge costs

To find a feedback on a FIRM graph, we need to compute the cost associated with the graph edges. Moreover, we

include the constraint set F into the planning with FIRM by computing the probability of violating the constraint $(x, u) \notin F$ along the graph edges. Let us denote the cost of taking controller μ^{ij} at node B^i by $C^g(B^i, \mu^{ij})$. Superscript g refers to the ‘global’ (or ‘graph-level’) quantities, as these quantities are used to find the global policy (or policy on the graph). Similarly, let $\mathbb{P}^g(B^j|B^i, \mu^{ij})$ and $\mathbb{P}^g(F|B^i, \mu^{ij})$ denote the probabilities of the transition to B^j and F under μ^{ij} , respectively. These quantities are rigorously defined in Section 6 and their connection with the original POMDP is established. However, in this subsection, we just give an example of how such costs and transition probabilities can be computed.

Transition probabilities: Computing transition probabilities $\mathbb{P}^g(\cdot|B^i, \mu^{ij})$ in general can be computationally expensive. Here, we utilize particle-based methods to approximate the distributions and thus compute the collision probabilities. Basically, we can approximate the failure and reachability probabilities based on the number of particles that violate the constraints (hit the set F) and based on the number of particles that can reach the target node (hit the set B^j). The method is described in more detail with the experiments in Section 8.1.4. The dependency of collision events on each other in different time steps, which is ignored in most collision probability computation methods in the POMDP literature, can be taken into account rigorously in particle-based methods. Owing to the offline construction of FIRM, the high computational burden of particle-based approaches can be tolerated. However, any other method for computing transition probabilities can also be adopted, such as that in Patil et al. (2012).

Edge costs: The FIRM edge costs in general and their derivation based on the one-step costs of the original POMDP problem are defined in Section 6. However, roughly speaking, we can define the cost $C^g(B^i, \mu^{ij})$ as the sum of all one-step costs along the edge until the system reaches the target node B^j or hits the failure set F . Depending on the application, one can define a variety of cost functions. Here, we form a cost function based on a linear combination of the estimation accuracy and edge traverse time. This cost function aims to find paths for which the estimator (and hence the controller) can perform well, and also to find faster paths. An indicator of estimation error is the trace of estimation covariance. Thus, we define $\Phi^{ij} = \mathbb{E}[\sum_{k=1}^T \text{tr}(P_k^{ij})]$ along the edge. In SLQG, the covariance matrix evolves deterministically and thus the expectation operator can be omitted. However, if the filter of choice in the edge controller is the extended Kalman filter (EKF), the covariance matrix evolution is also stochastic, and this measure can take into account its stochasticity. Let us denote the mean stopping time under controller μ^{ij} as \hat{T}^{ij} . Then, the total edge cost is considered as a linear combination of estimation accuracy and expected stopping time, with suitable coefficients α_1 and α_2 :

$$C^g(B^i, \mu^{ij}) = \alpha_1 \Phi^{ij} + \alpha_2 \hat{T}^{ij} \quad (19)$$

5.6. Graph feedback on SLQG-FIRM

Graph policy: Graph policy $\pi^g : \mathbb{V} \rightarrow \mathbb{M}$ is a function that returns an edge (local controller) for any given node of the graph. We denote the space of all graph policies by Π^g . To choose the best graph policy in Π^g we define the optimal graph cost-to-go J^g from every graph node.

Graph cost-to-go: The cost-to-go from a given node B^i is equal to the cost of the next taken controller, that is, $C^g(B^i, \pi^g(B^i))$, plus the expected cost-to-go from the next node or from the failure set. In other words, the DP equations for this graph are

$$J^g(B^i) = \min_{\mathbb{M}(i)} C^g(B^i, \mu^{ij}) + J^g(F) \mathbb{P}^g(F|B^i, \mu^{ij}) \\ + J^g(B^j) \mathbb{P}^g(B^j|B^i, \mu^{ij}) \quad (20a)$$

$$\pi^g(B^i) = \arg \min_{\mathbb{M}(i)} C^g(B^i, \mu^{ij}) + J^g(F) \mathbb{P}^g(F|B^i, \mu^{ij}) \\ + J^g(B^j) \mathbb{P}^g(B^j|B^i, \mu^{ij}) \quad (20b)$$

in which $J(F)$ is a suitably high user-defined cost-to-go for hitting the obstacles. The cost-to-go from goal node B^{goal} is defined to be zero, in other words, $J^g(B^{goal}) = 0$.

Solving SLQG-FIRM DP: The DP in equation (20) is a tractable DP as it is defined on a finite number of graph nodes. Computing the transition costs and probabilities offline, this DP can be solved online using standard techniques, such as value/policy iteration methods, for any submitted query. As a result, FIRM is indeed a multi-query roadmap in belief space. Moreover, if the goal node is fixed and only the starting point of the query changes, then this DP can be solved offline and π^g can be stored as a look-up table.

Offline construction of SLQG-FIRM: Algorithm 1 details the construction of SLQG-FIRM with a given goal node.

5.7. Planning with SLQG-FIRM (query phase)

Given that the FIRM graph is computed offline, the online phase of planning (and replanning) on the roadmap becomes very efficient, and thus feasible in real time. In this section, we assume that the goal node is fixed and we just input the start point as the query. However, as discussed in the previous subsection, one can easily submit queries with different goal locations by solving the DP online. If the initial belief b_0 of the submitted query does not belong to any B^i , we create a singleton set $B_0 = \{b_0\}$ as the initial FIRM node. To connect B_0 to the FIRM graph, we go back into the state space, where the underlying PRM is constructed. There, we add a new PRM node to the graph \mathbf{v}_0 , which is the expected value of the robot state, in other words, $\mathbf{v}_0 = \mathbb{E}[x_0]$. Then, we connect \mathbf{v}_0 to the underlying PRM graph based on the connecting function of the adopted PRM. We denote the set of newly added edges originating from \mathbf{v}_0 by $\mathcal{E}(0)$. Then, corresponding to each edge in $\mathcal{E}(0)$, we design a local controller and call the set of them $\mathbb{M}(0)$.

Algorithm 1: Offline construction of SLQG-FIRM

```

1 input : Free space map,  $X_{free}$ 
2 output : FIRM graph  $\mathcal{G}$ 
3 Sample PRM nodes  $\mathcal{V} = \{\mathbf{v}^i\}_{i=1}^N$  and construct its edges
    $\mathcal{E} = \{\mathbf{e}^{ij}\}$ ;
4 forall the PRM nodes  $\mathbf{v}^i \in \mathcal{V}$  do
5   Design the node controller (SLQG)  $\mu_s^i$  about the
   node  $\mathbf{v}^i$  using Equation (11);
6   Compute associated  $b_c^i$  using Equation (17);
7   Construct FIRM node  $B^i$  using Equation (18);
8 Construct  $\mathbb{V} = \{B^i\}$ ;
9 forall the PRM edges  $\mathbf{e}^{ij} \in \mathcal{E}$  do
10  Design the edge controller (time-varying LQG)  $\bar{\mu}_k^{ij}$ 
   along the edge  $\mathbf{e}^{ij}$  (detailed in Appendix B);
11  Construct the local controller  $\mu^{ij}$  by concatenating
   edge controller  $\bar{\mu}_k^{ij}$  and node controller  $\mu_s^i$ ;
12  Set  $b_0 = b_c^i$ ;
13  Generate sample belief paths  $b_{0:\mathcal{T}}$  and ground truth
   paths  $x_{0:\mathcal{T}}$  induced by controller  $\mu^{ij}$  invoked at  $B^i$ ;
14  Compute the transition probabilities  $\mathbb{P}^g(F|B^i, \mu^{ij})$ 
   and  $\mathbb{P}^g(B^j|B^i, \mu^{ij})$  and transition cost  $C^g(B^i, \mu^{ij})$ ;
15 Construct  $\mathbb{M} = \{\mu^{ij}\}$ ;
16 Compute the cost-to-go  $J^g$  and feedback  $\pi^g$  over the
   FIRM nodes by solving the DP in Equation (20);
17  $\mathcal{G} = (\mathbb{V}, \mathbb{M}, J^g, \pi^g)$ ;
18 return  $\mathcal{G}$ ;

```

Finally, we choose the best initial controller among the local controllers in $\mathbb{M}(0)$ using

$$\begin{aligned} \mu_0^*(\cdot) = \arg \min_{\mu \in \mathbb{M}(0)} \{ & C^g(B_0, \mu) \\ & + \mathbb{P}^g(B(\mu)|B_0, \mu) J^g(B(\mu)) + \mathbb{P}^g(F|B_0, \mu) J^g(F) \} \end{aligned} \quad (21)$$

where $B(\mu)$ is the target node of the controller μ . Under the controller μ_0^* , belief evolves and enters one of FIRM nodes, if no collision occurs. From this FIRM node, a combination of the global graph policy π^g and the local edge policies $\{\mu^{ij}\}$ can take the belief to the goal node, as explained below.

Merging global and local feedbacks: After computing a global graph feedback π^g and local edge feedbacks $\{\mu^{ij}\}$, we can construct a full feedback π . Actually, at every time instance, π is equal to one of the local feedbacks, which is chosen by the global feedback in the last visited node. In other words, given the current FIRM node, we use policy π^g defined on FIRM nodes to find μ^* and pick μ^* to move the robot into $B(\mu^*)$. This process is continued until the system reaches the goal region or hits the failure set. Algorithm 2 illustrates this procedure.

Kidnapped robot problem: In robotics, the kidnapped robot problem commonly refers to a situation where an autonomous robot in operation is carried to an arbitrary

Algorithm 2: Online phase algorithm (planning or replanning with SLQG-FIRM)

```

1 input : Initial belief  $b_0$ , FIRM graph  $\mathcal{G}$ 
2 if  $\exists B^i \in \mathbb{V}$  such that  $b_0 \in B^i$  then
3   compute  $\mu^{ij} = \pi^g(B^i)$ ;
4 else
5   Compute  $\mathbf{v}_0 = \mathbb{E}[x_0]$  based on  $b_0$ , and connect  $\mathbf{v}_0$ 
   to the PRM. Let  $\mathcal{E}(0)$  denote the set of outgoing
   edges from  $\mathbf{v}_0$ ;
6   Set  $B_0 = \{b_0\}$ ; design local controllers associated
   with edges in  $\mathcal{E}(0)$ . Call the set of these local
   controllers  $\mathbb{M}(0)$ ;
7   forall the  $\mu \in \mathbb{M}(0)$  do
8     Generate sample belief paths  $b_{0:\mathcal{T}}$  and ground
     truth paths  $x_{0:\mathcal{T}}$  induced by controller  $\mu$ 
     invoked at  $b_0$ ;
9     Compute the transition probabilities
      $\mathbb{P}^g(F|B_0, \mu)$  and  $\mathbb{P}^g(B(\mu)|B_0, \mu)$  and
     transition costs  $C^g(B_0, \mu)$ ;
10    Set  $i = 0$  and choose the best initial local controller
      $\mu^{ij}$  within the set  $\mathbb{M}(0)$  using Equation (21);
11 while  $B^i \neq B^{goal}$  do
12   Set  $B^j$  as the target node of  $\mu^{ij}$ ;
13   while  $b_k \notin B^j$  and ‘no collision’ do
14     Apply the control  $u_k = \mu^{ij}(b_k)$  to the system;
15     Get the measurement  $z_{k+1}$  from sensors;
16     if Collision happens then return Collision;
17     Update belief as  $b_{k+1} = \tau(b_k, \mu^{ij}(b_k), z_{k+1})$ ;
18   Set  $B^i = B^j$ , then compute  $\mu^{ij} = \pi^g(B^i)$ ;

```

location (Choset et al., 2005). Consider a kidnapped robot problem in a known environment. Just after the robot is kidnapped it would be risky to apply any control, because the robot may be close to an obstacle. Thus, in such a scenario, we first initialize the system belief with a Gaussian with large covariance and go into an ‘information gathering’ mode, where we do not apply any control signal and only gather measurements until the covariance shrinks to a reasonable covariance or it remains unchanged for a significant amount of time (i.e. when there is no additional information to reduce the uncertainty). Afterwards, we connect the resulting belief to the FIRM nodes and continue applying the FIRM policy to move the robot toward the goal region. A more efficient approach of handling this problem is detailed in Agha-mohammadi et al. (2013c) using innovation signals.

6. General FIRM framework

The goal of this section is to construct a general FIRM framework, assuming that there exists a mechanism to guarantee belief reachability. As a result, if for a certain class of systems one comes up with a controller that can accomplish

belief reachability, a graph in belief space directly follows according to this general framework.

To construct the general FIRM, we start by defining elements and assumptions needed in the FIRM construction. Accordingly, we transform the original intractable POMDP problem into an SMDP problem in belief space, inspired by sampling-based methods. Then, we construct an arbitrarily good approximation to the solution of this belief SMDP over finite subsets of belief space (FIRM nodes). Doing so, we end up with a tractable MDP, the so-called FIRM MDP. We discuss this derivation first for the obstacle-free case and then we add the obstacles to the planning framework. We characterize the quality of the solution obtained by FIRM via its success probability and provide a generic algorithm for planning with FIRM.

6.1. Feedback controllers and reachability

Belief transition probability: As discussed in Section 4, in partially observable environments the available data for decision-making at time step k can be compressed as the information state or belief b_k . As discussed, using dynamic estimation schemes, belief can be propagated as $b_{k+1} = \tau(b_k, u_k, z_{k+1})$ (see Equation (1)), which can be presented as a one-step transition pdf $p(b_{k+1}|b_k, u_k)$ or a one-step transition probability $\mathbb{P}(B|b_k, u_k) = \int_B p(b_{k+1}|b_k, u_k)$, where $B \subset \mathbb{B}$.

Feedback controllers and induced transition probability: In partially observable environments, at each stage, the decision-making process is performed based on the belief at that stage. Therefore, a controller is a mapping from the belief space to the control space, in other words, $\mu(\cdot) : \mathbb{B} \rightarrow \mathbb{U}$. Accordingly, a controller μ induces a Markov chain with the one-step transition probability $\mathbb{P}_1(B|b, \mu) := \mathbb{P}(B|b, \mu(b))$ over the belief space.

Hitting time: Let $\mathcal{T}(D|b, \mu) \in [0, \infty]$ denote the hitting time on the set $D \subset \mathbb{B}$ under the controller μ starting from belief b . Formally it is defined as

$$\mathcal{T}(D|b, \mu) := \min\{k \geq 0, b_k \in D | b_0 = b, \mu\} \quad (22)$$

Stopping region: We call region $B \subset \mathbb{B}$ a stopping region of the controller μ if we force the controller to stop executing as the belief reaches the region B ; in other words, for all $b \in B$, we impose $\mathbb{P}_1(B|b, \mu) = 1$.

n-step transition probability: We define the n -step transition probability as the probability of landing in the stopping region B in at most n steps:

$$\mathbb{P}_n(B|b, \mu) := \Pr(\mathcal{T}(B|b, \mu) \leq n) \quad (23)$$

Stationary transition probability: Consider the controller μ that starts executing from belief b and stops executing when the belief enters region B . Thus, we can define $\mathbb{P}(B|b, \mu)$ as the transition probability from b to B induced by μ , when the controller stops executing; in other words,

$\mathbb{P}(B|b, \mu)$ would be the probability of landing in the stopping region B in a finite amount of time:

$$\mathbb{P}(B|b, \mu) := \Pr(\mathcal{T}(B|b, \mu) < \infty) \quad (24)$$

Reachability and accessibility: The stopping region B is called reachable under a controller μ starting from b if $\mathbb{P}(B|b, \mu) = 1$. The stopping region B is called accessible under a controller μ from b , if $\mathbb{P}(B|b, \mu) > 0$.

αT -reachability: The stopping region B is called αT -reachable under a controller μ starting from b if $\mathbb{P}_T(B|b, \mu) = \Pr(\mathcal{T}(B|b, \mu) \leq T) > \alpha$, in other words, if the controller can drive the system into B in fewer than T steps with a probability greater than α .

Reachability basin: The reachability basin \check{B} associated with the pair (μ, B) is the set of all beliefs from which B is reachable under μ in the absence of constraints. The reachability (and αT -reachability) basins are thus defined respectively as follows:

$$\check{B} = \{b \in \mathbb{B} : \mathbb{P}(B|b, \mu) = 1\} \quad (25)$$

$$\check{B}(\alpha, T) = \{b \in \mathbb{B} : \mathbb{P}_T(B|b, \mu) > \alpha\} \quad (26)$$

Clearly, $B \subset \check{B}$, and in practical cases, B is much smaller than \check{B} .

6.2. FIRM graph

In this section, we assume that there are no constraints (i.e. $F = \emptyset$), and we reduce planning over the entire belief space to planning over a representative graph constructed within the belief space. Doing so, we can reduce the MDP problem in (4) over the continuous space into a tractable MDP problem defined over the graph nodes.

Stabilizer sampling: The first step in the construction of the proposed framework is to sample a set of stabilizers $\{\mu^j\}$, where each stabilizer $\mu^j(\cdot)$ is a mapping from the belief space to the control space. Typically, every stabilizer is characterized by a d_v -vector of parameters $\mathbf{v} \in \mathbb{R}^{d_v}$; in other words, we can denote the j th stabilizer more rigorously as $\mu^j(\cdot; \mathbf{v}^j) : \mathbb{B} \rightarrow \mathbb{U}$. As a result, we can sample the parameters $\mathcal{V} = \{\mathbf{v}^j\}$ and then construct a stabilizer corresponding to each parameter. One can view the set \mathcal{V} as a set of underlying PRM nodes in the parameter space.

Sampling FIRM nodes: FIRM nodes $\{B^j\}$ are disjoint sets in the belief space, where the j th node has to be chosen such that it is reachable under the j -stabilizer; in other words, $\mathbb{P}(B^j|b, \mu^j) = 1$, with a sufficiently large \check{B} . We discuss the size of \check{B} further below. Note that, for practical purposes, the reachability condition can be replaced by αT -reachability if needed.

Connecting samples: Consider a set of N samples $\{(\mu^i, B^i)\}_{i=1}^N$, where the reachability basin of the i th sample is denoted by \check{B}^i . Now, consider $\{B^i\}_{i=1}^N$ as the nodes of a graph. The node B^i is connected to the node B^j if, starting from any $b \in B^i$, we can reach B^j using μ^j . In other words

B^i is connected to the node B^j if $B^i \subset \check{B}^j$. Again, the reachability condition can be replaced by the αT -reachability condition.

Checking connection condition: For simple systems (linear with Gaussian noise) and some controllers (such as SLQG), the connection condition can be checked analytically. However, in general, checking this connection condition analytically may be very difficult. In such cases, the Markov chain induced by the controller can be simulated numerically (e.g. using particle-based methods). Accordingly, we can approximate the reachability (or αT -reachability) probability and check if the condition is true or not. Since this process is done offline, the computational burden can be tolerated. However, as we will see further below, in many cases designing suitable edge controllers in practice increases the reachability probability such that practically one can assume the reachability is satisfied, and so there is no need to propagate the probability distribution.

Stopping region: By definition, the graph node B associated with the controller μ acts as the stopping region of the controller. However, if the process under the stabilizer hits another graph node before its corresponding graph node, we can stop the controller and pick the best controller from this intermediate node. Therefore, we can extend the stopping region for all controllers to the union of all nodes $\Psi := \cup_{i=1}^N B^i$. As a result, we will not necessarily have $\mathbb{P}(B^i|b, \mu^i) = 1$ since the process may hit some other node before B^i . However, we will have $\mathbb{P}(\Psi|b, \mu^i) = 1$ for all i in the absence of constraints.

Local controllers (simplified connecting strategy): To ease the connection step, and to have more distant nodes, we can precede each stabilizer by a time-varying controller (referred to as the edge controller). To illustrate this idea, consider two nodes B^i and B^j , where $B^i \not\subset \check{B}^j$; that is, B^i cannot be connected to B^j through μ^j . In this case, we can connect the underlying state nodes \mathbf{v}^i and \mathbf{v}^j in the state space by a finite trajectory e^{ij} (say of length ι) and then design a time-varying controller $\bar{\mu}_k^{ij}$, for $k = 0, 1, \dots, \iota$, to track this finite trajectory. Therefore, if the node B^i is in the basin of reachability of the pair $(\bar{\mu}_k^{ij}, \check{B}^j)$, then obviously B^i would be in the basin of reachability of the pair (μ^{ij}, B^j) , where the controller $\mu^{ij} = \{\bar{\mu}_{0:\iota}^{ij}, \mu^j\}$. We call μ^{ij} the (i, j) th local controller, as it connects the node B^i to the node B^j .

