Task-Motion Planning in Belief Space

Antony Thomas, Sunny Amatya, Fulvio Mastrogiovanni, 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, sunnyamatya@gmail.com, fulvio.mastrogiovanni@unige.it, marco.baglietto@unige.it}

***Abstract*—We discuss an ongoing research in the area of integrated task-motion planning (TMP). Autonomous robotic agents operating in real world complex scenarios require planning in the discrete (task) space and the continuous (motion) space. We propose a framework for integrating belief space reasoning within a hybrid task planner. The expressive power of PDDL+ combined with external advisors performs rigorous belief prediction and update calculations. The underlying methodology for the development of the hybrid planner is discussed, providing promising future directions to be explored in the context of single and multi-robot planning systems.**

I. Introduction

Autonomous robots operating in complex real world scenarios require different levels of planning to execute their tasks. High-level (task) planning helps to break down a given set of tasks into a sequence of sub-tasks, depending on the required level of abstraction. Actual execution of each of these sub-tasks would require low-level control actions (i.e., motions). Hence, planning should be performed in the task-motion or the discrete-continuous space. In recent years, combining high-level task planning with the low-level motion planning has been a subject of great interest among the Robotics and Artificial Intelligence (AI) community. This is only natural as one of the ultimate goals in robotics is to create autonomous agents accepting high-level task descriptions to execute them without further human intervention. AI planning frameworks like the Planning Domain Definition Language (PDDL) [6] mainly focus on high-level task planning supposing that the geometric preconditions (e.g., available grasp configurations) for the robot motion are achievable to carry out such tasks (e.g., pick up task for the grasping). However, in reality, such an assumption can be catastrophic as an action or sequence of actions generated by the task planning algorithm might turn out to be unfeasible at the controller execution level.

Complex scenarios often induce uncertainties. Such uncertainties arise due to insufficient knowledge about the environment (partial observability), inexact robot motion or imperfect sensing. In such scenarios, the robot poses or other variables of interest (states) can only be dealt with in terms of probabilities. Planning is therefore done in the *belief* space, which is the probability distributions over the possible agent states. Such a problem falls under the category of partially observable Markov decision processes (POMDPs). To perform efficiently in these scenarios the robot needs to plan for actions, which when executed help gain more information, thereby providing a better belief estimate. Hence the task planner should be capable of reasoning in the belief space while synthesizing a plan. The motion planner might encounter unexpected scenarios notwithstanding the plan provided. This calls for a re-plan, updating the task planner with the new belief, resulting in a cyclic interdependency. Consequently, both task and motion planning are interdependent and should not be considered as separate processes.

In this paper, we discuss an ongoing research work in the area of TMP. PDDL+ [4] is used to model the planning task, providing the robot a sequence of actions (a plan) that can be passed on to the low-level motion planner for execution. Assuming Gaussian models, synthesizing an efficient plan requires performing the belief prediction and update within PDDL+ effects. Nonetheless, PDDL is restricted, as it is incapable of handling rigorous numerical calculations. Most approaches perform such calculations via an external module or semantic attachments, e.g. [3]. These semantic attachments are explicit external calls that perform numerical calculations during planner runtime. Yet, the effects returned by these semantic attachments are not exploited in identifying *helpful actions* and hence do not provide any heuristic guidance.

Recently Bernardini et al. [1] developed a PDDL based POPF-TIF planner to implicitly trigger such external calls via external advisors. They classify variables into direct, indirect and free variables. Direct (free) variables are the normal PDDL function variables whose values are changed in the action effects, in accordance with the PDDL semantics. The indirect variables are affected by the changes in the direct variables. A change in a direct variable triggers the external advisor which in turn updates the indirect variables. This planner is based on a temporal extension of the metric-FF planner [5]. An intriguing feature of the POPF-TIF planner is its ability to approximate the values of the indirect variables at the Temporal Relaxed Plan Graph (TRPG) construction stage. Approximate effects of the indirect variables are used while constructing the TRPG. Since the TRPG construction at each state is performed to extract the heuristic value for that state, the approximation values help in an efficient goal-directed search. During the forward state space search, the external advisor is called, updating the indirect variables with the exact values.

