

Towards Multi-Robot Task-Motion Planning for Navigation in Belief Space

Antony Thomas, Fulvio Mastrogiovanni and Marco Baglietto
Department of Informatics, Bioengineering, Robotics, and Systems Engineering,
University of Genoa, Via All’Opera Pia 13, 16145 Genoa, Italy.
antony.thomas@dibris.unige.it, fulvio.mastrogiovanni@unige.it, marco.baglietto@unige.it

Abstract

Autonomous robots operating in large knowledge-intensive domains require planning in the discrete (task) space and the continuous (motion) space. In knowledge-intensive domains, on the one hand, robots have to reason at the highest-level, for example the regions to navigate to or objects to be picked up and their properties; on the other hand, the feasibility of the respective navigation tasks have to be checked at the controller execution level. Moreover, employing multiple robots offer enhanced performance capabilities over a single robot performing the same task. To this end, we present an integrated multi-robot task-motion planning framework for navigation in knowledge-intensive domains. In particular, we consider a distributed multi-robot setting incorporating mutual observations between the robots. The framework is intended for motion planning under motion and sensing uncertainty, which is formally known as belief space planning. The underlying methodology and its limitations are discussed, providing suggestions for improvements and future work. We validate key aspects of our approach in simulation.

1 Introduction

One of the essential requirements of autonomous navigation is accurate localization. Uncertainty is prevalent in most problem domains and inferring robot pose (or other variables of interest) precisely is a fundamental requirement. Moreover, in large-scale problem domains, employing multiple robots offer enhanced performance capabilities over a single robot performing the same tasks. In addition, complex real-world scenarios present the need for planning at different levels to accomplish a given set of tasks. High-level (task) planning helps break down a given set of tasks into a sequence of sub-tasks. Finding the appropriate robot motions to execute these sub-tasks requires determining the appropriate set of low-level control actions. Hence, planning should be performed in the task-motion or the discrete-continuous space [LDG⁺18].

Over the past few years, Task-Motion Planning (TMP) for manipulation has attracted significant interest among the research community [DGTN09, CAG09, SFR⁺14, KLP13, GLPK18, DKCK18]. Robot-based manipulation domain calls for discrete and continuous reasoning to execute the required action reliably. The execution of these discrete actions require continuous reasoning in the robot configuration space to generate appropriate motions. Yet, a discrete action might turn out to be unfeasible due to the end-effector’s reachability workspace. This suppicates for a tight coupling between task planning and motion planning that enables an interface for

efficient interaction between the symbolic and geometric layers. In [DGTN09] task and motion planning is interleaved by checking individual high-level action feasibility using *semantic attachments* [Wey80]. A combined search in the logical and geometric spaces is performed in [CAG09] using a state composed of both the symbolic and geometric paths. Srivastava *et al.* [SFR⁺14] implicitly incorporate geometric variables, performing symbolic-geometric mapping using a planner-independent interface layer. In [KLP13] the current state uncertainty is incorporated, modeling the planning problem in the belief space. FFRob [GLPK18] directly conduct task planning by performing search over a sampled finite set of poses, grasps and configurations. The authors of [GLPK18] extend the FF heuristics, incorporating geometric and kinematic planning constraints that provide a tight estimate of the distance to the goal. Iteratively Deepened Task and Motion Planning (IDTMP) is a constraint-based task planning approach that incorporates geometric information to account for the motion feasibility at the task planning level [DKCK18]. In our architecture, the task action costs are computed using a motion planner, similar to the motion planner information that guides the IDTMP task planner. IDTMP performs task-motion interaction using abstraction-refinement functions whereas we use *semantic attachments*.