Graph: Formally, we define the constructed graph with the set of nodes $\mathbb{V} = \{B^i\}_{i=1}^N$ and the set of edges (or local controllers) $\mathbb{M} = \{\mu^{ij}\}$. The set of controllers available at B^i is denoted by $\mathbb{M}(i)$ (i.e. the set of edges starting from B^i). Similar to PRM, in which the path (final solution) is constructed as a concatenation of edges on the roadmap, in FIRM, the policy is constructed by the concatenation of the local policies. However, it is worth noting that with this construction we still perform planning in a continuous space and do not discretize the control space.

Local controllers versus macro-actions: By the term ‘macro-action’ we mean a sequence of controls (actions) (He et al., 2010, 2011). In other words, a macro-action is a sequence of open-loop policies. It is important to note that a

local controller is *not* a macro-action, but rather a sequence of policies (macro-policy), each of which is a mapping from belief space to the continuous control space. Using macro-actions results in an open-loop policy, which cannot compensate for the belief-state deviation from the planned path. However, under local controllers (macro-policies), the effect of noise can be compensated for, due to the feedback nature of the controllers, and thus, the belief can be steered towards a stopping region.

6.3. Belief SMDP

In this section, we reduce planning over the entire belief space into planning over a subset of belief space, which is actually the union of FIRM graph nodes; that is, $\Psi = \cup_j B^j$.

SMDP transition costs: First, we generalize the concept of one-step cost $c(b, u): \mathbb{B} \times \mathbb{U} \rightarrow \mathbb{R}_{\geq 0}$ to the one-step SMDP cost $C^s(b, \mu): \mathbb{B} \times \mathbb{M} \rightarrow \mathbb{R}_{\geq 0}$, which represents the cost of invoking the local controller $\mu(\cdot)$ at belief state b ; in other words,

$$C^s(b, \mu) := \sum_{t=0}^{\mathcal{T}} c(b_t, \mu(b_t) | b_0 = b) \quad (27)$$

where $\mathcal{T} := \mathcal{T}(\Psi|b, \mu)$.

Belief SMDP: According to the above definitions, the original POMDP, formulated using DP in Equation (6), can be reduced to an SMDP (Sutton et al., 1999) in the belief space, referred to as a *belief SMDP*:

$$J^s(b) = \min_{\mu \in \mathbb{M}(i)} C^s(b, \mu) + \int_{\Psi} p(b'|b, \mu) J^s(b') db', \forall b \in B^i, \forall i \quad (28)$$

The integration over the entire belief space in Equation (6) is reduced to integration over the sampled nodes (that is, Ψ) in Equation (28), as μ stops executing.

6.4. FIRM MDP

Graph transitions: The DP in (28), though computationally more tractable than the original POMDP, is defined on the continuous neighborhoods B^i and thus is still formidable to solve. However, for sufficiently small B^i and sufficiently smooth cost functions, the cost-to-go of all beliefs in B^i are approximately equal. Thus, we can define the graph-level transition cost and probabilities $C^g: \mathbb{V} \times \mathbb{M} \rightarrow \mathbb{R}$ and $\mathbb{P}^g: \mathbb{V} \times \mathbb{V} \times \mathbb{M} \rightarrow [0, 1]$ on the FIRM graph, in other words, over the finite space \mathbb{V} , such that $\mathbb{P}^g(B^j|B^i, \mu)$ is the transition probability from B^i to B^j under the local planner μ . Similarly, $C^g(B^i, \mu)$ denotes the cost of invoking local planner μ at the FIRM node B^i . Accordingly, $J^g: \mathbb{V} \rightarrow \mathbb{R}$ is the cost-to-go function over the FIRM nodes. These roadmap-level quantities are defined using the following ‘piecewise constant approximation’, which is an arbitrarily good approximation for smooth enough functions and

sufficiently small B^i :

$$\forall b \in B^i, \forall i \begin{cases} J^g(B^i) := J^s(b_c^i) \approx J^s(b) \\ C^g(B^i, \mu) := C^s(b_c^i, \mu) \approx C^s(b, \mu) \\ \mathbb{P}^g(\cdot|B^i, \mu) := \mathbb{P}(\cdot|b_c^i, \mu) \approx \mathbb{P}(\cdot|b, \mu) \end{cases} \quad (29)$$

where b_c^i is a representative point in B^i . For example, if B^i is a ball, the typical value for b_c^i is the center of B^i . This approximation essentially states that any belief in the region B^i is represented by b_c^i for the purpose of decision-making.

Obstacle-free FIRM MDP: Given the approximation in Equation (29), the DP equation in (28) becomes

$$\begin{aligned} J^g(B^i) &= J^s(b_c^i) = \min_{\mu \in \mathbb{M}(i)} C^s(b_c^i, \mu) + \int_{\Psi} p(b'|b_c^i, \mu) J^s(b') db' \\ &= \min_{\mu \in \mathbb{M}(i)} C^s(b_c^i, \mu) + \sum_j \int_{B^j} p(b'|b_c^i, \mu) J^s(b') db' \\ &\approx \min_{\mu \in \mathbb{M}(i)} C^g(B^i, \mu) + \sum_j \int_{B^j} p(b'|b_c^i, \mu) J^g(B^j) db' \\ &= \min_{\mu \in \mathbb{M}(i)} C^g(B^i, \mu) + \sum_j J^g(B^j) \mathbb{P}(B^j|b_c^i, \mu) \\ &= \min_{\mu \in \mathbb{M}(i)} C^g(B^i, \mu) + \sum_j J^g(B^j) \mathbb{P}^g(B^j|B^i, \mu), \quad \forall i \end{aligned} \quad (30)$$

Accordingly, we can get the graph feedback $\pi^g : \mathbb{V} \rightarrow \mathbb{M}$ through the following DP:

$$J^g(B^i) = \min_{\mu \in \mathbb{M}(i)} C^g(B^i, \mu) + \sum_j \mathbb{P}^g(B^j|B^i, \mu) J^g(B^j), \quad \forall i \quad (31a)$$

$$\pi^g(B^i) = \arg \min_{\mu \in \mathbb{M}(i)} C^g(B^i, \mu) + \sum_j \mathbb{P}^g(B^j|B^i, \mu) J^g(B^j), \quad \forall i \quad (31b)$$

Thus, the original POMDP over the entire belief space becomes a finite N -state MDP in Equation (31) defined on the finite set of FIRM nodes $\mathbb{V} = \{B^i\}_{i=1}^N$. We call the MDP in Equation (31) the FIRM MDP in the absence of obstacles. It is worth noting that $J^g(\cdot) : \mathbb{V} \rightarrow \mathbb{R}$ is the cost-to-go function over the FIRM nodes, which assigns a cost-to-go for every FIRM node B^i , and the mapping $\pi^g(\cdot) : \mathbb{V} \rightarrow \mathbb{M}$ is a mapping over the FIRM graph from FIRM nodes into the set of local controllers that returns the optimal local controller that has to be taken at any FIRM node. Given $C^g(B, \mu)$ for all (B, μ) pairs, the DP in Equation (31) can be solved offline using standard techniques such as the value/policy iteration to yield a feedback policy π^g over FIRM nodes $\{B^i\}$.

6.5. Incorporating obstacles into FIRM MDP

In the presence of obstacles (i.e. state or control constraints), we may not assume that the local controller $\mu^{ij}(\cdot)$

can drive any $b \in B^i$ into B^j with probability one. Instead, we have to specify the failure probabilities that the robot collides with an obstacle (hits the failure set F).

Let us generalize the transition probabilities by defining $\mathbb{P}(F|b, \mu)$ as the probability of hitting failure set F before hitting stopping region Ψ under μ starting from b . Similarly, we generalize \mathbb{P}^g such that $\mathbb{P}^g(F|B^i, \mu) := \mathbb{P}(F|b_c^i, \mu)$. Finally, we generalize the cost-to-go function by adding F to its input set, that is, $J^g : \{\mathbb{V}, F\} \rightarrow \mathbb{R}_{\geq 0}$, such that $J^g(F)$ is a user-defined suitably high cost for hitting obstacles. Note that the cost-to-go from the goal node is zero, that is, $J^g(B^{goal}) = 0$. Therefore, we can modify Equation (31) to incorporate constraints by repeating the procedure in the previous subsection to get the FIRM MDP in the presence of obstacles:

$$\begin{aligned} J^g(B^i) &= \min_{\mu \in \mathbb{M}(i)} C^g(B^i, \mu) + J^g(F) \mathbb{P}^g(F|B^i, \mu) \\ &\quad + \sum_j J^g(B^j) \mathbb{P}^g(B^j|B^i, \mu) \end{aligned} \quad (32a)$$

$$\begin{aligned} \pi^g(B^i) &= \arg \min_{\mu \in \mathbb{M}(i)} C^g(B^i, \mu) + J^g(F) \mathbb{P}^g(F|B^i, \mu) \\ &\quad + \sum_j J^g(B^j) \mathbb{P}^g(B^j|B^i, \mu) \end{aligned} \quad (32b)$$

All that is required to solve the above DP equation is the values of the costs $C^g(B^i, \mu)$ and the transition probability functions $\mathbb{P}^g(\cdot|B^i, \mu)$. Thus, the main difference from the obstacle-free case is the addition of a ‘failure’ state to the FIRM MDP along with associated probabilities of failure from various nodes B^i .

6.6. Overall policy π

The overall feedback $\pi : \mathbb{B} \rightarrow \mathbb{U}$ is generated by combining the global policy π^g on the graph and local policies $\{\mu^{ij}\}$. Suppose at the k th time step the active local controller is shown by μ_k^* . It remains unchanged ($\mu_{k+1}^* = \mu_k^*$) and keeps generating control signals based on the belief b_k at each time step, until the belief reaches the corresponding stopping region, Ψ . Once the belief enters the stopping region $\Psi = \cup_j B^j$, it is in a graph node, say $B_k^* \in \mathbb{V}$. Accordingly, the global policy π^g chooses the next local controller, that is, $\mu_{k+1}^* = \pi^g(B_k^*)$. Thus, this hybrid policy is stated as follows:

$$u_k = \pi(b_k) = \begin{cases} \mu_k^*(b_k), & \mu_k^* = \pi^g(B_{k-1}^*), & \text{if } b_k \in B_{k-1}^* \\ \mu_k^*(b_k), & \mu_k^* = \mu_{k-1}^*, & \text{if } b_k \notin \Psi \end{cases} \quad (33)$$

Initial controller: Given the initial belief b_0 , if b_0 is in one of the graph nodes, then we just choose the best local controller using π^g . However, if b_0 does not belong to any of the graph nodes, we first make a singleton set $B^0 = \{b_0\}$ and connect it to the graph nodes based on the connect methods discussed in Section 6.2. Denoting the outgoing edges (local controllers) from B^0 by $\mathbb{M}(0)$, we compute

the transition cost $C^g(B^0, \mu)$, the transition probabilities $\mathbb{P}^g(B^j|B^0, \mu)$ for all j , and failure probability $\mathbb{P}(F|B^0, \mu)$ for invoking local controllers $\mu \in \mathbb{M}(0)$ at B^0 . Then, we choose the best initial controller μ_0^* as

$$\mu_0^* = \begin{cases} \arg \min_{\mu \in \mathbb{M}(0)} \{C^g(B^0, \mu) + \mathbb{P}^g(F|B^0, \mu)J^g(F) \\ + \sum_j \mathbb{P}^g(B^j|B^0, \mu)J^g(B^j)\}, & \text{if } b_0 \notin \Psi \\ \pi^g(B^r), & \text{if } \exists r, \text{ s.t. } b_0 \in B^r \end{cases} \quad (34)$$

It is worth noting that computing μ_0^* is the only part of the computation that depends on the initial belief b_0 and that has to be performed online; in other words, if a large deviation occurs, μ_0^* is the only part that needs to be reproduced for the new initial point. After μ_0^* drives the system to a graph node, from there on the optimal policy is already known. Computing μ_0^* is feasible in real time as $\mathbb{M}(0)$ contains a limited number of edges.

6.7. Success probability

We would also like to quantify the quality of the solution π in the presence of obstacles. To this end, we require the probability of success of the policy π^g at the higher-level Markov chain on FIRM nodes given by Equation (32b). Without loss of generality let us assume that the first node B^1 is the goal node B^{goal} . The DP in Equation (32) has $N + 1$ states $\{F, B^{goal}, B^2, \dots, B^N\}$ that can be decomposed into three disjoint classes: the failure class $\{F\}$, the goal class $\{B^{goal}\}$, and the transient class $\{B^2, B^3, \dots, B^{N+1}\}$. The goal and failure classes are absorbing recurrent classes of this Markov chain. As a result, the transition probability matrix of this higher-level $N + 1$ state Markov chain can be decomposed as follows (Norris, 1997):

$$\mathcal{P} = \begin{bmatrix} \mathcal{P}_f & 0 & 0 \\ 0 & \mathcal{P}_{goal} & 0 \\ \mathcal{R}_f & \mathcal{R}_{goal} & \mathcal{Q} \end{bmatrix} \quad (35)$$

where $\mathcal{P}_{goal} = \mathbb{P}^g(B^1|B^1, \cdot) = 1$ and $\mathcal{P}_f = \mathbb{P}^g(F|F, \cdot) = 1$, since goal and failure classes are the absorbing recurrent classes; in other words, the system stops once it reaches the goal or it fails. \mathcal{Q} is a matrix that represents the transition probabilities between transient nodes in the transient class, whose (i, j) th element is $\mathcal{Q}[i, j] = \mathbb{P}^g(B^{i+1}|B^{j+1}, \pi^g(B^{j+1}))$. Vectors \mathcal{R}_{goal} and \mathcal{R}_f are $(N - 1) \times 1$ vectors that represent the probability of transient nodes $\mathbb{V} \setminus B^{goal}$ getting absorbed into the goal and failure node, respectively; that is, $\mathcal{R}_{goal}[j] = \mathbb{P}^g(B^1|B^{j+1}, \pi^g(B^{j+1}))$ and $\mathcal{R}_f[j] = \mathbb{P}^g(F|B^{j+1}, \pi^g(B^{j+1}))$. Then, it can be shown that the success probability from any desired node $B^i \in \mathbb{V} \setminus B^{goal}$ is as follows (Norris, 1997):

$$\begin{aligned} \mathbb{P}(\text{success}|B^i, \pi^g) &:= \mathbb{P}(B^{goal}|B^i, \pi^g) \\ &= \Gamma_{i-1}^T (I - \mathcal{Q})^{-1} \mathcal{R}_{goal}, \quad \forall i \geq 2 \end{aligned} \quad (36)$$

where Γ_i is a column vector with all elements equal to zero except the i th element, which is set to one. Note that the vector $\mathcal{P}^s = (I - \mathcal{Q})^{-1} \mathcal{R}_{goal}$ includes the success probability from every graph node.

In the next section, we will discuss the success probability in more detail in the context of probabilistic completeness. However, according to the computed $\mathbb{P}(\text{success}|B^i, \pi^g)$, one can compute the success probability from any given initial belief b_0 as

$$\mathbb{P}(\text{success}|b_0, \pi) = \sum_j \mathbb{P}(B^j|b_0, \mu_0^*) \mathbb{P}(\text{success}|B^j, \pi^g) \quad (37)$$

where μ_0^* is given by Equation (34). Then, this success probability is compared with a minimum acceptable success probability, denoted by p_{min} . If the condition $\mathbb{P}(\text{success}|b_0, \pi) > p_{min}$ is not satisfied, then the number of nodes in the graph has to be increased until the condition is satisfied. If, from the initial point b_0 , a successful policy in the class of admissible policies exists, then this procedure will eventually find a successful policy by increasing the number of nodes, due to the probabilistic completeness of the method, which is discussed in Section 7.1.

6.8. Generic FIRM algorithms

The generic algorithms for the offline construction of FIRM and online planning with FIRM are presented in Algorithms 3 and 4, respectively. Concrete instantiations of these algorithms for SLQG-FIRM are given in Algorithms 1 and 2, respectively.

Algorithm 3: Generic construction of the FIRM graph (offline)

- 1 Sample a set of stabilizer parameters $\mathcal{V} = \{\mathbf{v}^i\}$ and construct stabilizers $\mathbb{M} = \{\mu^i\}$ accordingly;
 - 2 Sample a set of belief nodes $\mathbb{V} = \{B^i\}$ such that they satisfy the reachability condition;
 - 3 Connect the belief nodes using local controllers μ^{ij} ;
 - 4 For each B^i and $\mu \in \mathbb{M}(i)$, compute the transition cost $C^g(B^i, \mu)$ and transition probabilities $\mathbb{P}^g(B^j|B^i, \mu)$ and $\mathbb{P}^g(F|B^i, \mu)$ associated with invoking μ at B^i ;
 - 5 Solve the graph DP in Equation (32) to compute feedback π^g over graph nodes, and compute π accordingly;
-

Single-query versus multi-query: As mentioned earlier, most approaches for planning in belief space in continuous state, action, and observation spaces result in query-dependent plans. However, one of the contributions of FIRM is that its construction does not depend on the query. In Algorithms 3 and 4, it is assumed that the goal is fixed for all queries; in this case in the planning phase we are only robust to changes in the starting point of the query. However, to make the algorithms also robust to changes in the

Algorithm 4: Generic planning (or replanning) on FIRM (online)

- 1 Given an initial belief b_0 , invoke the controller $\mu_0(\cdot)$ in Equation (34) to take the robot into some FIRM node B ;
 - 2 **while** $B \neq B^{goal}$ **do**
 - 3 Given the system is in FIRM node B , invoke the global feedback policy π^g to choose the local feedback policy $\mu(\cdot) = \pi^g(B)$;
 - 4 Let the local controller $\mu(\cdot)$ execute until the robot is absorbed into a FIRM node B' or until it hits the failure set;
 - 5 **if** Collision happens **then return** Collision;
 - 6 Update current node $B \leftarrow B'$;
-

goal belief, one can just move the last line of Algorithm 3 to the first line of Algorithm 4. Note that the computationally expensive part of Algorithm 3 is the computation of edge costs, which is independent of the start and goal locations of the submitted query.

6.9. Discussion

In summary, in FIRM we aim to transform the original POMDP problem into a belief SMDP problem and solve it on a subset of belief space. Given the smoothness of the cost function and transition probabilities, the solution of the FIRM MDP is arbitrarily close to the solution of the belief SMDP over FIRM nodes. The important characteristic of FIRM is that it is solved offline and thus performing the online phase of planning (or replanning) is computationally feasible in real time. To exploit the generic FIRM framework, one has to find (B, μ) pairs, where B is reachable (or αT -reachable) under μ , as FIRM nodes and edges. Also, transition costs and probabilities need to be computed. Finally, the corresponding FIRM MDP needs to be solved, which provides a global feedback policy on the graph that can be used in planning, as detailed in Algorithm 4. SLQG-FIRM, presented in Section 5, is an instance of FIRM, in which the design of local controllers μ^j and FIRM nodes B^i is based on the properties of SLQG controllers.

7. Probabilistic completeness under uncertainty

In this section, we extend the concept of probabilistic completeness of planning algorithms for deterministic systems to the concept of probabilistic completeness of planning algorithms under uncertainty, based on Agha-mohammadi et al. (2012b). Accordingly, in the next subsection, we discuss the probabilistic completeness of the FIRM-based algorithms. We start by reviewing the definition of success and probabilistic completeness in the deterministic case, and then we extend these definitions to the stochastic case.

Success in the deterministic case: In the deterministic case, such as conventional PRM, the outcome of the planning algorithm is a path. Thus, success is defined for paths:

for a given initial and goal point, a successful path is a path connecting the start point to the goal point which lies entirely in the obstacle-free space.

Probabilistic completeness in the deterministic case: In the absence of uncertainty, a sampling-based motion-planning algorithm is probabilistically complete if by increasing the number of samples, the probability of finding a successful path, if one exists, asymptotically approaches one.

A difference between the deterministic and the probabilistic case: In the presence of uncertainty, success cannot be defined for a path and has to be defined for a policy. Indeed, on a given path, different policies may result in different success probabilities. Moreover, under uncertainty, one can only assign a probability to reaching the goal. Thus, to define success for a policy we consider a threshold $p_{min} \in [0, 1]$ and decide about success or failure accordingly.

Successful policy: In the presence of uncertainty, the solution of the planning algorithm is a function, called a closed-loop policy or feedback. Therefore, success is defined for policies: for a given initial belief b_0 and goal region B^{goal} , a successful policy is a policy under which the probability of reaching the goal from the given initial point is greater than some predefined threshold p_{min} . In other words, π is successful for a given b_0 if $\mathbb{P}(\text{success}|b_0, \pi) := \mathbb{P}(B^{goal}|b_0, \pi) > p_{min}$.

Policy in sampling-based methods: In sampling-based methods, a policy is parametrized by a set of samples. These samples can be in the state or belief space, based on the algorithm. Let us denote these samples in a generic space by $\{\gamma_1, \gamma_2, \dots, \gamma_N\}$. Thus, we can highlight the dependency of the sampling-based policy on the samples by the notation $\pi(\cdot; \{\gamma_1, \gamma_2, \dots, \gamma_N\})$. The number of samples is denoted by N .