II. Problem Definition

We consider a mobile robot in a known environment (the map is given) with uncertainty in its initial pose. The set of landmarks in the environment are given by $\mathbf{l} = \{l_1, l_2, ..., l_n\}$. Currently, to emphasize our approach, we assume a simplistic cost function, the goal state robot pose uncertainty to be within a threshold (e.g., a bound on the trace of the final pose uncertainty). In addition, we also require the number of landmarks visited to be a minimum, i.e., $|\mathbf{l}_{visited}|$ is minimized. $|\mathbf{l}_{visited}| = 0$ can also be a plan, depending on the robot's initial pose uncertainty and the distance to goal. The robot motion model used is given by

$$x_{k+1} = f(x_k, u_k) + w_k \ , \ w_k \sim \mathcal{N}(0, \Sigma_w) \qquad (1)$$

where w_k and Σ_w are the Gaussian noise and process covariance respectively. We also assume a general observation model with Gaussian noise,

$$z_k = h(x_k) + v_k \ , \ v_k \sim \mathcal{N}(0, \Sigma_v) \qquad (2)$$

The approximate changes in the pose covariance at each time-step are computed within the *process* effect of the PDDL+ planner. The gradual increase in uncertainty during robot motion leads to a positive change and the reduction in uncertainty while observing a landmark is incorporated by an approximate negative change in the covariance. The task planner first constructs the staged RPG levels making use of the approximate covariance effects. This approximation helps the relaxed graph in identifying the *helpful actions* (the landmarks that can be visited in the next state), besides providing an efficient heuristic function to select the next best landmark to visit. The external advisor is then called for during the forward search to update the resulting states with the exact covariance values. The resulting plan is then fed to a motion planner for execution. Currently, this task-motion interaction is handled using the ROSPlan framework [2] which uses a PRM based motion planner. Yet, obstacles are much harder to model within the task planner, calling for a re-plan each time the motion planner encounters an obstacle that is newly added or occluded during the mapping phase. Nevertheless, we believe that an efficient approximation approach combined with a coherent plan serialization will help in minimizing the number of required re-plans, resulting in a much robust TMP planner.

III. Contribution

Building upon the POPF-TIF planner and leveraging the ROSPlan framework, we develop hybrid planning framework capable of reasoning in a robot belief space, while synthesizing a plan. Our hybrid planner is built on the state-of-the-art DiNo planner [7] which is capable of handling the hybrid nature of our problem domain. Conventional motion planning approaches in the belief space need to consider all the different possible plans for the next L (look-ahead) steps, selecting the one that is optimal. Cost computation also needs to be performed at each of these L look-ahead steps. This can be computationally very expensive while operating in larger domains with hundreds of possible actions to select from. We reduce the burden at the motion planning level by performing an efficient search at the task planning level. Techniques like, planning graphs, planning as satisfiability and heuristic-search planning have led to dominant task planning approaches like fast downward and FF planning systems. These planners are easily adaptable to larger domains and are capable of synthesizing highly efficient plans within limited time.

IV. Conclusion

We discuss an ongoing research for mobile robot localization in the area of TMP, equipping a hybrid task planner with the capability of reasoning in the belief space of the robot. Expressive power of PDDL+ combined with external advisors perform rigorous belief prediction and update calculations. For now, the framework is being handled within the ROSPlan architecture, for communication between the environment, accumulated state estimation, task planner and the low-level controller. The underlying methodology for the development of such a hybrid planner has been discussed using a simple example. Upon development, we envision to extend this framework to be able to solve the well-known benchmark problems like blocks-world, kitchen domain etc. Extending the approach to collaborative multi-robot TMP is another promising future direction.

References

[1] 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.

[2] Michael Cashmore, Maria Fox, Derek Long, Daniele Magazzeni, Bram Ridder, Arnau Carrera, Narcis Palomeras, Natalia Hurtos, and Marc Carreras. ROSPlan: Planning in the Robot Operating System. In *International Conference on Automated Planning and Scheduling*, pages 333–341, Jerusalem, Israel, June 2015.

[3] Christian Dornhege, Patrick Eyerich, Thomas Keller, Sebastian Trüg, Michael Brenner, and Bernhard Nebel. Semantic Attachments for Domain-Independent Planning Systems. In *Towards Service Robots for Everyday Environments*, pages 99–115. Springer Berlin Heidelberg, 2012.

[4] Maria Fox and Derek Long. Modelling Mixed Discrete-Continuous Domains for Planning. *Journal of Artificial Intelligence Research*, 27(1):235–297, 2006.

[5] Jörg Hoffmann. The Metric-FF Planning System: Translating "Ignoring Delete Lists" to Numeric State Variables. *Journal of Artificial Intelligence Research*, 20:291–341, 2003.

[6] 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.

[7] Wiktor Mateusz Piotrowski, Maria Fox, Derek Long, Daniele Magazzeni, and Fabio Mercorio. Heuristic Planning for PDDL+ Domains. In *AAAI Workshop: Planning for Hybrid Systems*, Phoenix, Arizona, USA, July 2016.