The scope of TMP is not limited to manipulation problems alone. TMP for navigation is rather ubiquitous in most real world scenarios. Real-world planning problems in large scale environments often require solving several sub-problems. For example, while navigating to a goal, the robot might have to visit other places of interests. Visiting these places of interest is a high-level task that can be addressed using traditional task planners. As such, TMP for navigation essentially involves selecting discrete actions to navigate to different regions, objects or locations of interest in the environment and deciding the order of these visits. Synthesizing the best set of discrete actions for a given objective requires computing the navigation costs for each of these actions. Hence it is inevitable that motion planning be interleaved with task planning to compute the motion costs for each of the discrete actions. Though it can be argued that the motion costs can be approximated a priori and fed to the task planner, in large knowledge-intensive domains such an assumption can be harder to justify, especially in the presence of localization and map uncertainty. UP2TA [MRMB16] develops a unified path planning and task planning framework for mobile robot navigation. An interesting feature of UP2TA is its task planner heuristic, which is a combination of the FF heuristic [Hof03] and the Euclidean distance between the waypoints. Wong *et al.* [WYYG18] develop a task planning approach that takes into account the optimal traversal costs to synthesize a plan. Similar to UP2TA, they define tasks that are to be performed at different waypoints. However, they pre-compute the optimal path for all pairs of waypoints. In contrast, our recent work [TMB19] consider a general approach where the robot has to reason at a high-level about different objects or locations or regions to navigate to. The objects/locations/regions are instantiated to their geometric counterpart, by considering a set of sampled poses. For example, if a robot has to reach a location close to a chair, the geometric instantiations of this symbolic goal would correspond to a set of poses around the chair. PETLON [LZS18] introduces a TMP approach for navigation that is task-level optimal. However, the action costs returned by their motion planner is the trajectory length and they assume completely observable domains. The approach in [TMB19], in addition to being task-level optimal is more general since it is not limited to any particular cost function and can be easily adapted to support any general formulation. Moreover, in [TMB19] planning is performed in the belief space—a probability distribution function over the robot states.

The above discussed TMP approaches focus on single robot scenarios. This paper contributes to the literature in the following directions: (1) We propose a distributed multi-robot task-motion planning for navigating in large knowledge-intensive domains. Motion planning is performed in partially-observable state-spaces with motion and sensing uncertainty and hence the approach falls under the category of multi-robot belief space planning. (2) Our approach is task-level optimal, that is, the task plan cost returned by our approach is lower than any of the other possible task plans. (3) Finally, the TMP framework that embeds a motion planner within a task planner through an interface layer is probabilistically complete.

2 Preliminaries and Problem Definition

2.1 Task-Motion Planning

Task planning or classical planning can be defined as the process of finding a discrete sequence of actions from the current state to a desired goal state [GNT16]. The Planning Domain Definition Language (PDDL) [MGH⁺98] being the *de facto* standard syntax for task planning, we resort to the same for modeling our task domain. Motion planning finds a sequence of collision free poses from a given initial/start pose (position and orientation) to a desired goal pose [Lat91]. Task-Motion Planning (TMP) essentially involves combining discrete and continuous decision-making to facilitate efficient interaction between the two domains. Starting from an initial state, TMP

synthesizes a plan to a goal state by a concurrent or interleaved set of discrete actions and continuous collision-free motions.

Definition 1. A *task-motion domain* is a tuple $\Psi = (C, \Omega, \phi, \xi, q_0)$ where:

- C is the configuration space or the space of possible robot poses;
- $\Omega = (S, A, \gamma, s_0, S_g)$ is the task domain with S being a finite set of states, A a finite set of actions, $\gamma : S \times A \rightarrow S$ is the state transition function, $s_0 \in S$ is the initial state and $S_g \subseteq S$ is the set of goal states;
- $\phi : S \rightarrow 2^C$, is a function mapping states to the configurations. For example, if s represents the task state—the robot is in a corridor, then $\phi(s)$ corresponds to all configurations such that the robot is in the corridor;
- $\xi : A \rightarrow 2^C$, is a function mapping actions to motion plans. Note that motion plan essentially computes a collision free poses in C ;
- q_0 is the initial configuration.

Definition 2. The *TMP problem* for the TMP domain Ψ is to find a sequence of actions a_0, \dots, a_n such that $s_{i+1} = \gamma(s_i, a_i)$, $s_{n+1} \in S_g$ and to find a sequence of motion plans τ_0, \dots, τ_n such that for $i = 0, \dots, n$, it holds that (i), $\tau_i(0) \in \phi(s_i)$ and $\tau_i(1) \in \phi(s_{i+1})$, (ii) $\tau_{i+1}(0) = \tau_i(1)$ and (iii) $\tau_i \in \xi(a_i)$.

2.2 Problem Definition

We consider a distributed multi-robot TMP framework where robots operating in a known environment (map given) can observe each other, thereby facilitating collaborative multi-robot localization. When one robot detects another, the resulting localization uncertainty for both the robots is less than when there is no such mutual observations [RB02]. This stems from the integration of multi-robot constraints into the joint robot beliefs. In this work we only consider robots mutually observing themselves at the same time. However, it should be noted that multi-robot constraints can also be formulated for different robots observing the same environment at different time instances [Ind18]. At any time k , we denote the robot pose (or configuration q_k) by $x_k \doteq (x, y, \theta)$, the acquired measurement is denoted by z_k and the applied control action is denoted as u_k . We consider a standard motion $x_{k+1} = f(x_k, u_k) + w_k$, $w_k \sim \mathcal{N}(0, R_k)$, where w_k is the random unobservable noise, modeled as a zero mean Gaussian. We also consider an observation model with Gaussian noise, $z_k = h(x_k) + v_k$, $v_k \sim \mathcal{N}(0, Q_k)$.