Strong probabilistic completeness under uncertainty (SPCUU): Suppose there exists a successful policy $\tilde{\pi}$. Then a sampling-based motion-planning algorithm is SPCUU if increasing the number of samples without bound causes the probability of finding a successful policy to approach one. In other words, if there exists a successful policy $\tilde{\pi}$, then we have the following property for the sampling-based policy π :

$$\lim_{N \rightarrow \infty} \mathbb{P}(B^{goal}|b_0, \pi) > p_{min} \quad (38)$$

where N is the number of samples in the sampling-based method.

Achieving an algorithm that is SPCUU requires searching in the entire space of policies, which is a computationally intractable task. Usually in solving POMDPs the space of admissible policies is restricted to a sufficiently rich subset of policy space, denoted by Π , within which the method searches for the best policy. Restricting the successful policy to the set Π , we define below a weaker notion of PCUU.

PCUU: Suppose that there exists a successful policy $\tilde{\pi} \in \Pi$. Then, a sampling-based motion-planning algorithm is PCUU if, when increasing the number of samples without bound, the probability of finding a successful policy approaches one. In other words, if there exists a successful policy $\tilde{\pi} \in \Pi$, then for the sampling-based policy π , we have $\lim_{N \rightarrow \infty} \mathbb{P}(B^{goal} | b_0, \pi) > p_{min}$.

As discussed in Section 6, in FIRM, inspired by the sampling-based PRM framework, this reduction from the entire function space to the restricted set of policies Π is performed by sampling feedback local planners and concatenating them. Therefore, the structure of local planners defines the set Π . Each local planner μ^{ij} is parametrized by its corresponding parameter \mathbf{v}^i . However, as mentioned in Section 6.2, we can consider the set $\mathcal{V} = \{\mathbf{v}^i\}$ as the set of underlying PRM nodes. Thus, any policy $\pi \in \Pi$ is parametrized by the set of underlying PRM nodes $\mathcal{V} = \{\mathbf{v}^i\}_{i=1}^N$. We highlight this dependency explicitly through the notation $\pi(\cdot; \mathcal{V})$. Therefore, the PCUU condition for FIRM can be written more explicitly as

$$\lim_{N \rightarrow \infty} \mathbb{P}(B^{goal} | b_0, \pi(\cdot; \mathcal{V})) > p_{min} \quad (39)$$

For a concrete instantiation of FIRM, we can explicitly characterize the set Π . For example, in SLQG-FIRM, Π is the set of all possible policies that can be generated by concatenating LQG controllers.

7.1. Probabilistic completeness of FIRM

Obviously, FIRM-based methods are not SPCUU algorithms. However, in this section, we show that under mild practical conditions, FIRM-based methods are PCUU algorithms. We first provide an analysis of the local planners in belief space, and then state the assumptions more rigorously.

Notation: The norm $\|\cdot\|$ is the supremum norm when it is applied to functions. The norm $\|\cdot\|_{op}$ is applied on operators and it stands for the operator norm (Keener, 2000). It is worth noting that in this section, by the word ‘continuous’, we mean ‘Lipschitz continuous’. Finally, we assume that \mathbb{X}_{free} is a compact set.

Hyper-state: $\mathcal{X} = (x, b) \in \mathbb{X}_h$ is referred to as hyper-state (or h-state), which is a state-belief pair. The space of all h-states is called hyper-state space (h-state space) $\mathbb{X}_h = \mathbb{X} \times \mathbb{B}$. Further, $p^\mu(\mathcal{X}' | \mathcal{X})$ denotes the one-step transition pdf induced by the local controller, μ , over the h-state space. Also, let $\mathbb{P}_n(S | \mathcal{X}, \mu)$ denote the transition probability from h-state \mathcal{X} into the set $S \subset \mathbb{X}_h$ in at most n steps.

Local planner and extended stopping region: The role of the (i, j) th local planner or local controller is to drive the belief from the region B^i to its stopping region B^j in the belief space (for ease of notation, we ignore the case where the controller can stop in any FIRM node, and we restrict its stopping region to B^j). In the presence of obstacles, we extend the concept of stopping region to include obstacles

also. The stopping regions $\{B^j\}$ in the belief space and the stopping region F in the state space can both be extended to the h-state space, respectively denoted by $\{\mathcal{B}^j\}$ and \mathcal{F} , where $\mathcal{B}^j \subset \mathbb{X}_h$ and $\mathcal{F} \subset \mathbb{X}_h$ are defined as

$$\mathcal{B}^j := \{(X, b) | X \in \mathbb{X}_{free}, b \in B^j\} \quad (40)$$

$$\mathcal{F} := \{(X, b) | X \in F, b \in \mathbb{B}\} \quad (41)$$

$$\mathcal{S}^j := \mathcal{B}^j \cup \mathcal{F}, \quad \overline{\mathcal{S}}^j := \mathbb{X}_h \setminus \mathcal{S}^j \quad (42)$$

where \mathcal{S}^j and $\overline{\mathcal{S}}^j$, respectively, denote the entire stopping region and transient region under the local controller μ^{ij} .

Absorption probability of local planners: If, under the dynamics induced by the local planner, the system reaches the target node \mathcal{B}^j , the local planner is considered to be successful; if the system hits an obstacle, the local planner is considered to have failed. The success probability of local planners (i.e. the absorption probability into FIRM nodes) is computed by solving the following integral equation that results from the law of total probability:

$$\begin{aligned} \mathbb{P}(\mathcal{B}^j | \mathcal{X}, \mu^{ij}) &= \int_{\mathbb{X}_h} p^{\mu^{ij}}(\mathcal{X}' | \mathcal{X}) \mathbb{P}(\mathcal{B}^j | \mathcal{X}', \mu^{ij}) d\mathcal{X}' \\ &= \int_{\mathcal{B}^j} p^{\mu^{ij}}(\mathcal{X}' | \mathcal{X}) d\mathcal{X}' \\ &\quad + \int_{\overline{\mathcal{S}}^j} p^{\mu^{ij}}(\mathcal{X}' | \mathcal{X}) \mathbb{P}(\mathcal{B}^j | \mathcal{X}', \mu^{ij}) d\mathcal{X}' \end{aligned} \quad (43)$$

where the second equality in Equation (43) follows from substituting the following conditions, inherited from FIRM construction, into the first integral:

$$\mathbb{P}(\mathcal{B}^j | \mathcal{X}, \mu^{ij}) = \begin{cases} 1, & \text{if } \mathcal{X} \in \mathcal{B}^j \\ 0, & \text{if } \mathcal{X} \in \mathcal{F} \end{cases} \quad (44)$$

Henceforth, we drop indices i and j to unclutter expressions. Thus, we can write

$$\begin{aligned} \mathbb{P}(\mathcal{B} | \mathcal{X}, \mu) &= \int_{\mathcal{B}} p^\mu(\mathcal{X}' | \mathcal{X}) d\mathcal{X}' \\ &\quad + \int_{\overline{\mathcal{S}}} p^\mu(\mathcal{X}' | \mathcal{X}) \mathbb{P}(\mathcal{B} | \mathcal{X}', \mu) d\mathcal{X}' \\ &= R(\mathcal{X}) + \mathbf{T}_S [\mathbb{P}(\mathcal{B} | \cdot, \mu)](\mathcal{X}) \end{aligned} \quad (45)$$

where the operator \mathbf{T}_S and the function $R(\mathcal{X})$ are defined as

$$\begin{aligned} \mathbf{T}_S[f(\cdot)](\mathcal{X}) &:= \int_{\overline{\mathcal{S}}} p^\mu(\mathcal{X}' | \mathcal{X}) f(\mathcal{X}') d\mathcal{X}', \\ R(\mathcal{X}) &:= \int_{\mathcal{B}} p^\mu(\mathcal{X}' | \mathcal{X}) d\mathcal{X}' \end{aligned} \quad (46)$$

The solution of the integral equation in Equation (45) is expressed in the following as a Liouville–Neumann

series (Keener, 2000), similar to the solution of the inhomogeneous Fredholm equation of second type (Keener, 2000):

$$\mathbb{P}(\mathcal{B}|\mathcal{X}, \mu) = \sum_{n=1}^{\infty} \mathbf{T}_{\mathcal{S}}^n [R(\cdot)] (\mathcal{X}) \quad (47)$$

We show that the series in Equation (47) is a convergent series by resorting to the following assumption, which is a weaker version of the aforementioned FIRM condition on the design of nodes and local controllers.

Assumption 1. We assume that there exists some time step N at which the controller stops with a positive probability. Mathematically, there exists an $N < \infty$ and a $\beta > 0$ such that $\mathbb{P}_N(\mathcal{S}^i|\mathcal{X}, \mu^i) \geq \beta > 0$, for all \mathcal{X} .

This assumption is almost always true, as it rephrases the role of the controller in driving the system toward the target region. For example, if we have Gaussian noise (as is the case in the SLQG-FIRM), the assumption is true at $N = 1$ regardless of the utilized controller.

Lemma 4. Given Assumption 1, we have

$$\begin{cases} \|\mathbf{T}_{\mathcal{S}}^n\|_{op} \leq 1, & n < N \\ \|\mathbf{T}_{\mathcal{S}}^n\|_{op} \leq 1 - \beta < 1, & n \geq N \\ \sum_{n=0}^{\infty} \|\mathbf{T}_{\mathcal{S}}^n\|_{op} \leq c < \infty \end{cases} \quad (48)$$

Proof. See Appendix E. \square

Corollary 1. The series $\sum_{n=0}^{\infty} \mathbf{T}_{\mathcal{S}}^n [R]$ is a convergent series, and therefore we can define the resolvent operator $(I - \mathbf{T}_{\mathcal{S}})^{-1} [R] = \sum_{n=0}^{\infty} \mathbf{T}_{\mathcal{S}}^n [R]$, where $\|(I - \mathbf{T}_{\mathcal{S}})^{-1}\|_{op} \leq c < \infty$.

Proof. See Appendix F. \square

According to Corollary 1, the success probability of the local controller μ can be written, using the defined resolvent operator, as

$$\mathbb{P}(\mathcal{B}|\mathcal{X}, \mu) = (I - \mathbf{T}_{\mathcal{S}})^{-1} [R(\cdot)] (\mathcal{X}) \quad (49)$$

As the first result of this section (Proposition 1), we aim to show that this absorption probability varies continuously with respect to changes in parameters of the local planner. However, we will first state two assumptions.

Assumption 2. We assume the local planning law and induced transition probabilities are smooth; in other words, we have assume the following.

- Local control laws are continuous in their parameters, that is, for the (i,j) th local controller, mapping $\mu^{ij}(\cdot; \mathbf{v}^j): \mathbb{B} \rightarrow \mathbb{U}$ is a continuous function in its parameter \mathbf{v}^j .
- The transition pdf on h -state, that is, $p(\mathcal{X}'|\mathcal{X}, u)$, is a continuous function of the control u ; in other words, there exists a $c_1 < \infty$ such that $\|p(\mathcal{X}'|\mathcal{X}, u) - p(\mathcal{X}'|\mathcal{X}, \check{u})\| \leq c_1 \|u - \check{u}\|$.

Finally, we state the following assumption, in which we emphasize the fact that, as $\mathbf{v} \rightarrow \check{\mathbf{v}}$, the transition probability induced by the local controller $\mu(\cdot; \mathbf{v})$ into the sets \mathcal{B} and $\check{\mathcal{B}}$ also has to converge, which is a reasonable assumption for a smooth control law.

Assumption 3. Consider the controllers $\mu(\cdot; \mathbf{v})$ and $\check{\mu}(\cdot; \check{\mathbf{v}})$, whose corresponding extended absorption regions are denoted by \mathcal{B} and $\check{\mathcal{B}}$, respectively. We assume that there exist real numbers $r > 0$ and $c' < \infty$ such that for $\|\mathbf{v} - \check{\mathbf{v}}\| \leq r$, we have

$$\|\mathbb{P}_1(\mathcal{B} \ominus \check{\mathcal{B}}|\mathcal{X}, \mu)\| \leq c' \|\mathbf{v} - \check{\mathbf{v}}\| \quad (50)$$

where \ominus is the symmetric difference operator; in other words, $\mathcal{B} \ominus \check{\mathcal{B}} = (\mathcal{B} \setminus \check{\mathcal{B}}) \cup (\check{\mathcal{B}} \setminus \mathcal{B})$.

Now we state the following proposition on the continuity of the success probability of local planners.

Proposition 1 (Continuity of absorption probabilities). Given Assumptions 1, 2, and 3, the absorption probability $\mathbb{P}(\mathcal{B}^i|b, \mu^i)$ is continuous in parameter \mathbf{v}^i for all i, j , and b .

Proof. See Appendix G. \square

Now we present the main result regarding the probabilistic completeness of FIRM-based methods.

Theorem 1. Given Assumptions 1, 2, and 3, any planning algorithm under uncertainty that is generated based on the FIRM framework (i.e. guarantees belief node reachability and induces a roadmap in the belief space with independent edge costs) is PCUU.

Proof. See Appendix H. \square

The basic idea of probabilistic completeness under uncertainty stems from an idea similar to the one in the path-isolation-based analysis for planners in deterministic systems. Roughly speaking, in the path isolation argument for sampling-based planners in the absence of uncertainty, if there is a successful path and a non-zero neighborhood of this path, in which every path is successful, we can eventually find a path in this neighborhood by increasing the number of samples unboundedly. Similarly, in the presence of uncertainty, if there is a successful policy, it is parametrized by some parameters (set of PRM nodes, in FIRM). Thus, if there exists a non-zero measure neighborhood of these parameters, within which selected parameters lead to a successful policy, we can eventually reach a successful policy by increasing the number of samples unboundedly and choosing samples in the target neighborhoods.

8. Experimental results

In this section, we first illustrate theoretical results from the previous sections on a planar robot in a small 3D planning domain. Then, we present planning results on a larger 3D state space. Finally, we report the results of the method on

a dynamical model of an eight-arm manipulator (sixteen-degree-of-freedom state space). This section is followed by a brief comparison with other state-of-the-art methods in Section 9.

8.1. Planar 3D omnidirectional robot: Illustrating steps in construction and planning with SLQG-FIRM

In this subsection, we focus on an omni-directional robot. Its state is composed of its 2D position in the plane and its heading angle. The goal in this section is to illustrate the steps of constructing SLQG-FIRM and planning with it.

8.1.1. Motion model. A three-wheel omnidirectional mobile robot is used in experiments with the nonlinear kinematic model given in Kalmár-Nagy et al. (2004). The state vector is composed of a 2D location and heading angle $x = [{}^1x, {}^2x, \theta]^T$ in a global world frame. Further, $u = [{}^1u, {}^2u, {}^3u]^T$ is the vector of controls, where ${}^i u$ is the linear velocity of the i th wheel. The motion model for this robot, in its original continuous form, is (Kalmár-Nagy et al., 2004)

$$\dot{x} = \mathbf{f}_c(x, u, w) = T(x)u + w \quad (51)$$

where w is the motion noise, which is drawn from a zero-mean Gaussian distribution, and

$$T(x) = \begin{pmatrix} -\frac{2}{3} \sin(\theta) & -\frac{2}{3} \sin(\frac{\pi}{3} - \theta) & \frac{2}{3} \sin(\frac{\pi}{3} + \theta) \\ \frac{2}{3} \cos(\theta) & -\frac{2}{3} \cos(\frac{\pi}{3} - \theta) & -\frac{2}{3} \cos(\frac{\pi}{3} + \theta) \\ \frac{1}{3r} & \frac{1}{3r} & \frac{1}{3r} \end{pmatrix} \quad (52)$$

where r is the distance of the wheels from the robot's center of mass. The discrete motion model is shown by

$$x_k = \mathbf{f}(x_{k-1}, u_{k-1}, w_{k-1}) \quad (53)$$

where $w_k \sim \mathcal{N}(0, \mathbf{Q})$ is the motion noise at the k th time step, which is drawn from a zero-mean Gaussian distribution with covariance matrix \mathbf{Q} . It can be shown that if we linearize this system, the linearized motion model satisfies the controllability condition in Property 1.

8.1.2. Observation model. In experiments, the robot is equipped with exteroceptive sensors that provide range and bearing measurements from existing landmarks (radio beacons) in the environment. The 2D location of the j th landmark is denoted by L_j . Measuring L_j can be modeled as follows:

$$\begin{aligned} {}^j z &= {}^j h(x, {}^j v) \\ &= [\|{}^j \mathbf{d}\|, \text{atan2}({}^j d_2, {}^j d_1) - \theta]^T + {}^j v, \quad {}^j v \sim \mathcal{N}(\mathbf{0}, {}^j \mathbf{R}) \end{aligned} \quad (54)$$

where ${}^j \mathbf{d} = [{}^j d_1, {}^j d_2]^T := [{}^1x, {}^2x]^T - L_j$. The vector ${}^j v$ is a state-dependent observation noise, with covariance

$${}^j \mathbf{R} = \text{diag}((\eta_r \|{}^j \mathbf{d}\| + \sigma_b^r)^2, (\eta_\theta \|{}^j \mathbf{d}\| + \sigma_b^\theta)^2) \quad (55)$$

In other words, the uncertainty (standard deviation) of the sensor reading increases as the robot gets farther from the landmarks; $\eta_r = \eta_\theta = 0.3$ determines this dependence, and $\sigma_b^r = 0.01$ m and $\sigma_b^\theta = 0.5^\circ$ are the bias standard deviations. A similar model for range sensing is used in Prentice and Roy (2009). We assume the robot observes all N_L landmarks at all times and their observation noises are independent. Thus, the total measurement vector is denoted by $z = [{}^1z^T, {}^2z^T, \dots, {}^{N_L}z^T]^T$, and, due to the independence of measurements of different landmarks, the observation model for all landmarks can be written as

$$z = h(x) + v, \quad v \sim \mathcal{N}(\mathbf{0}, \mathbf{R}), \quad \mathbf{R} = \text{diag}({}^1\mathbf{R}, \dots, {}^{N_L}\mathbf{R}) \quad (56)$$

It is straightforward to show that the linearized version of this observation model satisfies the observability condition in Property 1. Therefore, this entire system model (motion and sensing models) satisfies Property 1 and thus the SLQG-FIRM can be used for planning.

8.1.3. Construction of SLQG-FIRM nodes and edges. Figure 4(a) shows a sample environment, including obstacles, landmarks, and enumerated nodes in $({}^1x, {}^2x, \theta)$ space. Nodes are shown by blue triangles, which encode the position $({}^1x, {}^2x)$ and heading angle θ of the robot. Landmarks are shown by black stars. The corresponding FIRM nodes are computed and shown in Figure 4(b). All elements in Figure 4(b) are defined in $({}^1x, {}^2x, \theta)$ space but only the $({}^1x, {}^2x)$ portion of them is shown. Each $b_c^j \equiv (v^j, P_s^j)$ is illustrated by a red dot representing v^j , and a green ellipse representing the 3σ ellipse of covariance P_s^j . Each FIRM node B^j is a neighborhood around b_c^j . In the experiments, we define the node region using the component-wise version of Equation (18) to handle the error scale difference in position and orientation variables:

$$B^j = \{b \equiv (x, P) \mid |x - v^j| \prec \epsilon, |P - P_s^j| \prec \Delta\} \quad (57)$$

where $|\cdot|$ and \prec stand for the absolute value and component-wise comparison operators, respectively. We also set $\epsilon = [0.07 \text{ m}, 0.07 \text{ m}, 1^\circ]^T$ and $\Delta = \epsilon \epsilon^T$ to quantify B^j . The projection of B^j onto the space of estimation mean, that is, $B_x^j = \{\hat{x}^+ : |\hat{x}^+ - v^j| \prec \epsilon\}$, is a neighborhood around v^j , which is shown by a cyan rectangle centered at v^j . Projection of B^j onto the space of estimation covariances, that is, $B_p^j = \{P : |P - P_s^j| \prec \Delta\}$, is a neighborhood around P_s^j . However, in a 2D plot B_p^j cannot be shown due to its high dimension. Thus, we partially illustrate it only by two dashed green ellipses that represent 3σ covariances of $P_s^j - \Delta_d$ and $P_s^j + \Delta_d$, where Δ_d is the matrix Δ , whose off-diagonal elements are set to zero. For illustration purposes, both of these neighborhoods (i.e. B_x^j and B_p^j) are magnified 5x in Figure 4(b).

8.1.4. Transition costs and probabilities. After designing FIRM nodes and local controllers, the transition costs and

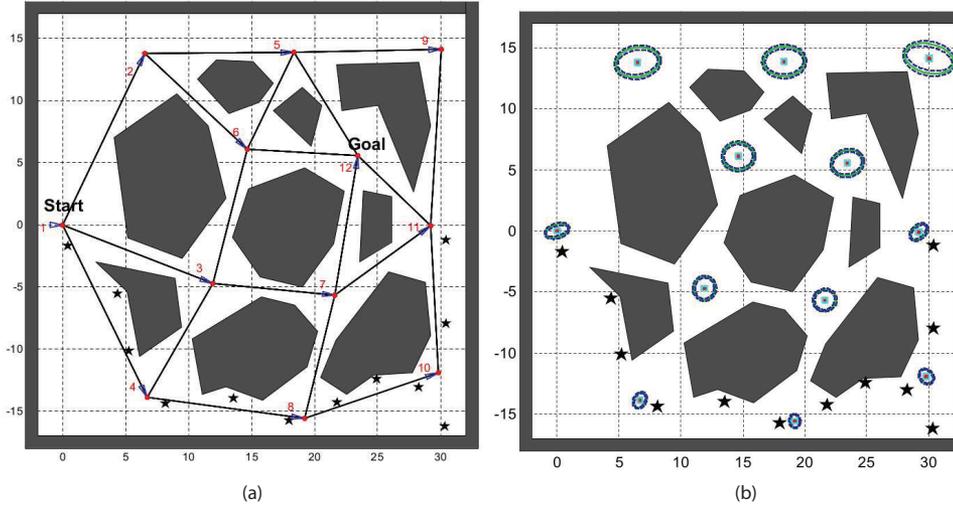


Fig. 4. (a) The underlying PRM graph. Gray polygons are the obstacles and black stars represent the landmarks' locations. (b) FIRM nodes corresponding to PRM nodes.

probabilities need to be computed. Based on the given task and needed accuracy, different approaches can be taken. Here, we use a particle-based approximation of the distribution to compute these quantities, where we use $M = 100$ particles. In other words, for every (B, μ) pair, we perform 100 runs. At every run, a sample path of state x , a sample path of estimation mean \hat{x}^+ , and a sample path of estimation covariance P is generated. If the filter of choice in the edge controller is the linearized Kalman filter (LKF) (Crassidis and Junkins, 2004; Simon, 2006), the covariance evolution is deterministic and there is no need to generate 100 different sample covariance paths. However, if the filter of choice in the edge controller is the EKF (Crassidis and Junkins, 2004; Simon, 2006), then we need to generate the sample covariance paths too, to take into account the stochasticity of the covariance matrix. Figure 5(a) depicts sample paths of the true state x and estimation mean \hat{x}^+ in green and dark red, respectively, for $M = 100$ particles. Note that when a true state path (green path) collides with an obstacle, the process stops and failure happens. However, in this figure, for illustration purposes, we continue the process and ignore the obstacles to better show the uncertainty tube and information availability in different parts of the space. As seen in Figure 5(a), the behavior of the true state on the edges which have access to more accurate observations is remarkably close to the planned behavior. In contrast, on the edges that get less informative observations, the controller cannot effectively compensate for deviations of the ground truth from the nominal path, which can lead to collision with obstacles.

To avoid clutter, Figure 5(b) depicts sample estimation covariance evolution only for a single particle. In this figure, we let the process and observation noise be zero, to keep the centers of the ellipses (i.e. estimation means) on the planned points. However, note that in general the estimation mean is affected by the noise (as is

seen in Figure 5(a)). Indeed, Figure 5(b) can be seen as the maximum-likelihood estimation uncertainty tube over the roadmap.

Let us denote the q th sample path for the true state by $x_{0:T^q}^{(q)}$, for the estimation mean by $\hat{x}_{0:T^q}^{(q)}$, and for the estimation covariance by $P_{0:T^q}^{(q)}$, where T^q is the stopping time of the q th particle in executing μ at B . Moreover, one can assign a weight to each particle q based on the probability of its occurrence. There are different ways proposed to compute these weights in the sequential Monte Carlo literature (Doucet et al., 2001). However, the main condition is that they have to sum to one, in other words, $\sum_{q=1}^M w^{(q)} = 1$. Here we simply consider $w^{(q)} = M^{-1}$. Note that if we particle μ^{ij} at B^i , all these quantities also need to have an ij superscript. Having these sample paths, we can compute the transition costs and probabilities associated with invoking μ^{ij} at B^i . For the collision probability, we have

$$\mathbb{P}^g(F|B^i, \mu^{ij}) = \mathbb{E}[\mathbb{I}_F|B^i, \mu^{ij}] \approx \sum_{q=1}^M w^{(q)} \mathbb{I}_F(x_{0:T^q}^{(q)}) \quad (58)$$

$$\mathbb{P}^g(B^j|B^i, \mu^{ij}) = 1 - \mathbb{P}^g(F|B^i, \mu^{ij}) \quad (59)$$

where \mathbb{I}_F is the failure indicator. $\mathbb{I}_F(x_{0:T^q}^{(q)})$ is one if there exists a time step $k \leq T^q$ such that $x_k \in F$; otherwise it is zero. Further, T^q , or more rigorously $T^{ij^{(q)}}$, is the stopping time of the q th particle in executing μ^{ij} at B^i . To compute $T^{ij^{(q)}}$, we only need to check the condition $b \in B^j$ at every time step and find the first time step where belief b enters the stopping region B^j . Thus, we can compute the mean stopping time as

$$\hat{T}^{ij} = \mathbb{E}[T^{ij}] \approx \sum_{q=1}^M w^{(q)} T^{ij^{(q)}} \quad (60)$$

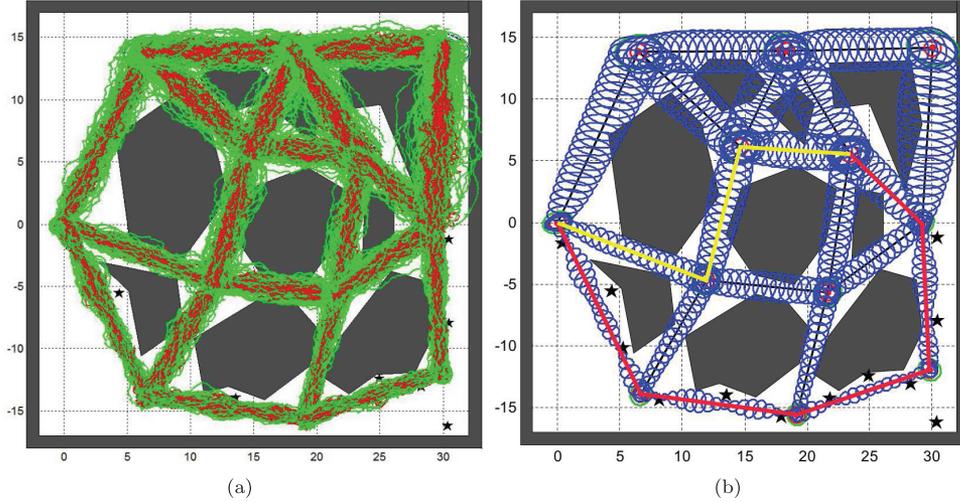


Fig. 5. Sample paths induced by controllers invoked at different nodes. (a) For $M = 100$ particles, sample ground truth paths and sample estimation mean paths are shown in green and dark red, respectively. (b) The most likely path under the optimal policy and the shortest path are shown in red and yellow respectively. The 3σ ML estimation uncertainty tube is drawn in blue.

To compute the filtering cost defined in Section 5.5, again we use the particle-based representation of belief,

$$\Phi^{ij} = \mathbb{E} \left[\sum_{k=1}^{\mathcal{T}^{ij}} \text{tr}(P_k) | B^i, \mu^{ij} \right] \approx \sum_{q=1}^M \sum_{k=1}^{\mathcal{T}^q} w^{(q)} \text{tr}(P_k^{(q)}) \quad (61)$$

where $P_k^{(q)}$ is the estimation covariance at the k th time step of the q th particle. Finally, the cost of taking μ^{ij} at B^i is as follows:

$$C^g(B^i, \mu^{ij}) = \alpha_1 \Phi^{ij} + \alpha_2 \widehat{\mathcal{T}}^{ij}$$

where we used the coefficients $\alpha_1 = 0.95$ and $\alpha_2 = 0.05$. Table 1 shows these quantities for several (B^i, μ^{ij}) pairs in FIRM corresponding to Figure 5.

8.1.5. Planning and replanning on FIRM. Plugging the computed transition costs and probabilities into Equation (32), we can solve the DP and compute the graph policy π^g . This process is performed once offline if the goal location is fixed. Figure 6(a) shows the policy π^g on the constructed FIRM in this example. Indeed, at every FIRM node B^i , the policy π^g decides which local controller has to be taken, which in turn aims to take the robot to the next FIRM node. Thus, the online part of the planning is significantly efficient and reduces to executing the controller and generating the control signal, which is almost an instantaneous computation.

Replanning: An important consequence of this framework is that replanning can be performed using FIRM efficiently. Suppose due to some unmodeled large disturbance, the robot's belief deviates significantly from the planned path; in other words, for some appropriate norm $\|\cdot\|$ on belief space we have $\|b_k - \mathbb{E}[b_k^p]\| > \varrho$, where b_k^p is the planned belief at the k th time step, and ϱ is the threshold for deciding if replanning is needed or not. In such cases,

replanning occurs based on Algorithm 2. In Figure 6(b), we illustrate a simple replanning process. In this figure it is assumed that an unmodeled large disturbance affects the system such that the estimation mean significantly deviates from the planned path. The deviated mean is shown on the figure as the 'restart point'. Thus, based on Algorithm 2, we connect this point to PRM. In Figure 6(b) the newly added PRM edges (i.e. $\mathcal{E}(0)$) are shown by dashed green lines. Then, for every edge in $\mathcal{E}(0)$, we design a local controller. Call the set of newly constructed local controllers $\mathbb{M}(0)$. For every $\mu \in \mathbb{M}(0)$ compute corresponding transition costs and probabilities. Finally, according to Bellman's principle of optimality, we use the precomputed costs-to-go $J^g(\cdot)$ to decide which controller has to be taken at the 'restart point' using Equation (34). Taking this controller, the belief enters into a FIRM node, and from there again we can use the precomputed π^g to control the robot toward the goal region.

We show the most likely path under π^g in red in Figure 5(b). The shortest path is also illustrated in Figure 5(b) in yellow. It can be seen that the 'most likely path under the best policy' detours from the shortest path to a path along which the filtering uncertainty is smaller and it is easier for the controller to avoid collisions.

8.2. Larger environment

In this section, we consider the same omnidirectional robot with the same observation model, and we perform planning in a larger environment (shown in Figure 7) whose size is almost 10,000 m². Every grid square is a 10 × 10 area. The standard deviation of the process noise is assumed to be 1 m for the positional degrees of freedom and 7° for the angular degree of freedom. We start with a five-node FIRM and at every step we randomly sample five more nodes until

Table 1. Computed costs for several node-controller pairs in FIRM using 100 particles.

(B^i, μ^{ij}) pair	$B^1, \mu^{1,4}$	$B^4, \mu^{4,8}$	$B^8, \mu^{8,10}$	$B^{10}, \mu^{10,11}$	$B^{11}, \mu^{11,12}$	$B^1, \mu^{1,3}$	$B^3, \mu^{3,6}$	$B^6, \mu^{6,12}$
$\mathbb{P}^g(B^i B^i, \mu^{ij})$	97%	95%	99%	77%	79%	87%	55%	79%
Φ^{ij}	18.5967	11.2393	6.8229	15.1148	26.2942	23.6183	48.8189	43.6207
$\mathbb{E}[T^{ij}]$	238.2	193.0	150.0	209.6	170.8	200.3	242.4	219.2
$\sigma[T^{ij}]$	21.8	28.7	15.1	24.5	22.6	22.7	30.1	26.7

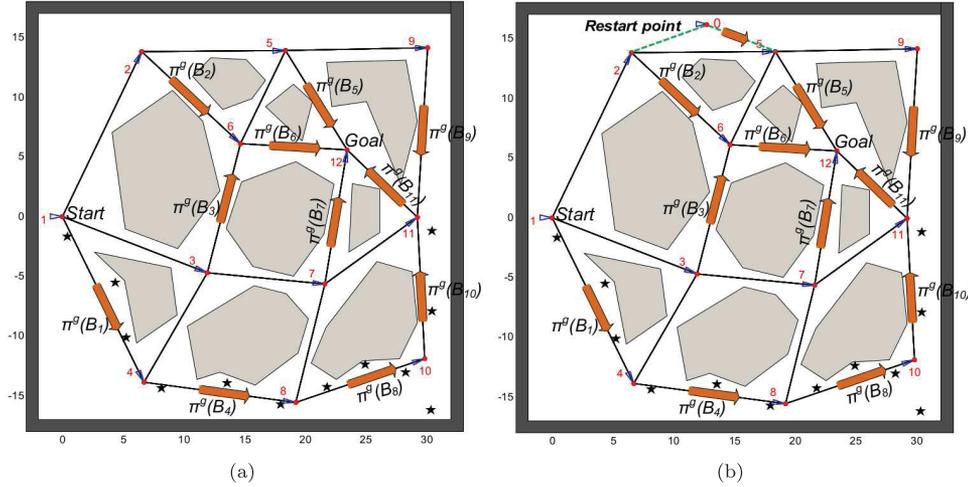


Fig. 6. Planning and replanning on FIRM. (a) Policy π^g resulted from solving the DP in Equation (31) is shown by red arrows. Indeed, for every FIRM node, the policy π^g tells us which controller has to be taken. (b) In this figure it is assumed that an unmodeled large disturbance affects the system such that the estimation mean significantly deviates from the planned path. The deviated mean is denoted by ‘restart point’ in the figure.

we reach 500 nodes. Thus, overall, we construct 100 FIRM graphs in this environment, for each of which we measure the construction time (cumulative) and compute the success probability. Plots in Figure 8 show these quantities as a function of the number of nodes for a sample run on an Intel i5 dual-core 1.7 GHz machine with 4 GB memory. Further, 50 particles are used for collision-checking, and every node in the underlying PRM is connected to its three nearest neighbors.

Basically, FIRM construction is an anytime algorithm in the sense that one can increase the number of nodes and stop enlarging the graph when a termination condition is satisfied such as: (i) achieving a desirable success probability or a desirable cost-to-go, (ii) no change being observed in the success probability or in the cost-to-go for a significant time, or (iii) exceeding the maximum time allowed for offline computation.

Again, as is seen in Figure 7, the highest-likelihood path under the optimal policy detours from the shortest path toward the more informative regions in the environment. As a result, it reduces the collision probability and at the same time increases the estimation accuracy and controller efficiency. However, it is important to note that the returned solution is not a single path, but it is a feedback law over the entire space. For the video of executing this plan (with fewer nodes to unclutter the video), see Extension 1.

We also conducted a simulation to illustrate the robustness of the method to large deviations. In this simulation, the robot is pushed away from the roadmap several times by some large disturbances, and replanning is performed online based on Algorithm 2. The video of this simulation is also available (see Extension 2).

8.3. Eight-arm manipulator

On a given graph, the number of paths between two given points grows exponentially with the size of graph. Thus, in the direct propagation of uncertainty on a roadmap, the number of edge costs and transition probabilities that need to be computed is exponential in the number of underlying PRM nodes (see Section 9 for a detailed analysis). As a result, when we deal with high-dimensional state spaces, where PRM needs to have many edges and nodes, it is not feasible to use the methods that perform direct uncertainty propagation. However, using FIRM, we only need to compute the costs and transition probabilities for as many edges as the underlying PRM has. Thus, we can easily increase the dimension to the level that PRM can handle, and the complexity of the algorithm is increased only by a constant factor (involving computation of costs and transition probabilities of a single edge). In the following experiment, we verify the effectiveness of FIRM in handling

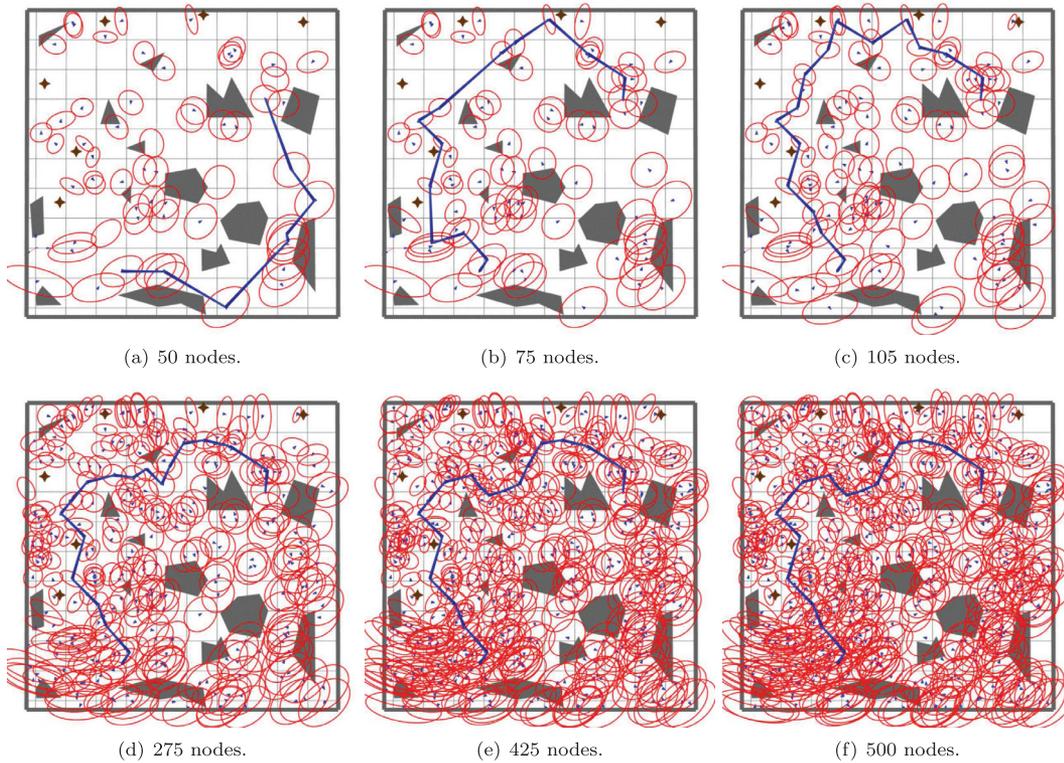


Fig. 7. (a)–(f) Different snapshots of the roadmap for 50, 75, 105, 275, 425, and 500 nodes, respectively. The most likely paths under the optimal plan are also shown in blue. Stars show the landmarks. The means and covariances of the FIRM node centers are shown by small blue triangles and their associated red ellipses, respectively. Also, see Extensions 1 and 2 regarding the video of planning with FIRM in this environment.

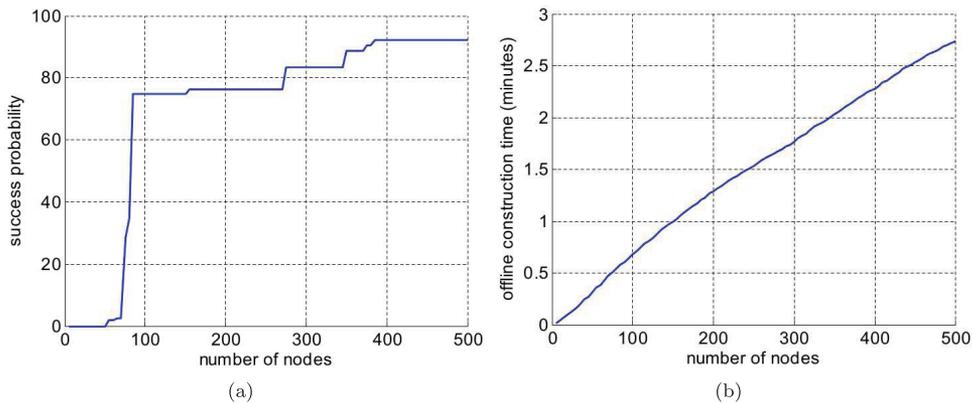


Fig. 8. A sample run showing (a) the success probability of the generated plan versus the number of nodes, as well as (b) the construction time (offline) for the plan.

high-dimensional systems through a simple example of an eight-arm manipulator. To the best of our knowledge, this is the first belief-space planner that can provide a plan over an entire roadmap for an eight-dimensional system while incorporating expensive costs in planning, such as computing collision probabilities. This experiment shows that FIRM can be used as a practical tool in many real-world problems.

8.3.1. Motion model. We consider an eight-arm manipulator with eight revolute joints in the plane. The state of the system is described by the angles of joints and their velocities $x = (\theta_1, \dots, \theta_8, \dot{\theta}_1, \dots, \dot{\theta}_8)^T$, and the available control is considered to be the angular acceleration (or torque) of joints $u = (\alpha_1, \alpha_2, \dots, \alpha_8)$. The process noise $w = (w_1, w_2, \dots, w_8)$ is modeled as a zero-mean Gaussian noise on angular accelerations. Therefore, the continuous

motion model for every link is $\ddot{\theta}_i = \alpha_i + w_i$, whose discrete version for the entire state can be written as

$$x_{k+1} = Ax_k + Bu_k + Gw_k \quad (62)$$

where

$$A = \begin{pmatrix} I_8 & I_8\delta t \\ 0_8 & I_8 \end{pmatrix}, \quad B = \begin{pmatrix} 0_8 \\ I_8\delta t \end{pmatrix}, \quad G = \begin{pmatrix} 0_8 \\ I_8\sqrt{\delta t} \end{pmatrix} \quad (63)$$

where δt is the time interval between two consecutive time steps, and I_n and 0_n are the identity matrix and square zero matrix of dimension n , respectively.