The motion and observation models can be written probabilistically as $p(x_{k+1}|x_k, u_k)$ and $p(z_k|x_k)$, respectively. Given an initial distribution $p(x_0)$, the motion and observation models, the posterior probability distribution at time k can be written

$$p(x_k|Z_{0:k}, U_{0:k-1}) = \eta p(z_k|x_k) \int p(x_k|x_{k-1}, u_{k-1}) b[x_{k-1}] \quad (1)$$

where $Z_{0:k} \doteq \{z_0, \dots, z_k\}$, $U_{0:k-1} \doteq \{u_0, \dots, u_{k-1}\}$ and $b[x_{k-1}] \sim \mathcal{N}(\mu_{k-1}, \Sigma_{k-1})$ is the posterior probability distribution or *belief* at time $k-1$. Similarly, given an action u_k , the propagated belief can be written as $b[x_{k+1}] = \int p(x_{k+1}|x_k, u_k) b[x_k]$. Given the current belief $b[x_k]$ and the control u_k , the propagated belief parameters can be computed using the standard Extended Kalman Filter (EKF) prediction. The Jacobian of $f(\cdot)$ with respect to x_k will be denoted by F_k . Upon receiving a measurement z_k , the posterior belief $b[x_{k+1}]$ is computed using the EKF update equations.

We now formulate the multi-robot localization problem. For simplicity we consider only two robots r and r' , but the formulation can be trivially expanded to incorporate R robots. At any time k , we denote the pose of robot r by x_k^r , the acquired measurement is denoted by z_k^r and the applied control action is denoted as u_k^r . We first consider the case in which there are no mutual observations between the robots. For two robots r and r' , the joint belief at time k is given by

$$\begin{aligned} p(x_k^r, x_k^{r'} | Z_{0:k}^r, Z_{0:k}^{r'}, U_{0:k-1}^r, U_{0:k-1}^{r'}) &= p(x_k^r | Z_{0:k}^r, U_{0:k-1}^r) p(x_k^{r'} | Z_{0:k}^{r'}, U_{0:k-1}^{r'}) \\ &= \eta p(z_k^r | x_k^r) \int p(x_k^r | x_{k-1}^r, u_{k-1}^r) b[x_{k-1}^r] \cdot p(z_k^{r'} | x_k^{r'}) \int p(x_k^{r'} | x_{k-1}^{r'}, u_{k-1}^{r'}) b[x_{k-1}^{r'}] \quad (2) \end{aligned}$$

As seen above the joint belief is factorized into individual beliefs of robot r and r' . Let $x_k = [x_k^r, x_k^{r'}]$ be the joint state, then the EKF prediction can be written as

$$\bar{\mu}_k = [f(\mu_{k-1}^r, u_{k-1}^r), f(\mu_{k-1}^{r'}, u_{k-1}^{r'})]; \quad \bar{\Sigma}_k = F_{k-1} \Sigma_{k-1} F_{k-1}^T + R_{k-1} \quad (3)$$

where F_{k-1} , Σ_{k-1} and R_{k-1} are diagonal matrices. This renders the predicted covariance matrix $\bar{\Sigma}_k$ diagonal. Since we do not consider mutual observations, the Kalman gain is also a diagonal matrix

$$F_{k-1} = \begin{bmatrix} F_{k-1}^r & 0 \\ 0 & F_{k-1}^{r'} \end{bmatrix}; \quad \Sigma_{k-1} = \begin{bmatrix} \Sigma_{k-1}^r & 0 \\ 0 & \Sigma_{k-1}^{r'} \end{bmatrix}; \quad R_{k-1} = \begin{bmatrix} R_{k-1}^r & 0 \\ 0 & R_{k-1}^{r'} \end{bmatrix}; \quad K_{k-1} = \begin{bmatrix} K_{k-1}^r & 0 \\ 0 & K_{k-1}^{r'} \end{bmatrix} \quad (4)$$

giving a diagonal covariance matrix Σ_k . As such this corresponds to performing the belief propagation and updates for each robot individually [RB02]. Now let us consider the case when robots can mutually observe each other. When robot r observes robot r' at time k , the measurement constraint will be denoted by $\zeta_k^{r,r'}$. It is assumed that a common reference frame is established so that the robots can communicate relevant information with each other. The joint belief at time k is given by