8.3.2. Observation model. We use the light–dark environment setting as the observation model, which is also used in Platt et al. (2010, 2011). In the light–dark environment, the accuracy of sensory readings is encoded by a gray level, in which the regions that have access to more accurate sensory readings are lighter than the regions that do not have access to such informative sensory readings. In this experiment, we assume that we measure the state of the system, but this measurement is more accurate as we get closer to the left wall on which our sensor is mounted. (This model is adopted from Platt et al. (2010).) Thus, we have $z = h(x) = [z^1, \dots, z^8]^T$, where

$$z^i = \theta_i + v^i, \quad v^i \sim \mathcal{N}(0, (\eta|x^i - l| + \sigma_b)^2) \quad (64)$$

where x^i is the x coordinate of the i th joint location, l is the location of the vertical wall, η defines the dependency of the noise standard deviation on the distance from the wall, and σ_b is the bias standard deviation. Figure 9 shows an example of such an environment, in which $l = -1.5$, $\eta = 0.1$, and $\sigma_b = 10^{-4}$. The full observation model can be written as

$$z_k = h(x_k) = Hx_k + Mv_k \quad (65)$$

where $H = [I_8, 0_8]$ and $M = I_8$.

8.3.3. Sampling stabilizer parameters. The described system is a controllable and observable system, and thus we adopt the SLQG controller as the stabilizing controller. Therefore, the parameters of the controller are points in the equilibrium space, as explained in Section 5. In other words, to generate sample nodes in the state space, we need to sample the configuration space $(\theta_1, \dots, \theta_8)$ and append zero angular velocities to it. To connect these samples in the state space we design simple trajectories between nodes, along which we accelerate the joints (angles) with constant acceleration until they are halfway to the next node, and after which we decelerate the joints until they reach the next node.

8.3.4. Construction of the SLQG-FIRM and planning with it. First, corresponding to sampled nodes in the state space,

we compute corresponding FIRM nodes and then design local controllers according to Algorithm 1. In a similar procedure to the one in the previous experiment, we compute the transition costs and probabilities.

To solve the DP, we need to characterize the goal nodes. In Figure 9, the goal region for the tip location of the manipulator is shown by a purple circle. We mark all PRM samples whose tip locations are within the goal region as goal nodes. Setting the cost-to-go to zero for all goal nodes, we solve the DP and compute the optimal feedback on the graph according to Algorithm 1. Finally, we execute the plan based on Algorithm 2 and we illustrate the propagation of the covariance of the manipulator tip in Figure 9 in red. As can be seen in Figure 9, there are two passages among the obstacles to reach the goal region. Although the right passage is closer to the initial configuration of the manipulator, the manipulator detours to a longer path through the left passage, because there is more accurate sensory information available in the left passage than the right one. As is seen in this example, the feedback plan minimizes the collision probability and picks the safest path, while being robust to deviations. In other words, if for any reason the manipulator deviates significantly from the underlying PRM, the feedback plan connects the deviated belief to the best neighboring FIRM node in real time, and continues the pre-computed plan from this node.

9. Comparison and limitations

In this section, we perform a short comparison of SLQG-FIRM against the two most related methods in the literature: BRM (Prentice and Roy, 2009) and LQG-MP on roadmaps (Van den Berg et al., 2011). Both methods are belief-space planners that exploit roadmap-based ideas. We compare the methods in terms of the offline construction and online planning complexity, and also in terms of some other properties, all listed in Table 2. In the following, we go over the complexity analysis that leads to the entries in this table. Afterwards, we discuss limitations of the SLQG-FIRM.

Offline construction complexity: In a general graph, the number of paths between two given nodes is exponential in the number of nodes N . For example, if each node in a graph is connected to k nearest neighbor nodes on the graph, for a search depth of d edges on the graph, the corresponding search tree contains k^d paths. Notice that each of these paths has d edges on it. Thus, if we directly (without using belief stabilizers) propagate the uncertainty on a roadmap for a depth of d , we have to evaluate the cost on dk^d edges. So, the asymptotic complexity of the overall problem is of the order $\mathcal{O}(Nk^N)$. Now, if computing the cost and transition probabilities associated with each edge under uncertainty is a constant multiplier $\mathcal{O}(c)$ of computing its cost in a deterministic case, the overall complexity of the methods based on direct belief propagation is $\mathcal{O}(cNk^N)$. On the other hand, in any variant of FIRM, due to the edge independence, only the cost of $\mathcal{O}(Nk)$ edges needs to be constructed as in PRM,

Table 2. Belief space roadmap-based method comparison (without using a heuristic in search algorithms).

Algorithm	Offline construction complexity (no heuristic)	Replanning (online planning) complexity	Future observations	System requirement	Valid region of plan	Collision probabilities
Generic PRM	$\mathcal{O}(Nk)$	$\mathcal{O}(k)$	————	Assumes a controller exists to drive the system from node to node	On the graph only	————
BRM	$\mathcal{O}(cNk^N)$	$\mathcal{O}(c\frac{N}{l}k^N)$ or $\mathcal{O}(cNk^N)$	Maximum likelihood observation	Well linearizable systems	Vicinity of the nominal path	Not considered
LQG-MP on roadmaps	$\mathcal{O}(cNk^N)$	$\mathcal{O}(cNk^N)$	All observations	Well linearizable systems	Vicinity of the nominal path	Simplified measures are used
Generic FIRM	$\mathcal{O}(cNk)$	$\mathcal{O}(ck)$	————	Assumes a controller exists to drive the system from node to node	Union of convergence regions of local controllers	————
SLQG-FIRM	$\mathcal{O}(cNk)$	$\mathcal{O}(ck)$ or $\mathcal{O}(1)$	All observations	Well linearizable, and linear controllable and observable systems	Vicinity of whole PRM (entire space for a dense PRM)	Computed

and thus the overall complexity of offline construction of FIRM is $\mathcal{O}(cNk)$.

Online planning (replanning) complexity: If the system deviates from the valid region of the plan, in direct propagation methods, edge costs need to be recomputed for all edges. So, in BRM and LQG-MP on roadmaps, the replanning complexity will be of the order $\mathcal{O}(Nk^N)$. If the cost of each edge is defined in such a way that it only depends on the belief at the start and end of the edge (i.e. does not depend on the belief along the edge), BRM can reduce the computation complexity to $\mathcal{O}(c(N/l)k^N)$ through covariance factorization techniques, where l is assumed to be the length (number of steps) of each edge. In FIRM, in the case of replanning (submitting a query with new starting point), it is only necessary to connect the deviated belief to k neighboring FIRM nodes. Thus, we only need to compute the cost for the k new edges. It is worth noting that if the underlying PRM is dense enough that the valid region of the local controllers covers the space, edge-cost computation in the replanning phase reduces to zero, because if the system deviates out of a valid region of a local planner, it will fall into the valid region of some other planner.

To reduce the complexity of the search algorithm in BRM and LQG-MP on roadmaps, it is assumed that the costs on different edges of the roadmap are independent. This heuristic can reduce the complexity of the algorithm, but it may still be significantly high compared to the PRM or FIRM. Moreover, this heuristic (edge-independent assumption) is not true without having belief stabilizers, and thus search algorithms relying on such a heuristic may result in solutions arbitrarily different from the true solution of the search algorithm. Assuming that no such heuristic is used

in the search algorithm, Table 2 summarizes the complexity of these algorithms.

The huge reduction in the computational complexity of the planning algorithm (in particular, in the online phase) opens many possibilities in utilizing POMDP solvers in real-world applications. Moreover, due to its sampling-based nature, it ameliorates the curse of dimensionality just as PRM does in the deterministic case. In other words, if the dimension of the system increases, we need a greater number of nodes N in the underlying PRM to capture the free space connectivity, in which case we cannot use direct methods due to their complexity. However, FIRM can tolerate the increase in the dimension since its complexity is only a constant multiplier of the PRM complexity.

9.1. Limitations of the SLQG-FIRM and future directions

In this section, we discuss limitations of the proposed method. It is important to distinguish which limitation is associated with the generic FIRM framework, and which limitation is associated with the particular presented instantiation of the FIRM, that is, the SLQG-FIRM. In some cases, we also propose ways to remedy these limitations as future research directions.

Stabilization time: The FIRM framework introduces the usage of belief stabilizers. However, the time needed for the belief stabilization procedure is added to the overall execution time. If the number of time steps along the nominal path is l , and the number of time steps needed for stabilization is τ , the extra time τ is usually negligible compared

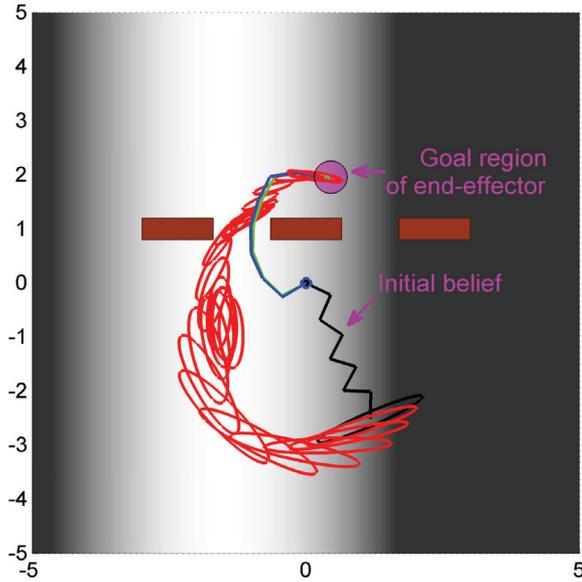


Fig. 9. A result of executing the FIRM plan for an eight-arm manipulator in a light–dark (sensing) environment. The manipulator is attached to the origin $(0, 0)$ and the purple region is the goal region for the manipulator tip. To unclutter the figure, we only show the uncertainty of the manipulator tip (end-effector). The initial mean and covariance are shown in black, and the evolution of the tip covariance during the plan execution is shown in red. The final estimation mean and the true configuration of the manipulator are shown in blue and green, respectively. Obstacles are shown in brown. The manipulator follows a longer but safer path to the goal region through the left passage, compared to the shorter but risky (with high collision probability) path through the right passage.

to l . However, τ can increase as the sensing uncertainty increases. In such a situation, one can consider two cases: if obstacles are close to the robot, it is indeed unsafe to move with a poor estimate, and it is indeed better to lose some time to gain more information, and then start moving. On the other hand, if there is no obstacle close to the robot, then one can increase the size of the corresponding FIRM node, and thus decrease the extra stopping time. Moreover, efficient sampling-based methods, which are aware of available information at different locations of the environment, and thus aware of the mean stabilization time, can be used to efficiently sample the nodes in the locations with lower mean stabilization times. These issues open up new directions for future research. However, if an application is very sensitive to the extra time, FIRM may not be a good choice for it, and methods such as BRM or LQG-MP can result in better guarantees on execution time.

Controllability and observability: As mentioned in Section 5, SLQG-FIRM works for systems that satisfy Property 1, which basically requires the linearized system about the PRM nodes to be controllable and observable. Although this includes a large class of systems, it excludes some

important systems, such as non-holonomic systems that are not linearly controllable about any point. It is worth noting that this is not a limitation of the generic FIRM framework, but a limitation of the SLQG-FIRM. More recent instantiations of FIRM, such as PLQG-based FIRM (Agha-mohammadi et al., 2012c) or DFL-based FIRM (Agha-mohammadi et al., 2012a), aim to relax the controllability requirements in Property 1 and thus can include non-holonomic systems as well. However, relaxing the observability assumption is still an open problem.

Gaussian beliefs: The reachability argument in the SLQG-FIRM is restricted to Gaussian beliefs. In other words, we cannot guarantee reachability to some predefined non-Gaussian beliefs with SLQG controllers. This issue is a subject of future research.

Increasing the uncertainty: Although it may rarely happen in practice, it is possible to have a situation that leads to an uncertainty growth during the belief-stabilization process. However, this issue can be addressed easily. Notice that FIRM nodes are known a priori. Thus, at the beginning of each stabilization procedure, we can compare the current belief with the stationary belief of the stabilizer. If the current belief has more information than the stationary belief (e.g. if all eigenvalues of the estimation covariance are strictly less than the corresponding eigenvalues of the stationary estimation covariance), we replan from the current belief based on Algorithm 2. Therefore, uncertainty will not be increased during the stabilization procedure.

Locally linearizable systems: If a linear representation of the system of interest cannot be obtained (e.g. if the system state lives in a discrete set of states), the class of methods that use the linearized system as a local approximation of the true system will not work. In this case, another class of methods can be adopted which can handle these systems much better, such as those in Smith and Simmons (2005), Kurniawati et al. (2008), and Kurniawati et al. (2011). Coming up with belief stabilizers that work in discrete state space settings to design a discrete-state variant of FIRM is also an area for future research.

Velocity reduction in dynamical systems: To apply SLQG-FIRM to dynamical systems, the underlying PRM samples need to be selected from the equilibrium space, in other words, they need to have zero velocity. As a result a reduction in the system's velocity is expected while trying to reach the FIRM nodes. However, in many applications, reducing the speed at nodes to gain the robustness, reliability, and scalability offered by FIRM may be a useful trade-off. Nevertheless, this reduction in speed may not be desirable for some applications where the system cannot (or should not) decrease its velocity. For such systems, Agha-mohammadi et al. (2013a) propose a FIRM variant based on periodic controllers which does not require a reduction in the system's velocity. However, designing more efficient variants of FIRM that can sample points with non-zero velocities without introducing periodicity in the system's motion is an interesting future research direction.

10. Conclusion

In this paper, we have proposed the FIRM framework for solving the motion-planning problem under motion and sensing uncertainties. This problem is originally a POMDP, whose solution is computationally intractable. Exploiting feedback controllers, we reduced it to a tractable FIRM MDP that can be solved using standard DP techniques. FIRM utilizes feedback controllers to create reachable node regions in belief space. An important consequence is that FIRM preserves the optimal substructure property on the roadmap and thus overcomes the curse of history in the original POMDP problem. Finally, by computing the collision probabilities, obstacles are also appropriately taken into account in planning on FIRM. We showed an instantiation of the abstract FIRM framework using SLQG controllers and illustrated the construction and planning results on it. By extending the probabilistic completeness concept to planners under uncertainty, we also showed that FIRM is probabilistically complete under uncertainty. We believe that FIRM provides an important step toward solving POMDPs and utilizing them as a practical tool for robot motion planning under uncertainty.

Acknowledgements

The authors are grateful to the anonymous reviewers for their helpful suggestions as well as Aditya Mahadevan and Daniel Tomkins for fruitful discussions and their help with experiments.

Funding

This work was supported in part by the NSF (award number RI-1217991). In addition, the work of Agha-mohammadi and Chakravorty is supported in part by the AFOSR (grant FA9550-08-1-0038), and the work of Agha-mohammadi and Amato is supported in part by the NSF (award numbers CNS-0551685, CCF-0833199, CCF-0830753, IIS-0917266, IIS-0916053 and EFRI-1240483), by the NSF/DNDO (award number 2008-DN-077-ARI018-02), by the NIH (award number NCI R25 CA090301-11), by the DOE (award numbers DE-FC52-08NA28616, DE-AC02-06CH11357, B575363 and B575366), by the THECB (NHARP award number 000512-0097-2009), by Samsung, by Chevron, by IBM, by Intel, by Oracle/Sun, and by the King Abdullah University of Science and Technology (award number KUS-C1-016-04).

References

Agha-mohammadi A, Chakravorty S and Amato N (2011) FIRM: Feedback controller-based Information-state RoadMap – a framework for motion planning under uncertainty. In: *International conference on intelligent robots and systems (IROS)*.
 Agha-mohammadi A, Chakravorty S and Amato N (2012a) Non-holonomic motion planning in belief space via dynamic feedback linearization-based FIRM. In: *International conference on intelligent robots and systems (IROS)*.

Agha-mohammadi A, Chakravorty S and Amato N (2012b) On the probabilistic completeness of the sampling-based feedback motion planners in belief space. In: *IEEE international conference on robotics and automation (ICRA)*.
 Agha-mohammadi A, Chakravorty S and Amato N (2012c) Periodic-feedback motion planning in belief space for non-holonomic and/or non-stoppable robots. Technical Report no. TR12-003, Texas A&M University, USA.
 Agha-mohammadi A, Chakravorty S and Amato N (2013a) Online replanning in belief space for dynamical systems: Towards handling discrete changes of goal location. In: *IEEE international conference on robotics and automation (ICRA): Workshop on combining task and motion planning*.
 Agha-mohammadi A, Chakravorty S and Amato N (2013b) Sampling-based stochastic control with constraints: A unified approach in state and information spaces. In: *The American control conference (ACC)*.
 Agha-mohammadi A, Agarwal S, Mahadeval A, et al. (2013c) Dynamic real-time replanning in belief space: an experimental study on physical mobile robots. Technical report TR 13-007, Texas A&M University, USA.
 Alterovitz R, Siméon T and Goldberg K (2007) The stochastic motion roadmap: A sampling framework for planning with Markov motion uncertainty. In: *Proceedings of robotics: Science and systems (RSS)*.
 Amato N, Bayazit B, Dale L, et al. (1998) OBPRM: An obstacle-based PRM for 3D workspaces. In: *International workshop on the algorithmic foundations of robotics*, pp. 155–168.
 Astrom K (1965) Optimal control of Markov decision processes with incomplete state estimation. *Journal of Mathematical Analysis and Applications* 10: 174–205.
 Bai H, Hsu D, Lee WS, et al. (2010) Monte Carlo value iteration for continuous-state POMDPs. In: Hsu D, Isler V, Latombe J-C, et al. (eds) *Algorithmic Foundations of Robotics IX (Springer Tracts in Advanced Robotics, vol. 68)*. New York, NY: Springer, pp. 175–191.
 Bertsekas D (1976) *Dynamic Programming and Stochastic Control*. Waltham, MA: Academic Press.
 Bertsekas D (2007) *Dynamic Programming and Optimal Control*, 3rd edn. Cambridge, MA: Athena Scientific.
 Bohlin R (2002) *Robot path planning*. PhD Thesis, Chalmers University of Technology, Sweden.
 Bry A and Roy N (2011) Rapidly-exploring random belief trees for motion planning under uncertainty. In: *International conference on robotics and automation (ICRA)*, pp. 723–730.
 Censi A, Calisi D, Luca AD, et al. (2008) A Bayesian framework for optimal motion planning with uncertainty. In: *Proceedings of the IEEE international conference on robotics and automation (ICRA)*, Pasadena, CA.
 Chakravorty S and Erwin RS (2011) Information space receding horizon control. In: *IEEE symposium on adaptive dynamic programming and reinforcement learning (ADPRL)*.
 Chakravorty S and Kumar S (2009) Generalized sampling based motion planners with application to nonholonomic systems. In: *Proceedings of the IEEE international conference on systems, man, and cybernetics*, San Antonio, TX.
 Chakravorty S and Kumar S (2011) Generalized sampling-based motion planners. *IEEE Transactions on Systems, Man, and Cybernetics, Part B* 41(3): 855–866.
 Choset H, Lynch KM, Hutchinson S, et al. (2005) *Principles of Robot Motion: Theory, Algorithms, and Implementations*. Cambridge, MA: MIT Press.