$$\begin{aligned} p(x_k^r, x_k^{r'} | Z_{0:k}^r, Z_{0:k}^{r'}, \zeta_k^{r,r'}, U_{0:k-1}^r, U_{0:k-1}^{r'}) &= p(x_k^r | x_k^{r'}, Z_{0:k}^r, \zeta_k^{r,r'}, U_{0:k-1}^r) p(x_k^{r'} | Z_{0:k}^{r'}, U_{0:k-1}^{r'}) \\ &= \eta p(z_k^r | x_k^r) p(x_k^r | x_k^{r'}, Z_{0:k-1}^r, \zeta_k^{r,r'}, U_{0:k-1}^r) \cdot p(z_k^{r'} | x_k^{r'}) p(x_k^{r'} | Z_{0:k-1}^{r'}, U_{0:k-1}^{r'}) \\ &= \eta p(z_k^r | x_k^r) p(\zeta_k^{r,r'} | x_k^r, x_k^{r'}) \int p(x_k^r | x_{k-1}^r, u_{k-1}^r) b[x_{k-1}^r] \cdot p(z_k^{r'} | x_k^{r'}) \int p(x_k^{r'} | x_{k-1}^{r'}, u_{k-1}^{r'}) b[x_{k-1}^{r'}] \end{aligned} \quad (5)$$

The measurement likelihood term $p(\zeta_k^{r,r'} | x_k^r, x_k^{r'})$ introduces cross correlations in Σ_k . This is because the measurement Jacobian is computed with respect to x_{k-1}^r and $x_{k-1}^{r'}$ [MPS05] unlike the previous scenario where the measurement Jacobian was computed separately for each r using its corresponding x_{k-1}^r . We assume that robot r measures the range and bearing of r' , that is, $\zeta_k^{r,r'} = [d_k^{r,r'}, \phi_k^{r,r'}]^T$ where

$$d_k^{r,r'} = \sqrt{(x_k^{r'}(1) - x_k^r(1))^2 + (x_k^{r'}(2) - x_k^r(2))^2}; \quad \phi_k^{r,r'} = \arctan\left(\frac{x_k^{r'}(2) - x_k^r(2)}{x_k^{r'}(1) - x_k^r(1)}\right) - x_k^r(3) \quad (6)$$

Thus the Jacobian H_{k-1} , the partial derivative of the measurement function with respect to the joint state is

$$H_{k-1} = \begin{bmatrix} -\frac{(x_k^{r'}(1) - x_k^r(1))}{d_k^{r,r'}} & -\frac{(x_k^{r'}(2) - x_k^r(2))}{d_k^{r,r'}} & 0 & \frac{(x_k^{r'}(1) - x_k^r(1))}{d_k^{r,r'}} & \frac{(x_k^{r'}(2) - x_k^r(2))}{d_k^{r,r'}} & 0 \\ \frac{(x_k^{r'}(2) - x_k^r(2))}{(d_k^{r,r'})^2} & -\frac{(x_k^{r'}(1) - x_k^r(1))}{(d_k^{r,r'})^2} & -1 & -\frac{(x_k^{r'}(1) - x_k^r(1))}{(d_k^{r,r'})^2} & \frac{(x_k^{r'}(2) - x_k^r(2))}{(d_k^{r,r'})^2} & 0 \end{bmatrix} \quad (7)$$

The first time when a mutual observation is incorporated (say at time k), the measurement Jacobian introduces cross correlations in Σ_k and thereafter the matrices are no longer diagonal.

3 Approach

PDDL-based planning frameworks are limited, as they are incapable of handling rigorous numerical calculations. Most approaches perform such calculations via external modules or *semantic attachments* [Wey80]. Recently, Bernardini *et al.* [BFLP17] developed a PDDL-based temporal planner to implicitly trigger such external calls via a specialized semantic attachments called *external advisors*. They classify variables into direct V^{dir} , indirect V^{ind} and free V^{free} . V^{dir} and V^{free} variables are the normal PDDL function variables whose values are changed in the action effects, in accordance with PDDL semantics. V^{ind} variables are affected by the changes in the V^{dir} variables. A change in a V^{dir} variable invokes the external advisor which in turn computes the V^{ind} variables. The Temporal Relaxed Plan Graph (TRPG) [CCFL10] construction stage of the planner incorporates the indirect variable values for heuristic calculation, thereby synthesizing an efficient goal-directed search. We employ this semantic attachment based approach for the task-motion interface.

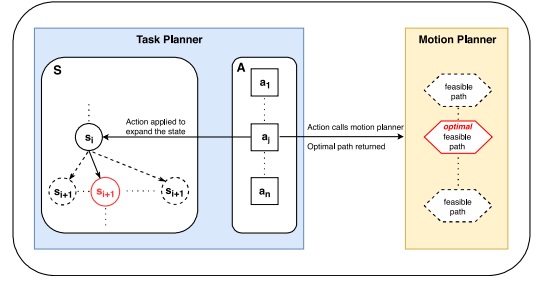
Task-Motion planning for Navigation: An overview of our approach is shown in Fig. 1b. We define $A = \{a_1, \dots, a_n\}$ as the finite set of symbolic/discrete actions available to the task planner. We use a sampling based Probabilistic Roadmap (PRM) [KSLO96] to instantiate robot poses for the task actions. To begin with, the initial mean and covariance of the robot's pose are assumed to be known. This means that, for each robot r the initial state s_0^r corresponds to a single pose instantiation q_0^r . The regions to be navigated to are also instantiated into poses, by sampling from the pose space within each region or task state. For example, let us consider a scenario where robots need to visit different rooms $L1, \dots, Ln$. The task state s_i^r might specify that the robot r is in room $L1$ and the goal state s_{i+1}^r can be for the robot r to reach room $L2$. In the considered scenario $\phi(s_i^r)$ and $\phi(s_{i+1}^r)$, that is, the mapping from states to configurations, correspond to all possible poses such that r is in rooms $L1$ and afterwards must be in $L2$, respectively. In other words, the pose instantiations are the poses that lie inside the rooms and are sampled once the map of the environment is available. Since the set of possible poses is infinite, we randomly sample a set of poses corresponding to each task state s_i . Note that these pose instantiations for each room are the same for all the robots. Furthermore, this sampling is an independent problem and the pose instantiations are incorporated while building the entire roadmap.

```

(:durative-action goto_room
:parameters(?from1 ?from2 ?to1 ?to2 - room ?r1 ?r2 - robot)
:duration (= ?duration 100)
:condition (and (at start (robot_in ?r1 ?from1)) (at start
connected ?from1 ?to1))(at start (robot_in ?r2 ?from2))
(at start (connected ?from2 ?to2)))
:effect (and (at start (not (robot_in ?r1 ?from1)))(at start
(not (robot_in ?r2 ?from2)))(at start (increase
(triggered ?r1 ?from1 ?to1 ?r2 ?from2 ?to2) 1))(at end
(robot_in ?r1 ?to1))(at end (robot_in ?r2 ?to2))(at end
(assign (triggered ?r1 ?from1 ?to1 ?r2 ?from2 ?to2) 0))
(at end (increase (act-cost) (external))))
(at end (visited ?to1)) (at end (visited ?to2)))

```

(a)



(b)

Figure 1: (a) A fragment of the PDDL room domain. (b) The discrete actions available to the planner are denoted by $A = \{a_1, a_2, a_3, \dots, a_n\}$. Different motion plans are generated for the action that requires appropriate robot motion via an external module. This module is essentially a motion planner. The optimal path among the feasible motion plans is then selected, returning the optimal cost to the task planner. The corresponding action and the optimal path is the task-motion plan for changing the task state of a robot from s_i to s_{i+1} .

A fragment of the corresponding PDDL domain is shown in Fig. 1a. The external module computes the V^{ind} values and is invoked only when a change occurs in the V^{dir} variables due to action effects. The PDDL keyword *increase* is overloaded to refer to an encapsulated object [PAFL15] and the external module is called if the PDDL action to be expanded has an *effect* of the form $(increase (v_i^{dir}) (v_j^{ind}))$, where $v_i^{dir} \in V^{dir}$ and $v_j^{ind} \in V^{ind}$. Once such an action a_i is expanded by the task planner, the corresponding start and goal states of robot r , that is, s_i^r and s_{i+1}^r are communicated to the motion planner. This is facilitated through the function $(triggered ?r1 ?from1 ?to1 ?r2 ?from2 ?to2) 1)$. This specifies that robot $r1$ is navigating from $?from1$ to $to1$ and that robot $r2$ is navigating from $?from2$ to $to2$, where $from1$, $from2$ and $to1$, $to2$ are free variables denoting the start and goal states (corresponds to different rooms) of robots $r1$ and $r2$ respectively. The function *triggered* is assigned the value of 1 each time the actions are expanded and re-initialized to 0 once the action duration is completed. This is performed so that the grounded variables are communicated to the motion planner. For each region s_i , the number of pose instantiations will be denoted by s_i^n and a particular instantiation by $s_i^{n_k}$. For each robot r , with the pose instantiation of s_i^r as the start node, for each pose instantiation of s_{i+1}^r , we simulate a sequence of controls and observations along each edge of the roadmap starting from s_i^{r,n_k} and ending in s_{i+1}^{r,n_j} , estimating the beliefs at the each of these nodes using (5). The s_{i+1}^{r,n_j} that corresponds to the minimum cost is then selected as the goal pose for robot r for the state s_{i+1}^r . Thereafter, this instantiation becomes the start node when an expansion is attempted from state s_{i+1}^r for robot r .