- Cormen TH, Leiserson CE, Rivest RL, et al. (2001) *Introduction to Algorithms*. 2nd edn. Cambridge, MA: MIT Press.
- Crassidis J and Junkins J (2004) *Optimal Estimation of Dynamic Systems*. Boca Raton, FL: Chapman & Hall/CRC.
- Doucet A, de Freitas J and Gordon N (2001) *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer.
- Guibas L, Hsu D, Kurniawati H, et al. (2008) Bounded uncertainty roadmaps for path planning. In: *International workshop on algorithmic foundations of robotics (WAFR)*.
- He R, Brunskill E and Roy N (2010) PUMA: Planning under uncertainty with macro-actions. In: *Proceedings of the twenty-fourth conference on artificial intelligence (AAAI)*, Atlanta, GA.
- He R, Brunskill E and Roy N (2011) Efficient planning under uncertainty with macro-actions. *Journal of Artificial Intelligence Research* 40: 523–570.
- Hsu D (2000) *Randomized single-query motion planning in expansive spaces*. PhD Thesis, Stanford University, CA.
- Huynh V and Roy N (2009) iLQG: Combining local and global optimization for control in information space. In: *IEEE international conference on robotics and automation (ICRA)*.
- Kaelbling LP, Littman ML and Cassandra AR (1998) Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101: 99–134.
- Kalmár-Nagy T, D'Andrea R and Ganguly P (2004) Near-optimal dynamics trajectory generation and control of an omnidirectional vehicle. *Robotics and Autonomous Systems* 46(1): 47–64.
- Karaman S and Frazzoli E (2011) Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research* 30(7): 846–894.
- Kavraki L, Kolountzakis M and Latombe J (1998) Analysis of probabilistic roadmaps for path planning. *IEEE Transactions on Robotics and Automation* 14: 166–171.
- Kavraki L, Latombe J, Motwani R, et al. (1995) Randomized query processing in robot motion planning. In: *Proceedings of the ACM symposium on the theory of computing*, pp. 353–362.
- Kavraki L, Švestka P, Latombe J, et al. (1996) Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation* 12(4): 566–580.
- Keener JP (2000) *Principles of Applied Mathematics: Transformation and Approximation*. 2nd edn. Boulder, CO: Westview Press.
- Kumar PR and Varaiya PP (1986) *Stochastic Systems: Estimation, Identification, and Adaptive Control*. Englewood Cliffs, NJ: Prentice-Hall.
- Kurniawati H, Bandyopadhyay T and Patrikalakis N (2012) Global motion planning under uncertain motion, sensing, and environment map. *Autonomous Robots* 33(3): 255–272.
- Kurniawati H, Du Y, Hsu D, et al. (2010) Motion planning under uncertainty for robotic tasks with long time horizons. *The International Journal of Robotics Research* 30: 308–323.
- Kurniawati H, Du Y, Hsu D, et al. (2011) Motion planning under uncertainty for robotic tasks with long time horizons. *The International Journal of Robotics Research* 30(3): 308–323.
- Kurniawati H, Hsu D and Lee W (2008) SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In: *Proceedings of robotics: Science and systems (RSS)*.
- Ladd A and Kavraki L (2004) Measure theoretic analysis of probabilistic path planning. *IEEE Transactions on Robotics and Automation* 20(2): 229–242.
- Lavelle S and Kuffner J (2001) Randomized kinodynamic planning. *The International Journal of Robotics Research* 20: 378–400.
- Lozano-Perez T (1983) Spatial planning: A configuration space approach. *IEEE Transactions on Computers* 100(2): 108–120.
- Madani O, Hanks S and Condon A (1999) On the undecidability of probabilistic planning and infinite-horizon partially observable Markov decision problems. In: *Proceedings of the sixteenth conference on artificial intelligence (AAAI)*, pp. 541–548.
- Melchior N and Simmons R (2007) Particle RRT for path planning with uncertainty. In: *International conference on intelligent robots and systems (IROS)*.
- Meyn S and Tweedie RL (2009) *Markov Chains and Stochastic Stability*, 2nd edn. Cambridge: Cambridge University Press.
- Missiuro P and Roy N (2006) Adapting probabilistic roadmaps to handle uncertain maps. In: *International conference on robotics and automation (ICRA)*.
- Nakhaei A and Lamiroux F (2008) A framework for planning motions in stochastic maps. In: *IEEE international conference on robotics and automation (ICRA)*.
- Norris JR (1997) *Markov Chains*. Cambridge: Cambridge University Press.
- Ong SCW, Png SW, Hsu D, et al. (2010) Planning under uncertainty for robotic tasks with mixed observability. *The International Journal of Robotics Research* 29(8): 1053–1068.
- Papadimitriou C and Tsitsiklis JN (1987) The complexity of Markov decision processes. *Mathematics of Operations Research* 12(3): 441–450.
- Patil S, van den Berg J and Alterovitz R (2012) Estimating probability of collision for safe planning under Gaussian motion and sensing uncertainty. In: *IEEE international conference on robotics and automation (ICRA)*.
- Pineau J, Gordon G and Thrun S (2003) Point-based value iteration: An anytime algorithm for POMDPs. In: *International joint conference on artificial intelligence*, pp. 1025–1032.
- Pineau J, Gordon G and Thrun S (2006) Anytime point based approximations for large POMDPs. *Journal of Artificial Intelligence Research* 27: 335–380.
- Platt R, Kaelbling L, Lozano-Perez T, et al. (2011) Efficient planning in non-Gaussian belief spaces and its application to robot grasping. In: *Proceedings of the international symposium of robotics research (ISRR)*.
- Platt R, Tedrake R, Kaelbling L, et al. (2010) Belief space planning assuming maximum likelihood observatoins. In: *Proceedings of robotics: Science and systems (RSS)*.
- Porta JM, Vlassis N, Spaan MTJ, et al. (2006) Point-based value iteration for continuous POMDPs. *Journal of Machine Learning Research* 7: 2329–2367.
- 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.
- Simon D (2006) *Optimal State Estimation: Kalman, H-infinity, and Nonlinear Approaches*. New York, NY: John Wiley and Sons.
- Smallwood RD and Sondik EJ (1973) The optimal control of partially observable Markov processes over a finite horizon. *Operations Research* 21(5): 1071–1088.

- Smith T and Simmons R (2005) Point-based POMDP algorithms: Improved analysis and implementation. In: *Proceedings of the international conference on uncertainty in artificial intelligence (UAI)*.
- Sniedovich M (2006) Dijkstra's algorithm revisited: The dynamic programming connexion. *Control and Cybernetics* 35(3): 599–620.
- Spaan M and Vlassis N (2005) Perseus: Randomized point-based value iteration for POMDPs. *Journal of Artificial Intelligence Research* 24: 195–220.
- Sutton R, Precup D and Singh S (1999) Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence* 112: 181–211.
- Švestka P and Overmars M (1997) Motion planning for car-like robots using a probabilistic learning approach. *The International Journal of Robotics Research* 16(2): 119–143.
- Thrun S, Burgard W and Fox D (2005) *Probabilistic Robotics*. Cambridge, MA: MIT Press.
- Toit ND and Burdick JW (2010) Robotic motion planning in dynamic, cluttered, uncertain environments. In: *International conference on robotics and automation (ICRA)*.
- Van den Berg J, Abbeel P and Goldberg K (2010) LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information. In: *Proceedings of robotics: Science and systems (RSS)*.
- Van den Berg J, Abbeel P and Goldberg K (2011) LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information. *The International Journal of Robotics Research* 30(7): 895–913.
- Van den Berg J and Overmars M (2007) Kinodynamic motion planning on roadmaps in dynamic environments. In: *IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pp. 4253–4258.
- Van den Berg J, Patil S and Alterovitz R (2011) Motion planning under uncertainty using differential dynamic programming in belief space. In: *Proceedings of the international symposium of robotics research (ISRR)*.
- Van den Berg J, Patil S and Alterovitz R (2012) Efficient approximate value iteration for continuous Gaussian POMDPs. In: *Proceedings of the AAAI conference on artificial intelligence (AAAI)*.
- Vitus MP and Tomlin CJ (2011) Closed-loop belief space planning for linear, Gaussian systems. In: *International conference on robotics and automation (ICRA)*, pp. 2152–2159.

Appendices

A. Index to Multimedia Extensions

The multimedia extension page is found at <http://www.ijrr.org>

Table of Multimedia Extensions

Extension Type	Description
1	Video Executing the FIRM plan in the environment shown in Figure 7
2	Video Real-time replanning with FIRM, which shows the robustness of the method to large disturbances

B. Time-varying LQG controller

The time-varying LQG controller is often used to track a pre-planned trajectory (also called a nominal, desired, or open-loop trajectory) in the presence of process and observation noise. In principal it is designed (and optimal) for linear systems with Gaussian noise, but it can also be utilized for stabilizing nonlinear systems locally around the planned trajectory. An LQG controller is composed of a KF as an estimator and an LQR as a controller. At every time step k , the KF provides the *a posteriori* distribution (belief) b_k over the system state, and the LQR generates control u_k based on b_k .

In this appendix, we first discuss the system linearization and planned nominal trajectory, and then discuss the KF, LQR, and LQG corresponding to this nominal trajectory. Consider the nonlinear partially observable state-space equations of the system as follows:

$$x_{k+1} = f(x_k, u_k, w_k), \quad w_k \sim \mathcal{N}(0, Q_k) \quad (66a)$$

$$z_k = h(x_k, v_k), \quad v_k \sim \mathcal{N}(0, R_k) \quad (66b)$$

A planned nominal trajectory for this system is a sequence of planned states $(x_k^p)_{k \geq 0}$ and planned controls $(u_k^p)_{k \geq 0}$ such that it is consistent with the noiseless dynamics model; in other words, we have

$$x_{k+1}^p = f(x_k^p, u_k^p, 0) \quad (67)$$

The planned trajectory can be a finite sequence of some length N . The role of a closed-loop stochastic controller, during the trajectory tracking, is to compensate for the robot's deviations from the planned trajectory and to keep the robot close to the planned trajectory in the sense of minimizing the following quadratic cost:

$$J = \mathbb{E} \left[\sum_{k \geq 0} (x_k - x_k^p)^T W_x (x_k - x_k^p) + (u_k - u_k^p)^T W_u (u_k - u_k^p) \right] \quad (68)$$

where W_x and W_u are positive-definite weight matrices for the state and control costs, respectively.

Since the state space is not fully observable and is only partially observable, we do not have access to the perfect value of the state x_k , and thus, we provide the estimate x_k^+ of the state x_k based on the available observations $z_{0:k}$ from the beginning up to the current time step. Then, based on the separation principle (Kumar and Varaiya, 1986; Bertsekas, 2007), it can be shown that in a linear system with Gaussian noise, the above minimization in terms of the error $x_k - x_k^p$ is equivalent to performing two separate minimizations based on the estimation error $x_k - \hat{x}_k^+$ and the controller error $\hat{x}_k^+ - x_k^p$, whose summation is the same as the original main error $x_k - x_k^p = (x_k - \hat{x}_k^+) + (\hat{x}_k^+ - x_k^p)$, where $\hat{x}_k^+ = \mathbb{E}[x_k^+] = \mathbb{E}[x_k | z_{0:k}]$. As a major consequence, the design of the stochastic controller with a partially observable state space (LQG) reduces to designing a controller

with a fully observable state (LQR) and designing an estimator (KF), separately. In the following, we first discuss the linearization of a nonlinear model. Then we discuss how a KF and an LQR can be designed for this linearized system. Finally, we combine them to construct a time-varying LQG controller.

Model linearization: Given a nominal trajectory $(x_k^p, u_k^p)_{k \geq 0}$, we linearize the dynamics and observation model in Equation (66) as follows:

$$x_{k+1} = f(x_k^p, u_k^p, 0) + A_k(x_k - x_k^p) + B_k(u_k - u_k^p) + G_k w_k, \quad w_k \sim \mathcal{N}(0, Q_k) \quad (69a)$$

$$z_k = h(x_k^p, 0) + H_k(x_k - x_k^p) + M_k v_k, \quad v_k \sim \mathcal{N}(0, R_k) \quad (69b)$$

where

$$\begin{aligned} A_k &= \frac{\partial f}{\partial x}(x_k^p, u_k^p, 0), \quad B_k = \frac{\partial f}{\partial u}(x_k^p, u_k^p, 0), \\ G_k &= \frac{\partial f}{\partial w}(x_k^p, u_k^p, 0), \quad H_k = \frac{\partial h}{\partial x}(x_k^p, 0), \\ M_k &= \frac{\partial h}{\partial v}(x_k^p, 0) \end{aligned} \quad (70)$$

Now, let us define the following errors:

- LQG error (main error): $e_k = x_k - x_k^p$
- KF error (estimation error): $\tilde{e}_k = x_k - \hat{x}_k^+$
- LQR error (estimation of LQG error): $\hat{e}_k^+ = \hat{x}_k^+ - x_k^p$

Note that these errors are linearly dependent: $e_k = \hat{e}_k^+ + \tilde{e}_k$. Also, defining $\delta u_k = u_k - u_k^p$ and $\delta z_k = z_k - z_k^p := z_k - h(x_k^p, 0)$, we can rewrite the above linearized models as follows:

$$e_{k+1} = A_k e_k + B_k \delta u_k + G_k w_k \quad (71a)$$

$$\delta z_k = H_k e_k + M_k v_k \quad (71b)$$

KF: In Kalman filtering, we aim to provide an estimate of the system's state based on the available partial information we have obtained up to time k , that is, $z_{0:k}$. The state estimate is a random vector denoted by x_k^+ , whose distribution is the conditional distribution of the state on the obtained observations so far, which is called belief and is denoted by b_k :

$$b_k = p(x_k^+) = p(x_k | z_{0:k}) \quad (72)$$

$$\hat{x}_k^+ = \mathbb{E}[x_k | z_{0:k}] \quad (73)$$

$$P_k = \mathbb{C}[x_k | z_{0:k}] \quad (74)$$

where $\mathbb{E}[\cdot]$ and $\mathbb{C}[\cdot]$ are the conditional expectation and conditional covariance operators, respectively. In the Gaussian case, we have $b_k = \mathcal{N}(\hat{x}_k^+, P_k)$; in other words, the belief can be characterized only by its mean and covariance, that is, $b_k \equiv (\hat{x}_k^+, P_k)$.

Kalman filtering consists of two steps at every time stage: a prediction step and an update step. In the prediction step, the mean and covariance of prior x_k^- are computed. For the system in Equation (71), the prediction step is

$$\hat{e}_{k+1}^- = A_k \hat{e}_k^- + B_k \delta u_k \quad (75)$$

$$P_{k+1}^- = A_k P_k^- A_k^T + G_k Q_k G_k^T \quad (76)$$

In the update step, the mean and covariance of posterior x_k^+ are computed. For the system in Equation (71), the update step is

$$K_k = P_k^- H_k^T (H_k P_k^- H_k^T + M_k R_k M_k^T)^{-1} \quad (77)$$

$$\hat{e}_{k+1}^+ = \hat{e}_{k+1}^- + K_{k+1} (\delta z_{k+1} - H_{k+1} \hat{e}_{k+1}^-) \quad (78)$$

$$P_{k+1}^+ = (I - K_{k+1} H_{k+1}) P_{k+1}^- \quad (79)$$

Note that

$$\hat{x}_k^+ = \mathbb{E}[x_k | z_{0:k}] = x_k^p + \hat{e}_k^+ = x_k^p + \mathbb{E}[e_k | z_{0:k}] \quad (80)$$

$$P_k = \mathbb{C}[x_k | z_{0:k}] = P_k^+ = \mathbb{C}[e_k | z_{0:k}] \quad (81)$$

LQR controller: Once we obtain the belief from the filter, a controller can generate an optimal control signal accordingly. In other words, we have a time-varying mapping μ_k from the belief space into the control space that generates an optimal control based on the given belief $u_k = \mu_k(b_k)$ at every time step k . The LQR controller is of this kind and it is optimal in the sense of minimizing the following cost:

$$\begin{aligned} J_{LQR} &= \mathbb{E} \left[\sum_{k \geq 0} (\hat{x}_k^+ - x_k^p)^T W_x (\hat{x}_k^+ - x_k^p) + (u_k - u_k^p)^T W_u (u_k - u_k^p) \right] \\ &= \mathbb{E} \left[\sum_{k \geq 0} (\hat{e}_k^+)^T W_x (\hat{e}_k^+) + (\delta u_k)^T W_u (\delta u_k) \right] \end{aligned} \quad (82)$$

The linear control law that minimizes this cost function for a linear system is of the form

$$\delta u_k = -L_k \hat{e}_k^+ \quad (83)$$

where the time-varying feedback gains L_k can be computed recursively as follows:

$$L_k = (B_k^T S_{k+1} B_k + W_u)^{-1} B_k^T S_{k+1} A_k \quad (84)$$

$$S_k = W_x + A_k^T S_{k+1} A_k - A_k^T S_{k+1} B_k L_k \quad (85)$$

If the nominal path is of length N , then $S_N = W_x$ is the initial condition of the above recursion, which is solved backwards in time. Note that the full control is $u_k = u_k^p + \delta u_k$.

LQG controller: Plugging the obtained LQR control law into the Kalman filtering equations, we obtain the following error dynamics for the defined errors:

$$\begin{aligned} \begin{pmatrix} e_{k+1} \\ \tilde{e}_{k+1} \end{pmatrix} &= \begin{pmatrix} A_k - B_k L_k & B_k L_k \\ 0 & A_k - K_{k+1} H_{k+1} A_k \end{pmatrix} \begin{pmatrix} e_k \\ \tilde{e}_k \end{pmatrix} \\ &+ \begin{pmatrix} G_k & 0 \\ G_k - K_{k+1} H_{k+1} G_k & -K_{k+1} M_{k+1} \end{pmatrix} \begin{pmatrix} w_k \\ v_{k+1} \end{pmatrix} \end{aligned} \quad (86)$$

or equivalently,

$$\begin{aligned} \begin{pmatrix} e_{k+1} \\ \hat{e}_{k+1}^+ \end{pmatrix} &= \begin{pmatrix} A_k & -B_k L_k \\ K_{k+1} H_{k+1} A_k A_k - B_k L_k - K_{k+1} H_{k+1} A_k \end{pmatrix} \begin{pmatrix} e_k \\ \hat{e}_k^+ \end{pmatrix} \\ &+ \begin{pmatrix} G_k & 0 \\ K_{k+1} H_{k+1} G_k & K_{k+1} M_{k+1} \end{pmatrix} \begin{pmatrix} w_k \\ v_{k+1} \end{pmatrix} \end{aligned} \quad (87)$$

Defining $\zeta_k := (e_k, \hat{e}_k^+)^T$ and $q_k := (w_k, v_{k+1})^T$, we can rewrite Equation (87) in a more compact form as

$$\begin{aligned} \zeta_{k+1} &= \bar{F}_k \zeta_k + \bar{G}_k q_k, \quad q_k \sim \mathcal{N}(0, \bar{Q}_k), \\ \bar{Q}_k &= \begin{pmatrix} Q_k & 0 \\ 0 & R_{k+1} \end{pmatrix} \end{aligned} \quad (88)$$

with appropriate definitions for \bar{F}_k and \bar{G}_k .

The above equation, along with the equation on estimation covariance propagation,

$$P_{k+1} = (I - K_{k+1}H_{k+1})(A_k P_k A_k^T + G_k Q_k G_k^T) \quad (89)$$

characterize the evolution of state x_k and belief $b_k \equiv (\hat{x}_k^+, P_k)$ under the time-varying LQG controller.

C. SLQG controller

The SLQG controller is often used to regulate (or stabilize) the system state to a pre-planned point (also called the set-point, nominal, or desired point) in the presence of process and observation noise. In principal it is designed (and optimal) for linear systems with Gaussian noise, but it can also be utilized for stabilizing nonlinear systems locally around the planned point. The SLQG controller is composed of a stationary Kalman filter (SKF) as an estimator and a stationary linear quadratic regulator (SLQR) as a controller. At every time step k , the SKF provides the *a posteriori* distribution (belief) b_k over the system state, and the SLQR generates control u_k based on b_k .

In this appendix, we first discuss the system linearization around the planned point, and then discuss the SKF, SLQR, and SLQG corresponding to this nominal point. Consider the nonlinear partially observable state-space equations of the system as follows:

$$x_{k+1} = f(x_k, u_k, w_k), \quad w_k \sim \mathcal{N}(0, Q_k) \quad (90a)$$

$$z_k = h(x_k, v_k), \quad v_k \sim \mathcal{N}(0, R_k) \quad (90b)$$

and consider a planned state point x^p , to whose vicinity the controller has to drive the system. If the system state reaches the point x^p , it is assumed that the system remains there with zero control, $u^p = 0$, in other words,

$$x^p = f(x^p, 0, 0) \quad (91)$$

The role of a closed-loop stochastic controller during the state regulation is to compensate for robot deviations from the desired point due to noise effects, and to drive the robot close to the desired point in the sense of minimizing the following quadratic cost:

$$J = \mathbb{E} \left[\sum_{k \geq 0} (x_k - x^p)^T W_x (x_k - x^p) + (u_k)^T W_u (u_k) \right] \quad (92)$$

where W_x and W_u are positive-definite weight matrices for the state and control costs, respectively.

Again, similar to the time-varying case, since we only have imperfect information about the state x_k , we have to make the estimate x_k^+ about the state based on the available observations $z_{0:k}$. Accordingly, the controller generates the control signal based on the estimated value of the state; i.e., belief. Based on the separation principle (Bertsekas, 2007), in a linear system with Gaussian noise, minimization of the cost in Equation (92) is equivalent to performing two separate minimizations that lead to the separate design of the SKF and SLQR. In the following, we first discuss the linearization of a nonlinear model, and then we discuss how the SKF and the SLQR can be designed for this linearized system, and finally, we combine them to construct an SLQG controller.