Though our formulation can be adapted to any generic cost function we use a standard cost function, $c \doteq M_u c_u + M_G c_G + M_\Sigma c_\Sigma$, where c_u is the control usage, c_G is the distance to the goal and c_Σ is the cost due to uncertainty, defined as $trace(\Sigma)$, where Σ is the state covariance associated with the robot belief. M_u , M_G and M_Σ are user-defined weights. The cost of the selected motion plan is then returned to the task planner as the cost of the corresponding action. The variable *external* returns the motion cost computed by the external module and achieves semantic attachment by passing its value to the task-level cost variable *act-cost* (see Fig. 1a). The task-motion plan for changing the task state of the robot from the state s_i to s_{i+1} is the ordered tuple of the action a_i and the corresponding optimal path. The tuple is appended for all the task-level actions to generate the complete task-motion plan.

Optimality: For a given roadmap, the plan synthesized by our approach is optimal at the task-level. This means that the task plan cost returned by our approach (c^*) is lower than any of the other possible task plan costs (c). Let us denote the optimal plan corresponding to c^* as π^* . Suppose that there exists a plan π with associated cost c such that $c < c^*$. If π and π^* have the same sequence of actions, this is not possible since the action costs are evaluated by the motion planner and for a given roadmap, the motion cost returned is the optimal for each action, giving $c^* \leq c$. If π and π^* have a different sequence of actions, the task planner ensures that the returned plan is optimal, giving $c^* \leq c$. Therefore, in both the case, we have $c^* \leq c$.

Completeness: We provide a sufficient condition under which the probability of our approach returning a plan approaches one exponentially with the number of samples used in the construction of the roadmap. A task planning problem, $\Omega = (S, A, \gamma, s_0, S_g)$ is complete if it does not contain any dead-ends [HN01], that is there are no states from which goal states cannot be reached. The PRM motion planner is probabilistically complete [KF11], that is the probability of failure decays to zero exponentially with the number of samples used in the construction of the roadmap. Therefore, if the motion planner terminates each time it is invoked then probability of finding a plan, if it exists, approaches one.




Figure 2: (a) Map of the corridor environment. The stars with different colors represent landmarks that aid the robots in better localization. (b) and (c) Pose covariance evolution for robots r (blue trajectory) and r' (green trajectory). The belief evolution for a single planning instantiation corresponding to a unique set of simulated observations are shown. Black dots represent the sampled poses and the covariance estimates (only (x,y) portion shown) are shown as red ellipses. (b) *config 1* incorporating mutual observations between the robots. (c) No mutual observations considered.




Figure 3: (a) Average position estimation errors. This corroborates the single instance of belief evolution as shown in Fig. 2. (b) Average planning time with increasing number of rooms to visit for 2 robots. The planning time is only about 5 seconds when 10 rooms are to be visited. (c) Average planning time in log-scale for different number of collaborating robots.

4 Empirical Evaluation

We evaluate our approach in a simulated corridor environment whose map is as shown in Fig. 2a. The robot’s can navigate to rooms $\mathbf{L} = L1, \dots, L10$ that are connected to one another through a corridor. These rooms have doors, which can either be closed or open, connecting them to the corridor. We assume that once the robot is near to a closed door that directly connects a room to the corridor, it is able to open the door— for example using human aid. Navigating to rooms can hence be encoded using a single high-level PDDL action `goto_room` as seen in Fig. 1a. The stars with different colors represent certain unique features assumed to be known and modeled like a printer or trash can that aid the robot’s in better localization. The performance are evaluated on an Intel® Core i7-6500U under Ubuntu 16.04 LTS.

We first validate our approach by considering a scenario in which robot r starting at room $L1$ has to visit room $L10$ and robot r' starting at $L10$ has to visit $L1$. Fig. 2 shows the planned trajectories with belief evolution (pose covariances) for robots r and r' . Multi-robot constraints are incorporated in Fig. 2(b) and correspond to *config 1* while *config 2* as seen in Fig. 2(c) does not consider mutual observations between the robots. Clearly, incorporating mutual observation constraints facilitate improved localization. We ran the same scenario for 25 different planning sessions, each time sampling the initial position of the robots r and r' from the known initial beliefs. The average position errors at each node along the planned trajectories are shown in Fig. 3(a). This performance evaluation shows the improved estimation accuracy for both the robots while incorporating multi-robot constraints. In particular, for robot r' in *config 2*, that is, without multi-robot constraints, it is seen that there is significant pose uncertainty along its path. This is attributed to the lack of landmarks, rendering inaccurate localization. However, incorporating multi-robot constraints significantly improves localization, with the worst case position norm error for r' reducing by about 90%.

Next, we test the scalability for an increasing number of rooms to visit. As the number of rooms to visit increase, the task-level complexity increase as the task completion requires more task-level actions. We ran *config 1* for five different scenarios that correspond to visiting 2, 4, 6, 8 and 10 rooms, respectively. For each scenario, 25 different planning sessions are conducted with the rooms to be visited being selected randomly at each run. The selected rooms are then given in the goal condition of the PDDL problem. The average planning time with two robots are shown in Fig. 3(b), the plans being computed in less than 5 seconds in all cases. As seen from the figure, the planning time increase almost linearly as there is only one single high-level action, namely, `goto_room`.

Currently, this high-level action models the navigation of both r and r' and therefore the complexity is directly dependent on the number of rooms.

Finally, we test the scalability for an increasing number of collaborating robots. In the considered scenario eight rooms are to be visited with 2, 4 and 6 different robots. For each run, the rooms to be visited are randomly selected and the average time for 25 different planning sessions are plotted in log-scale in Fig. 3(c). It is seen that planning time scales exponentially with an increasing number of collaborating robots. This is quite intuitive as planning is to be performed for all possible robot pairs.

5 Conclusion and Discussion

In this paper, we have introduced a multi-robot cooperative navigation approach for TMP under motion and sensing uncertainty. Task-motion interaction is facilitated by means of semantic attachments that return motion costs to the task planner. In this way, the action costs of the task plans are evaluated using a motion planner. The plan synthesized is optimal at the task-level since the overall action cost is less than that of other task plans generated for a given roadmap. It is also shown that our approach is probabilistically complete. Though our approach scales well with an increasing task-level complexity, there is an exponential increase in planning time as the number of cooperative robots that perform the task increases. Caching and reusing plans might help alleviate this complexity in some cases.

Presently, our approach fares well only when there are an even number of rooms to visit. For odd number of rooms to visit, the generated plan can force additional robot motions since our task-level action is defined for a pair of robots. Let us consider a scenario with 4 robots r_1, \dots, r_4 and rooms L_1, L_3 and L_7 to be visited. The synthesized plan might be that r_1 visits L_1 , r_2 visits L_3 and r_3 visits L_7 , r_4 visits L_i (where $i = 1, \dots, 10$). Robot r_4 visiting L_i is an additional room visit, even though it is not specified in the goal condition. This visit helps to obtain mutual observations between r_3 and r_4 but in practice we only need r_1 visiting L_1 , r_2 visiting L_3 and r_3 visiting L_7 . However, it is to be noted that this formulation still preserves the task-level optimality. Decoupling and defining the action for each robot can rescind the additional motions. Yet, the computational challenge associated with the semantic attachment architecture needs to be analyzed and it is an immediate work for future.

References

- [BFLP17] Sara Bernardini, Maria Fox, Derek Long, and Chiara Piacentini. Boosting Search Guidance in Problems with Semantic Attachments. In *International Conference on Automated Planning and Scheduling (ICAPS)*, pages 29–37, Pittsburgh, PA, USA, June 2017.
- [CAG09] Stéphane Cambon, Rachid Alami, and Fabien Gravot. A hybrid approach to intricate motion, manipulation and task planning. *The International Journal of Robotics Research*, 28(1):104–126, 2009.
- [CCFL10] Amanda Jane Coles, Andrew Ian Coles, Maria Fox, and Derek Long. Forward-chaining partial-order planning. In *Twentieth International Conference on Automated Planning and Scheduling*, 2010.
- [DGTN09] Christian Dornhege, Marc Gissler, Matthias Teschner, and Bernhard Nebel. Integrating symbolic and geometric planning for mobile manipulation. In *Safety, Security & Rescue Robotics (SSRR), IEEE International Workshop on*, pages 1–6. IEEE, 2009.
- [DKCK18] Neil T. Dantam, Zachary K. Kingston, Swarat Chaudhuri, and Lydia E. Kavraki. An Incremental Constraint-Based Framework for Task and Motion Planning. *International Journal of Robotics Research, Special Issue on the 2016 Robotics: Science and Systems Conference*, 37(10):1134–1151, 2018.
- [GLPK18] Caelan Reed Garrett, Tomas Lozano-Perez, and Leslie Pack Kaelbling. FFRob: Leveraging symbolic planning for efficient task and motion planning. *The International Journal of Robotics Research*, 37(1):104–136, 2018.
- [GNT16] Malik Ghallab, Dana Nau, and Paolo Traverso. *Automated planning and acting*. Cambridge University Press, 2016.