Model linearization: Given a desired point x^p , we linearize the dynamics and observation model in Equation (90) as follows:

$$\begin{aligned} x_{k+1} &= f(x^p, 0, 0) + A_s(x_k - x^p) + B_s(u_k - 0) \\ &\quad + G_s w_k, \quad w_k \sim \mathcal{N}(0, Q_s) \end{aligned} \quad (93a)$$

$$z_k = h(x^p, 0) + H_s(x_k - x^p) + M_s v_k, \quad v_k \sim \mathcal{N}(0, R_s) \quad (93b)$$

where

$$\begin{aligned} A_s &= \frac{\partial f}{\partial x}(x^p, 0, 0), \quad B_s = \frac{\partial f}{\partial u}(x^p, 0, 0), \quad G_s = \frac{\partial f}{\partial w}(x^p, 0, 0) \\ H_s &= \frac{\partial h}{\partial x}(x^p, 0), \quad M_s = \frac{\partial h}{\partial v}(x^p, 0) \end{aligned} \quad (94)$$

Now, let us define the following errors:

- SLQG error (main error): $e_k = x_k - x^p$
- SKF error (estimation error): $\tilde{e}_k = x_k - \hat{x}_k^+$, where $\hat{x}_k^+ = \mathbb{E}[x_k^+] = \mathbb{E}[x_k | z_{0:k}]$
- SLQR error (estimation of SLQG error): $\hat{e}_k^+ = \hat{x}_k^+ - x^p$

Note that these errors are linearly dependent: $e_k = \hat{e}_k^+ + \tilde{e}_k$. Defining $\delta u_k := u_k$ and $\delta z_k := z_k - z^p = z_k - h(x^p, 0)$, we can rewrite the above linearized models as follows:

$$e_{k+1} = A_s e_k + B_s \delta u_k + G_s w_k \quad (95a)$$

$$\delta z_k = H_s e_k + M_s v_k \quad (95b)$$

SKF: In SKF, we aim to provide an estimate of the system's state based on the available partial information we have obtained up to time k , that is, $z_{0:k}$. The state estimate is a random vector denoted by x_k^+ , whose distribution is the conditional distribution of the state on the obtained observations so far, which is called belief and is denoted by $b_k = p(x_k^+) = p(x_k | z_{0:k})$. In the Gaussian case, the belief can only be characterized by its mean and covariance, that is, $b_k \equiv (\hat{x}_k^+, P_k)$. Thus, in the Gaussian case, we can write

$$b_k = p(x_k^+) = p(x_k | z_{0:k}) = \mathcal{N}(\hat{x}_k^+, P_k) \Leftrightarrow b_k \equiv (\hat{x}_k^+, P_k) \quad (96)$$

$$\hat{x}_k^+ = \mathbb{E}[x_k | z_{0:k}], \quad P_k = \mathbb{C}[x_k | z_{0:k}] \quad (97)$$

where $\mathbb{E}[\cdot|\cdot]$ and $\mathbb{C}[\cdot|\cdot]$ are the conditional expectation and conditional covariance operators, respectively.

SKF consists of two steps at every time stage: a prediction step and an update step. In the prediction step, the mean and covariance of prior x_k^- are computed. For the system in Equation (95) the prediction step is

$$\widehat{e}_{k+1}^- = A_s \widehat{e}_k^+ + B_s \delta u_k \quad (98)$$

$$P_{k+1}^- = A_s P_k^+ A_s^T + G_s Q_s G_s^T \quad (99)$$

In the update step, the mean and covariance of posterior x_k^+ are computed. For the error system in Equation (95), the update step is

$$K_k = P_k^- H_s^T (H_s P_k^- H_s^T + M_s R_s M_s^T)^{-1} \quad (100)$$

$$\widehat{e}_{k+1}^+ = \widehat{e}_{k+1}^- + K_{k+1} (\delta z_{k+1} - H_s \widehat{e}_{k+1}^-) \quad (101)$$

$$P_{k+1}^+ = (I - K_{k+1} H_s) P_{k+1}^- \quad (102)$$

Note that

$$\widehat{x}_k^+ = x^p + \widehat{e}_k^+, \quad P_k = P_k^+ \quad (103)$$

In SKF, if (A_s, H_s) is an observable pair and (A_s, \check{Q}_s) is a controllable pair, where $G_s Q_s G_s^T = \check{Q}_s \check{Q}_s^T$, then the prior and posterior covariances P_k^- and P_k and the filter gain K_k all converge to their stationary values, denoted by P_s^-, P_s , and K_s , respectively (Bertsekas, 2007). P_s^- can be computed by solving the following DARE in Equation (104). Having P_s^- , the stationary gain K_s and estimation covariance P_s are computed as follows:

$$P_s^- = G_s Q_s G_s^T + A_s (P_s^- - P_s^- H_s^T (H_s P_s^- H_s^T + M_s R_s M_s^T)^{-1} H_s P_s^-) A_s^T \quad (104)$$

$$K_s = P_s^- H_s^T (H_s P_s^- H_s^T + M_s R_s M_s^T)^{-1} \quad (105)$$

$$P_s = (I - K_s H_s) P_s^- \quad (106)$$

SLQR controller: In the SLQR we have a stationary mapping μ_s from the belief space to the control space that generates an optimal control based on the given belief $u_k = \mu_s(b_k)$ at every time step k . The SLQR controller is optimal in the sense of minimizing the following cost:

$$J_{SLQR} = \mathbb{E} \left[\sum_{k \geq 0} (\widehat{x}_k^+ - x^p)^T W_x (\widehat{x}_k^+ - x^p) + (u_k)^T W_u (u_k) \right] \\ = \mathbb{E} \left[\sum_{k \geq 0} (\widehat{e}_k^+)^T W_x (\widehat{e}_k^+) + (\delta u_k)^T W_u (\delta u_k) \right] \quad (107)$$

If (A_s, B_s) is a controllable pair and (A_s, \check{W}_x) is an observable pair, where $\check{W}_x^T \check{W}_x = W_x$, then the stationary linear control law that minimizes the cost function J_{SLQR} for a linear system is of the form

$$\delta u_k = -L_s \widehat{e}_k^+ \quad (108)$$

where the stationary feedback gain L_s can be computed as follows:

$$L_s = (B_s^T S_s B_s + W_u)^{-1} B_s^T S_s A_s \quad (109)$$

$$S_s = W_x + A_s^T S_s A_s - A_s^T S_s B_s L_s \quad (110)$$

where the second equation is indeed a DARE that can be efficiently solved for S_s . Plugging S_s into Equation (109), we get the feedback gain L_s .

SLQG controller: Plugging the obtained control law of SLQR into the SKF equations, we can get the following stationary dynamics for the defined errors:

$$\begin{pmatrix} e_{k+1} \\ \tilde{e}_{k+1} \end{pmatrix} = \begin{pmatrix} A_s - B_s L_s & B_s L_s \\ 0 & A_s - K_s H_s A_s \end{pmatrix} \begin{pmatrix} e_k \\ \tilde{e}_k \end{pmatrix} + \begin{pmatrix} G_s & 0 \\ G_s - K_s H_s G_s & -K_s M_s \end{pmatrix} \begin{pmatrix} w_k \\ v_{k+1} \end{pmatrix} \quad (111)$$

or equivalently,

$$\begin{pmatrix} e_{k+1} \\ \widehat{e}_{k+1}^+ \end{pmatrix} = \begin{pmatrix} A_s & -B_s L_s \\ K_s H_s A_s & A_s - B_s L_s - K_s H_s A_s \end{pmatrix} \begin{pmatrix} e_k \\ \widehat{e}_k^+ \end{pmatrix} + \begin{pmatrix} G_s & 0 \\ K_s H_s G_s & K_s M_s \end{pmatrix} \begin{pmatrix} w_k \\ v_{k+1} \end{pmatrix} \quad (112)$$

Defining $\zeta_k := (e_k, \widehat{e}_k^+)^T$ and $q_k := (w_k, v_{k+1})^T$, we can rewrite Equation (112) in a more compact form as

$$\zeta_{k+1} = \bar{F}_s \zeta_k + \bar{G}_s q_k, \quad q_k \sim \mathcal{N}(0, \bar{Q}_s), \quad \bar{Q}_s = \begin{pmatrix} Q_s & 0 \\ 0 & R_s \end{pmatrix} \quad (113)$$

with appropriate definitions for \bar{F}_s and \bar{G}_s .

It can be shown that if \bar{F}_s is a stable matrix (i.e. $\lim_{k \rightarrow \infty} (\bar{F}_s)^k = 0$), ζ_k converges i.d. to $\zeta_s \sim \mathcal{N}(0, \mathcal{P}_s)$. Stationary covariance \mathcal{P}_s is the solution of the following Lyapunov equation:

$$\mathcal{P}_s = \bar{F}_s \mathcal{P}_s \bar{F}_s^T + \bar{G}_s \bar{Q}_s \bar{G}_s^T \quad (114)$$

Note that \mathcal{P}_s can be decomposed into four blocks,

$$\mathcal{P}_s = \begin{pmatrix} \mathcal{P}_{s,11} & \mathcal{P}_{s,12} \\ \mathcal{P}_{s,21} & \mathcal{P}_{s,22} \end{pmatrix} \quad (115)$$

such that $\mathcal{P}_{s,11} = \lim_{k \rightarrow \infty} \mathbb{C}[e_k]$ and $\mathcal{P}_{s,22} = \lim_{k \rightarrow \infty} \mathbb{C}[\widehat{e}_k^+]$. Therefore, since $\widehat{x}_k^+ = x^p + \widehat{e}_k^+$, the estimation mean also converges to a stationary random variable, denoted by \widehat{x}_s^+ :

$$\widehat{x}_s^+ := \lim_{k \rightarrow \infty} \widehat{x}_k^+ \sim \mathcal{N}(x^p, \mathcal{P}_{s,22}) \quad (116)$$

Due to the linear relation $e_k = \widehat{e}_k^+ + \tilde{e}_k$, we can also conclude $\lim_{k \rightarrow \infty} \mathbb{C}[\tilde{e}_k] = \mathcal{P}_{s,11} + \mathcal{P}_{s,22} - 2\mathcal{P}_{s,12}$. It can be proven that in SLQG, the stability of matrix \bar{F}_s is a direct consequence of the controllability of pair (A_s, B_s) and the observability of pair (A_s, H_s) (Bertsekas, 1976, 2007).

Thus, collecting all the conditions, if (A_s, B_s) and (A_s, \check{Q}_s) are controllable pairs, where $G_s Q_s G_s^T = \check{Q}_s \check{Q}_s^T$, and if (A_s, H_s) and (A_s, \check{W}_x) are observable pairs, where $W_x = \check{W}_x^T \check{W}_x$, then the belief b_k converges in distribution to a stationary belief under the SLQG:

$$b_s := \lim_{k \rightarrow \infty} b_k = \mathcal{N}(\hat{x}_s^+, P_s^+) \quad (117)$$

where P_s^+ is a deterministic quantity and we can characterize the distribution over the stationary belief as

$$b_s \equiv (\hat{x}_s^+, P_s^+) \sim \mathcal{N} \left(\begin{pmatrix} x^p \\ P_s^+ \end{pmatrix}, \begin{pmatrix} P_{s,22} & 0 \\ 0 & 0 \end{pmatrix} \right) \quad (118)$$

D. Proof of Lemma 3

Proof. Let us consider the state-space model of the linear system of interest as follows:

$$x_{k+1} = \mathbf{A}x_k + \mathbf{B}u_k + \mathbf{G}w_k, \quad w_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}) \quad (119a)$$

$$z_k = \mathbf{H}x_k + v_k, \quad v_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}) \quad (119b)$$

Based on Lemma 1, if (\mathbf{A}, \mathbf{B}) and $(\mathbf{A}, \check{\mathbf{Q}})$ are controllable pairs, where $\mathbf{G}\mathbf{Q}\mathbf{G}^T = \check{\mathbf{Q}}\check{\mathbf{Q}}^T$, and if (\mathbf{A}, \mathbf{H}) and $(\mathbf{A}, \check{\mathbf{W}}_x)$ are observable pairs, where $\mathbf{W}_x = \check{\mathbf{W}}_x^T \check{\mathbf{W}}_x$, then the estimation covariance deterministically tends to a stationary covariance P_s . Therefore, for any $\epsilon > 0$, after a deterministic finite time, P_k enters the ϵ -neighborhood of the stationary covariance, denoted by P_s .

The estimation mean dynamics, however, are stochastic and are as follows for the system in Equation (119):

$$\begin{aligned} \hat{x}_{k+1}^+ &= \mathbf{v} + (\mathbf{A} - \mathbf{B}\mathbf{L} - \mathbf{K}_{k+1}\mathbf{H}\mathbf{A})(\hat{x}_k^+ - \mathbf{v}) \\ &\quad + \mathbf{K}_{k+1}\mathbf{H}\mathbf{A}(x_k - \mathbf{v}) + \mathbf{K}_{k+1}\mathbf{H}\mathbf{G}w_k + \mathbf{K}_{k+1}v_{k+1} \\ &= \mathbf{v} - (\mathbf{A} - \mathbf{B}\mathbf{L})\mathbf{v} + (\mathbf{A} - \mathbf{B}\mathbf{L} - \mathbf{K}_{k+1}\mathbf{H}\mathbf{A})\hat{x}_k^+ \\ &\quad + \mathbf{K}_{k+1}\mathbf{H}\mathbf{A}x_k + \mathbf{K}_{k+1}\mathbf{H}\mathbf{G}w_k + \mathbf{K}_{k+1}v_{k+1} \end{aligned} \quad (120)$$

where the Kalman gain \mathbf{K}_k is

$$\mathbf{K}_k = P_k^- \mathbf{H}^T (\mathbf{H}P_k^- \mathbf{H}^T + \mathbf{R})^{-1} \quad (121)$$

Since \mathbf{K} is full rank (due to the condition on the rank of \mathbf{H}), and since v and w represent Gaussian noise, Equation (120) induces an irreducible Markov process over the state space (Meyn and Tweedie, 2009). Thus, if we have a stopping region for the estimation mean with size $\epsilon > 0$, the estimation mean process will hit this stopping region in finite time (Meyn and Tweedie, 2009), with probability one.

Based on the estimation mean dynamics in Equation (120) and the state dynamics in Appendix C, in the absence of a stopping region, if the estimation mean process and state process start from \hat{x}_0^+ and x_0 respectively, such that $\mathbb{E}[\hat{x}_0^+] = \mathbf{v}$ and $\mathbb{E}[x_0] = \mathbf{v}$ (which indeed is the case in FIRM due to the usage of edge controllers), ‘the mean of estimation mean’ remains on \mathbf{v} . That is, $\mathbb{E}[\hat{x}_k^+] = \mathbf{v}$, for all k . As a result, if we center the stopping region for the

estimation mean at \mathbf{v} , the probability of hitting the stopping region is maximized and the stopping time is minimized.

Combining the results for estimation covariance and estimation mean, if we define the region B as a set in the Gaussian belief space with a non-empty interior centered at (\mathbf{v}, P_s) , the belief $b_k \equiv (\hat{x}_k^+, P_k)$ enters region B in finite time with probability one. Thus, the pair (B, μ) is a proper pair over $\mathbb{G}\mathbb{B}$; in other words, B is reachable under μ starting from any Gaussian distribution. \square

E. Proof of Lemma 4

Before proving Lemma 4, we state and prove the following lemma.

Lemma 5. Consider the bounded function $0 \leq f(\mathcal{X}) \leq 1$, and kernel $k(\mathcal{X}', \mathcal{X}) \geq 0$. Then, for any set \mathcal{A} , we have

$$\left\| \int_{\mathcal{A}} k(\mathcal{X}', \mathcal{X}) f(\mathcal{X}') d\mathcal{X}' \right\| \leq \left\| \int_{\mathcal{A}} k(\mathcal{X}', \mathcal{X}) d\mathcal{X}' \right\| \quad (122)$$

Proof. Given the properties of $f(\cdot)$ and $k(\cdot, \cdot)$, we have $k(\mathcal{X}', \mathcal{X}) f(\mathcal{X}') \leq k(\mathcal{X}', \mathcal{X})$, for all \mathcal{X} and \mathcal{X}' . Taking the integral from both sides with respect to \mathcal{X}' and then taking the supremum norm with respect to \mathcal{X} , the result follows. \square

Now we prove Lemma 4.

Proof. If we denote the domain of operator \mathbf{T}_S by \mathcal{D} , we know that for all $f \in \mathcal{D}$ we have $0 \leq f(\mathcal{X}) \leq 1$, because $f(\mathcal{X})$ is the probability of reaching given set \mathcal{S} under some given controller invoked at point \mathcal{X} . Thus, it cannot be negative or greater than one, and based on Lemma 5, we have

$$\begin{aligned} \mathbf{T}_S[f] &= \int_{\bar{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) f(\mathcal{X}') d\mathcal{X}' \\ &\leq \int_{\bar{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) d\mathcal{X}' = \mathbb{P}_1(\bar{\mathcal{S}}|\mathcal{X}, \mu) \leq 1 \end{aligned} \quad (123)$$

Therefore, based on the definition of operator norm, we have

$$\|\mathbf{T}_S\|_{op} = \sup_{f(\cdot)} \{\|\mathbf{T}_S[f]\| : \forall f \in \mathcal{D}, \|f\| \leq 1\} \leq 1 \quad (124)$$

According to Assumption 1, there exists a finite number N such that

$$\inf_{\mathcal{X}} \mathbb{P}_n(\mathcal{S}|\mathcal{X}, \mu) = \beta > 0 \quad \forall n > N \quad (125)$$

where ‘inf’ and ‘sup’ denote the infimum and supremum, respectively. Thus, we have

$$\begin{aligned} \|\mathbb{P}_n(\bar{\mathcal{S}}|\mathcal{X}, \mu)\| &= \sup_{\mathcal{X}} (1 - \mathbb{P}_n(\mathcal{S}|\mathcal{X}, \mu)) \\ &= 1 - \inf_{\mathcal{X}} \mathbb{P}_n(\mathcal{S}|\mathcal{X}, \mu) \\ &= 1 - \beta < 1 \quad \forall n > N \end{aligned} \quad (126)$$

Let us denote the n th iterated kernel of \mathbf{T}_S as $p_n(\mathcal{X}'|\mathcal{X}, \mu)$. Since this iterated kernel is a pdf, we have $p_n(\mathcal{X}'|\mathcal{X}, \mu) \geq 0, \forall \mathcal{X}, \forall \mathcal{X}', \forall n$. We can write

$$\begin{aligned} \|\mathbf{T}_S^N[f]\| &= \left\| \int_{\bar{\mathcal{S}}} p_N(\mathcal{X}'|\mathcal{X}, \mu) f(\mathcal{X}') d\mathcal{X}' \right\| \\ &\leq \left\| \int_{\bar{\mathcal{S}}} p_N(\mathcal{X}'|\mathcal{X}, \mu) d\mathcal{X}' \right\| = \|\mathbb{P}_N(\bar{\mathcal{S}}|\mathcal{X}, \mu)\| \leq \alpha < 1 \end{aligned} \quad (127)$$

where $\alpha = 1 - \beta$, and similar to Equation (124), we get $\|\mathbf{T}_S^N\|_{op} \leq \alpha < 1$. From the operator norm properties, we have

$$\|\mathbf{T}_S^{N+1}\|_{op} \leq \|\mathbf{T}_S^N\|_{op} \|\mathbf{T}_S\|_{op} \leq \alpha < 1$$

and similarly for all $n \geq N$, we have

$$\|\mathbf{T}_S^n\|_{op} \leq \alpha < 1 \quad \forall n \geq N$$

Now, consider the series: $\sum_{i=1}^{\infty} \|\mathbf{T}_S^i\|_{op}$. We can split the sum into smaller pieces as follows:

$$\sum_{n=1}^{\infty} \|\mathbf{T}_S^n\|_{op} = \sum_{n=1}^N \|\mathbf{T}_S^n\|_{op} + \sum_{i=1}^{\infty} \sum_{n=iN+1}^{(i+1)N} \|\mathbf{T}_S^n\|_{op}$$

But because $\|\mathbf{T}_S^{n+1}\|_{op} \leq \|\mathbf{T}_S^n\|_{op}$ for all $n \geq N$, we have

$$\sum_{n=iN+1}^{(i+1)N} \|\mathbf{T}_S^n\|_{op} \leq N \|\mathbf{T}_S^{iN}\|_{op}$$

Also, we know

$$\|\mathbf{T}_S^{iN}\|_{op} \leq \|\mathbf{T}_S^N\|_{op}^i \leq \alpha^i$$

and thus, we have

$$\begin{aligned} \sum_{n=1}^{\infty} \|\mathbf{T}_S^n\|_{op} &= \sum_{n=1}^N \|\mathbf{T}_S^n\|_{op} + \sum_{i=1}^{\infty} \sum_{n=iN+1}^{(i+1)N} \|\mathbf{T}_S^n\|_{op} \\ &\leq N + \sum_{i=1}^{\infty} N\alpha^i = N + \frac{N}{1-\alpha} = c < \infty \end{aligned} \quad \square$$

F. Proof of Corollary 1

Proof. We know $\|R\| \leq 1$, and thus we can write

$$\left\| \sum_{n=0}^{\infty} \mathbf{T}_S^n[R] \right\| \leq \sum_{n=0}^{\infty} \|\mathbf{T}_S^n\|_{op} \|R\| \leq \sum_{n=0}^{\infty} \|\mathbf{T}_S^n\|_{op} \leq c < \infty$$

Thus, series $\sum_{n=0}^{\infty} \mathbf{T}_S^n[R]$ is a convergent series and we can define the operator $(I - \mathbf{T}_S)^{-1}[R] = \sum_{n=0}^{\infty} \mathbf{T}_S^n[R]$. We have

$$\|(I - \mathbf{T}_S)^{-1}\|_{op} = \left\| \sum_{n=0}^{\infty} \mathbf{T}_S^n \right\|_{op} \leq c < \infty \quad (128)$$

G. Proof of Proposition 1

We first state and prove the following lemma on the continuity of the transition probability in the local controller's parameter.