- [HN01] Jörg Hoffmann and Bernhard Nebel. The FF Planning System: Fast Plan Generation Through Heuristic Search. *Journal of Artificial Intelligence Research*, 14:253–302, 2001.
- [Hof03] Jörg Hoffmann. The Metric-FF Planning System: Translating “Ignoring Delete Lists” to Numeric State Variables. *Journal of Artificial Intelligence Research*, 20:291–341, 2003.
- [Ind18] Vadim Indelman. Cooperative multi-robot belief space planning for autonomous navigation in unknown environments. *Autonomous Robots*, 42(2):353–373, 2018.
- [KF11] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7):846–894, 2011.
- [KLP13] Leslie Pack Kaelbling and Tomás Lozano-Pérez. Integrated task and motion planning in belief space. *The International Journal of Robotics Research*, 32(9-10):1194–1227, 2013.
- [KSLO96] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE transactions on Robotics and Automation*, 12(4):566–580, 1996.
- [Lat91] Jean-Claude Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, 1991.
- [LDG⁺18] Fabien Lagriffoul, Neil T. Dantam, Caelan Garrett, Aliakbar Akbari, Siddharth Srivastava, and Lydia E. Kavraki. Platform-independent benchmarks for task and motion planning. *Robotics and Automation Letters*, 2018.
- [LZS18] Shih-Yun Lo, Shiqi Zhang, and Peter Stone. Petlon: Planning efficiently for task-level-optimal navigation. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 220–228. International Foundation for Autonomous Agents and Multiagent Systems, 2018.
- [MGH⁺98] Drew McDermott, Malik Ghallab, Adele Howe, Craig Knoblock, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. PDDL- The Planning Domain Definition Language. In *AIPS-98 Planning Competition Committee*, 1998.
- [MPS05] Agostino Martinelli, Frederic Pont, and Roland Siegwart. Multi-robot localization using relative observations. In *Proceedings of the 2005 IEEE international conference on robotics and automation*, pages 2797–2802. IEEE, 2005.
- [MRMB16] Pablo Muñoz, María D R-Moreno, and David F Barrero. Unified framework for path-planning and task-planning for autonomous robots. *Robotics and Autonomous Systems*, 82:1–14, 2016.
- [PAFL15] Chiara Piacentini, Varvara Alimisis, Maria Fox, and Derek Long. An extension of metric temporal planning with application to ac voltage control. *Artificial intelligence*, 229:210–245, 2015.
- [RB02] Stergios I Roulletiotis and George A Bekey. Distributed multirobot localization. *IEEE transactions on robotics and automation*, 18(5):781–795, 2002.
- [SFR⁺14] Siddharth Srivastava, Eugene Fang, Lorenzo Riano, Rohan Chitnis, Stuart Russell, and Pieter Abbeel. Combined task and motion planning through an extensible planner-independent interface layer. In *Robotics and Automation (ICRA), IEEE International Conference on*, pages 639–646. IEEE, 2014.
- [TMB19] Antony Thomas, Fulvio Mastrogiovanni, and Marco Baglietto. Task-Motion Planning for Navigation in Belief Space. In *The International Symposium on Robotics Research (ISRR)*, 2019.
- [Wey80] Richard W. Weyhrauch. Prolegomena to a theory of mechanized formal reasoning. *Artificial Intelligence*, 13, 1980.
- [WYYG18] Cuebong Wong, Erfu Yang, Xiu-Tian Yan, and Dongbing Gu. Optimal path planning based on a multi-tree T-RRT* approach for robotic task planning in continuous cost spaces. In *2018 12th France-Japan and 10th Europe-Asia Congress on Mechatronics*, pages 242–247. IEEE, 2018.