Lemma 6. *Given Assumption 2, there exists a $c_2 < \infty$ such that*

$$\|p(\mathcal{X}'|\mathcal{X}, \mu(b; \mathbf{v})) - p(\mathcal{X}'|\mathcal{X}, \check{\mu}(b; \check{\mathbf{v}}))\| \leq c_2 \|\mathbf{v} - \check{\mathbf{v}}\| \quad (129)$$

Proof. The result directly follows by combining the two parts of Assumption 2. \square

Now we are ready to prove Proposition 1.

Proof. To show $\mathbb{P}(\mathcal{B}|\mathcal{X}, \mu)$ is continuous in \mathbf{v} , we perturb \mathbf{v} to some $\check{\mathbf{v}}$, such that $\|\mathbf{v} - \check{\mathbf{v}}\| < r$. The local controller associated with node $\check{\mathbf{v}}$ is referred to as $\check{\mu}$, whose successful absorption region is denoted by $\check{\mathcal{B}}$ and whose stopping region is $\check{\mathcal{S}}$. Similarly, the corresponding transient operator and recurrent function are referred to as $\check{\mathbf{T}}_{\mathcal{S}}$ and \check{R} respectively. Finally, the success probability associated with the perturbed node $\check{\mathbf{v}}$ is $\mathbb{P}(\check{\mathcal{B}}|\mathcal{X}, \check{\mu})$. To shorten the statements, we refer to $\mathbb{P}(\mathcal{B}|\mathcal{X}, \mu)$ and $\mathbb{P}(\check{\mathcal{B}}|\mathcal{X}, \check{\mu})$ respectively as $\mathfrak{P}(\mathcal{X})$ and $\check{\mathfrak{P}}(\mathcal{X})$. As a result of node perturbation, the success probability is perturbed as

$$\begin{aligned} \mathbb{P}(\mathcal{B}|\mathcal{X}, \mu) - \mathbb{P}(\check{\mathcal{B}}|\mathcal{X}, \check{\mu}) &:= \mathfrak{P} - \check{\mathfrak{P}} = R + \mathbf{T}_S[\mathfrak{P}] - \check{R} - \check{\mathbf{T}}_{\mathcal{S}}[\check{\mathfrak{P}}] \\ &= R - \check{R} + \mathbf{T}_S[\mathfrak{P}] - \mathbf{T}_S[\check{\mathfrak{P}}] + \mathbf{T}_S[\check{\mathfrak{P}}] - \check{\mathbf{T}}_{\mathcal{S}}[\check{\mathfrak{P}}] + \check{\mathbf{T}}_{\mathcal{S}}[\check{\mathfrak{P}}] - \check{\mathbf{T}}_{\mathcal{S}}[\check{\mathfrak{P}}] \\ &= (R - \check{R}) + \mathbf{T}_S[\mathfrak{P} - \check{\mathfrak{P}}] + (\mathbf{T}_S - \check{\mathbf{T}}_{\mathcal{S}})[\check{\mathfrak{P}}] + (\check{\mathbf{T}}_{\mathcal{S}} - \mathbf{T}_S)[\check{\mathfrak{P}}] \end{aligned}$$

where

$$\mathbf{T}_{\check{\mathcal{S}}}[f(\cdot)](\mathcal{X}) := \int_{\check{\mathcal{S}}} p^{\mu}(\mathcal{X}'|\mathcal{X}) f(\mathcal{X}') d\mathcal{X}' \quad (130)$$

Let us define the operators $\mathbf{T}_{\Delta S} := (\mathbf{T}_S - \mathbf{T}_{\check{\mathcal{S}}})$ and $\Delta \mathbf{T}_{\check{\mathcal{S}}} := (\mathbf{T}_{\check{\mathcal{S}}} - \check{\mathbf{T}}_{\mathcal{S}})$. Now, based on Corollary 1, we can write

$$\mathfrak{P} - \check{\mathfrak{P}} = (I - \mathbf{T}_S)^{-1} \left[R - \check{R} + \mathbf{T}_{\Delta S}[\check{\mathfrak{P}}] + \Delta \mathbf{T}_{\check{\mathcal{S}}}[\check{\mathfrak{P}}] \right] \quad (131)$$

and thus the following inequality holds on the supremum norm of the perturbation of the absorption probability:

$$\begin{aligned} \|\mathfrak{P} - \check{\mathfrak{P}}\| &\leq \|(I - \mathbf{T}_S)^{-1}\|_{op} \left(\|R - \check{R}\| + \|\mathbf{T}_{\Delta S}[\check{\mathfrak{P}}]\| + \|\Delta \mathbf{T}_{\check{\mathcal{S}}}[\check{\mathfrak{P}}]\| \right) \\ &\leq c \left(\|R - \check{R}\| + \|\mathbf{T}_{\Delta S}[\check{\mathfrak{P}}]\| + \|\Delta \mathbf{T}_{\check{\mathcal{S}}}[\check{\mathfrak{P}}]\| \right) \\ &= c (\|K_1(\mathcal{X})\| + \|K_2(\mathcal{X})\| + \|K_3(\mathcal{X})\|) \end{aligned} \quad (132)$$

where $K_1(\mathcal{X}) := R(\mathcal{X}) - \check{R}(\mathcal{X})$, $K_2(\mathcal{X}) := \mathbf{T}_{\Delta S}[\check{\mathfrak{P}}(\cdot)](\mathcal{X})$, and $K_3(\mathcal{X}) := \Delta \mathbf{T}_{\check{\mathcal{S}}}[\check{\mathfrak{P}}(\cdot)](\mathcal{X})$. In the following we bound K_1 , K_2 , and K_3 , and thus bound $\|\mathfrak{P} - \check{\mathfrak{P}}\|$, accordingly. \square

G.1. Bound for $K_1(\mathcal{X})$

The supremum norm of $K_1(\mathcal{X})$ is

$$\begin{aligned}
\|K_1(\mathcal{X})\| &= \|R(\mathcal{X}) - \check{R}(\mathcal{X})\| \\
&= \left\| \int_{\mathcal{B}} p^\mu(\mathcal{X}'|\mathcal{X}) d\mathcal{X}' - \int_{\check{\mathcal{B}}} p^{\check{\mu}}(\mathcal{X}'|\mathcal{X}) d\mathcal{X}' \right\| \\
&= \left\| \int_{\mathcal{B} \cap \check{\mathcal{B}}} [p^\mu(\mathcal{X}'|\mathcal{X}) - p^{\check{\mu}}(\mathcal{X}'|\mathcal{X})] d\mathcal{X}' \right. \\
&\quad \left. + \int_{\mathcal{B} - \check{\mathcal{B}}} p^\mu(\mathcal{X}'|\mathcal{X}) d\mathcal{X}' - \int_{\check{\mathcal{B}} - \mathcal{B}} p^{\check{\mu}}(\mathcal{X}'|\mathcal{X}) d\mathcal{X}' \right\| \\
&\leq \int_{\mathcal{B} \cap \check{\mathcal{B}}} \|p^\mu(\mathcal{X}'|\mathcal{X}) - p^{\check{\mu}}(\mathcal{X}'|\mathcal{X})\| d\mathcal{X}' \\
&\quad + \left\| \int_{\mathcal{B} - \check{\mathcal{B}}} p^\mu(\mathcal{X}'|\mathcal{X}) d\mathcal{X}' + \int_{\check{\mathcal{B}} - \mathcal{B}} p^{\check{\mu}}(\mathcal{X}'|\mathcal{X}) d\mathcal{X}' \right\| \\
&\stackrel{\text{from (129)}}{\leq} \int_{\mathcal{B} \cap \check{\mathcal{B}}} c_2 \|\mathbf{v} - \check{\mathbf{v}}\| d\mathcal{X}' + \|\mathbb{P}_1(\mathcal{B} \ominus \check{\mathcal{B}}|\mathcal{X}, \mu)\| \\
&\quad + \|\mathbb{P}_1(\check{\mathcal{B}} \ominus \mathcal{B}|\mathcal{X}, \check{\mu})\| \\
&\stackrel{\text{from (50)}}{\leq} c'_2 \|\mathbf{v} - \check{\mathbf{v}}\| + 2c' \|\mathbf{v} - \check{\mathbf{v}}\| = \gamma_1 \|\mathbf{v} - \check{\mathbf{v}}\| \quad (133)
\end{aligned}$$

where $c'_2 < \infty$ and $\gamma_1 = c'_2 + 2c' < \infty$. In the penultimate inequality, we also used the fact that $\mathbb{P}_1(\check{\mathcal{B}} - \mathcal{B}|\mathcal{X}, \check{\mu}) \leq \mathbb{P}_1(\check{\mathcal{B}} \ominus \mathcal{B}|\mathcal{X}, \check{\mu})$ and $\mathbb{P}_1(\mathcal{B} - \check{\mathcal{B}}|\mathcal{X}, \mu) \leq \mathbb{P}_1(\mathcal{B} \ominus \check{\mathcal{B}}|\mathcal{X}, \mu)$ because $\check{\mathcal{B}} - \mathcal{B} \subseteq \check{\mathcal{B}} \ominus \mathcal{B}$ and $\mathcal{B} - \check{\mathcal{B}} \subseteq \mathcal{B} \ominus \check{\mathcal{B}}$.

G.2. Bound for $K_2(\mathcal{X})$

We have

$$\begin{aligned}
\|K_2(\mathcal{X})\| &= \|\mathbf{T}_{\Delta\mathcal{S}}[\check{\mathfrak{P}}]\| = \|\mathbf{T}_{\mathcal{S}}[\check{\mathfrak{P}}] - \mathbf{T}_{\check{\mathcal{S}}}[\check{\mathfrak{P}}]\| \\
&= \left\| \int_{\check{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) \check{\mathfrak{P}}(\mathcal{X}') d\mathcal{X}' - \int_{\check{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) \check{\mathfrak{P}}(\mathcal{X}') d\mathcal{X}' \right\| \\
&= \left\| \int_{\check{\mathcal{S}} - \check{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) \check{\mathfrak{P}}(\mathcal{X}') d\mathcal{X}' - \int_{\check{\mathcal{S}} - \check{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) \check{\mathfrak{P}}(\mathcal{X}') d\mathcal{X}' \right\| \\
&\leq \left\| \int_{\check{\mathcal{S}} - \check{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) \check{\mathfrak{P}}(\mathcal{X}') d\mathcal{X}' + \int_{\check{\mathcal{S}} - \check{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) \check{\mathfrak{P}}(\mathcal{X}') d\mathcal{X}' \right\| \\
&= \left\| \int_{\check{\mathcal{S}} \ominus \check{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) \check{\mathfrak{P}}(\mathcal{X}') d\mathcal{X}' \right\| \stackrel{\text{from (122)}}{\leq} \left\| \int_{\check{\mathcal{S}} \ominus \check{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) d\mathcal{X}' \right\| \\
&= \|\mathbb{P}_1(\check{\mathcal{S}} \ominus \check{\mathcal{S}}|\mathcal{X}, \mu)\| \leq \|\mathbb{P}_1(\check{\mathcal{B}} \ominus \check{\mathcal{B}}|\mathcal{X}, \mu)\| \\
&= \|\mathbb{P}_1(\mathcal{B} \ominus \check{\mathcal{B}}|\mathcal{X}, \mu)\| \stackrel{\text{from (50)}}{\leq} \gamma_2 \|\mathbf{v} - \check{\mathbf{v}}\| \quad (134)
\end{aligned}$$

where $\gamma_2 = c' < \infty$. The penultimate inequality and equality follow from the relations $\check{\mathcal{S}} \ominus \check{\mathcal{S}} \subseteq \check{\mathcal{B}} \ominus \check{\mathcal{B}}$ and $\check{\mathcal{B}} \ominus \check{\mathcal{B}} = \mathcal{B} \ominus \check{\mathcal{B}}$, respectively.

G.3. Bound for $K_3(\mathcal{X})$

We have

$$\begin{aligned}
\|K_3(\mathcal{X})\| &= \|\Delta\mathbf{T}_{\mathcal{S}}[\check{\mathfrak{P}}]\| = \|\mathbf{T}_{\mathcal{S}}[\check{\mathfrak{P}}] - \check{\mathbf{T}}_{\check{\mathcal{S}}}[\check{\mathfrak{P}}]\| \\
&= \left\| \int_{\check{\mathcal{S}}} p^\mu(\mathcal{X}'|\mathcal{X}) \check{\mathfrak{P}}(\mathcal{X}') d\mathcal{X}' - \int_{\check{\mathcal{S}}} p^{\check{\mu}}(\mathcal{X}'|\mathcal{X}) \check{\mathfrak{P}}(\mathcal{X}') d\mathcal{X}' \right\| \\
&= \left\| \int_{\check{\mathcal{S}}} (p^\mu(\mathcal{X}'|\mathcal{X}) - p^{\check{\mu}}(\mathcal{X}'|\mathcal{X})) \check{\mathfrak{P}}(\mathcal{X}') d\mathcal{X}' \right\| \\
&\leq \int_{\check{\mathcal{S}}} \|p^\mu(\mathcal{X}'|\mathcal{X}) - p^{\check{\mu}}(\mathcal{X}'|\mathcal{X})\| \|\check{\mathfrak{P}}(\mathcal{X}')\| d\mathcal{X}' \\
&\stackrel{\text{from (129)}}{\leq} \int_{\check{\mathcal{S}}} c_2 \|\mathbf{v} - \check{\mathbf{v}}\| d\mathcal{X}' = \gamma_3 \|\mathbf{v} - \check{\mathbf{v}}\| \quad (135)
\end{aligned}$$

where $\gamma_3 < \infty$.

Therefore, based on Equations (132)–(135), we can conclude that

$$\|\mathbb{P}(\mathcal{B}|\mathcal{X}, \mu) - \mathbb{P}(\check{\mathcal{B}}|\mathcal{X}, \check{\mu})\| \leq \gamma \|\mathbf{v} - \check{\mathbf{v}}\| \quad (136)$$

where $\gamma = c(\gamma_1 + \gamma_2 + \gamma_3) < \infty$, which completes the proof that the absorption probability under the controller μ is continuous in the PRM node \mathbf{v} . \square

H. Proof of Theorem 1

Before starting the proof of Theorem 1, we state the following proposition that concludes the continuity of the success probability of π (overall planner) given the continuity of the success probability of μ^{ij} (the individual local planners).

Proposition 2. (Continuity of success probability of π) The success probability $\mathbb{P}(\text{success}|b_0, \pi)$ is continuous in \mathcal{V} if the absorption probabilities $\mathbb{P}(B^i|b, \mu^{ij})$ are continuous in \mathbf{v}^j for all i, j , and b .

Proof. Given that $\mathbb{P}(B^i|b, \mu^{ij})$ is continuous in \mathbf{v}^j , for all i, j , we want to show that $\mathbb{P}(\text{success}|\pi, b_0)$ is continuous in all \mathbf{v}^j . First, let us look at the structure of the success probability:

$$\mathbb{P}(\text{success}|b_0, \pi) = \mathbb{P}(B(\mu_0)|b_0, \mu_0) \mathbb{P}(\text{success}|B(\mu_0), \pi^g) \quad (137)$$

where μ_0 is computed using Equation (34). The term $\mathbb{P}(B(\mu_0)|b_0, \mu_0)$ on the right-hand side of Equation (137) is continuous because the continuity of $\mathbb{P}(B^i|b, \mu^{ij})$ for all i, j is assumed in this proposition. Thus, we only need to show the continuity of the second term in Equation (137). Without loss of generality we can consider $B^i = B(\mu_0)$. Then, we need to show that $\mathbb{P}(\text{success}|B^i, \pi^g)$ is continuous in \mathbf{v}^i for all i .

As we saw in Section 6.7, the probability of success from the i th FIRM node is as follows:

$$\mathbb{P}(\text{success}|B^i, \pi^g) = \Gamma_i^T (I - \mathcal{Q})^{-1} \mathcal{R}_g \quad (138)$$

Moreover, we can consider $B^{goal} = B^N$ without loss of generality; then, the (i, j) th element of matrix \mathcal{Q} is $\mathcal{Q}[i, j] = \mathbb{P}(B^i | B^j, \pi^g(B^j))$, and the j th element of vector \mathcal{R}_g is $\mathcal{R}_g[j] = \mathbb{P}(B^N | B^j, \pi^g(B^j))$.

Since we considered B^j as the stopping region of the local controller μ^{ij} , we have

$$\mathbb{P}(B^l | B^j, \mu^{il}) = 0, \text{ if } l \neq j \quad (139)$$

Therefore, all non-zero elements in the matrices \mathcal{R}_g and \mathcal{Q} are of the form $\mathbb{P}(B^l | B^j, \mu^{lj})$. Thus, given the continuity of $\mathbb{P}(B^l | b, \mu^{lj})$, the transition probability $\mathbb{P}(B^l | B^j, \mu^{lj})$ is continuous and the matrices \mathcal{R}_g and \mathcal{Q} are continuous. Therefore, $\mathbb{P}(\text{success} | B^j, \pi^g)$ and thus $\mathbb{P}(\text{success} | b_0, \pi)$ are continuous in underlying PRM nodes. \square

Now we are ready to prove Theorem 1.

Proof. Based on the definition of probabilistic completeness under uncertainty, if there exists a successful policy $\check{\pi}$, FIRM has to find a successful policy π as the number of FIRM nodes increases unboundedly. Thus, we start by assuming that there exists a successful policy $\check{\pi} \in \Pi$ for a given initial belief b_0 . Since each policy in Π is parametrized by a PRM graph, there exists a PRM with nodes $\check{\mathcal{V}} = \{\check{\mathbf{v}}^i\}_{i=1}^N$ that parametrizes the policy $\check{\pi}$. Since $\check{\pi}$ is a successful policy, we know $\mathbb{P}(\text{success} | b_0, \check{\pi}) > p_{min}$. Thus, we can define $\epsilon^* = \mathbb{P}(\text{success} | b_0, \check{\pi}) - p_{min} > 0$.

Given Assumptions 1, 2, and 3, and based on Propositions 1 and 2, we know that $\mathbb{P}(\text{success} | b_0, \pi)$ is continuous with respect to the parameters of the local planners. In other words, for any $\epsilon > 0$, there exists a $\delta > 0$ such that if $\|\mathcal{V} - \check{\mathcal{V}}\| < \delta$, then $|\mathbb{P}(\text{success} | b_0, \pi(\cdot; \mathcal{V})) - \mathbb{P}(\text{success} | b_0, \check{\pi}(\cdot; \check{\mathcal{V}}))| < \epsilon$. The notation $\|\mathcal{V} - \check{\mathcal{V}}\| < \delta$ means that $\|\mathbf{v}^i - \check{\mathbf{v}}^i\| < \delta$, for all i , or equivalently, $\mathbf{v}^i \in \check{\Omega}_i$, for all i , where $\check{\Omega}_i$ is a ball with radius δ , centred at $\check{\mathbf{v}}^i$.

Therefore, for the introduced ϵ^* , there exists a δ^* and corresponding regions $\{\check{\Omega}_i\}_{i=1}^N$ such that if we have a PRM whose nodes (or a subset of whose nodes: a subset of nodes is sufficient, because the success probability is a non-decreasing function in terms of the number of nodes) satisfy the condition $\mathbf{v}_i^* \in \check{\Omega}_i$ for all $i = 1, \dots, N$, then the planner π parametrized by this PRM has a success probability greater than p_{min} , that is, $\mathbb{P}(\text{success} | b_0, \pi(\cdot; \mathcal{V})) > p_{min}$, and hence π is successful.

Since $\delta > 0$, the regions $\check{\Omega}_i$ have non-empty interiors. Consider a PRM with a sampling algorithm under which there is a non-zero probability of sampling in $\check{\Omega}_i$, such as uniform sampling. Then, starting with any PRM, if we increase the number of nodes, a PRM node will eventually be chosen at every $\check{\Omega}_i$ with probability one. Therefore the policy constructed based on these nodes will have a success probability greater than p_{min} ; in other words, we eventually get a successful policy if one exists. Thus, FIRM is probabilistically complete under uncertainty. \square