

A Reachability Tree-Based Algorithm for Robot Task and Motion Planning

Kanghyun Kim¹, Daehyung Park², and Min Jun Kim¹

Abstract—This paper presents a novel algorithm for robot task and motion planning (TAMP) problems by utilizing a reachability tree. While tree-based algorithms are known for their speed and simplicity in motion planning (MP), they are not well-suited for TAMP problems that involve both abstracted and geometrical state variables. To address this challenge, we propose a hierarchical sampling strategy, which first generates an abstracted task plan using Monte Carlo tree search (MCTS) and then fills in the details with a geometrically feasible motion trajectory. Moreover, we show that the performance of the proposed method can be significantly enhanced by selecting an appropriate reward for MCTS and by using a pre-generated goal state that is guaranteed to be geometrically feasible. A comparative study using TAMP benchmark problems demonstrates the effectiveness of the proposed approach.

I. INTRODUCTION

A number of studies are being conducted to enable robots to achieve autonomy and perform everyday tasks in an unstructured environment [1]. To this end, on top of low-level control algorithms, e.g., [2], [3], the robot should be able to plan an abstracted action sequence (by task planning/planner, TP), while ensuring the geometric feasibility of each action (by motion planning/planner, MP). It is worthwhile to note that the TP involves finding a discrete task-level action sequence, and the MP involves finding a collision-free path for a robot. Both TP and MP methods have already been studied extensively in the literature [4], [5].

To accomplish the robot manipulation task autonomously, both TP and MP problems need to be solved simultaneously, and this is referred to as task and motion planning (TAMP). Fig. 1 shows a motivating example of a TAMP problem called *kitchen* domain. In this problem, the robot can have both geometric (e.g. joint angles, object poses) and abstracted (e.g. `cooked[object]`, `washed[object]`) state variables, and the TAMP planner should generate actions that satisfy not only geometric constraints (e.g. collision-free) but also task-level constraints defined in the problem (e.g. all food blocks must be washed before cooking). A major challenge of this setup is that, when the *sink* and/or *stove* have tight geometric constraints, the planner is likely to require a large amount of samplings. TAMP problems generally become more complex when the geometry of

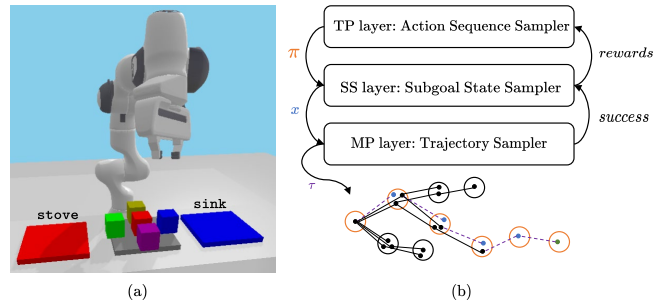


Fig. 1. Manipulation planning problems often involve tight geometric constraints as well as task-level planning. (a) In the kitchen domain, the robot needs to cook a given food block after washing it. (b) We propose an algorithm to build a reachability tree with hybrid states using three layers of planning hierarchies towards task and motion planning problems.

the problem narrows the solution space. In this case, the high-level TP tends to suggest an action sequence that is not geometrically valid. Consequently, the low-level MP performs unnecessary collision checks, making the entire planner inefficient.

Our planning approach utilizes a *reachability tree* that captures the connectivity of sampled states, like a rapidly-exploring random tree (RRT [6]). However, unlike MP problems, TAMP problems have a hybrid state space consisting of both abstracted and geometrical state variables, which poses a challenge in defining an effective distance metric between two states. For this reason, it is not trivial to utilize an RRT-like extension scheme, i.e. sampling a random node and extending the tree toward it, to TAMP problems.

To address this issue, we propose a hierarchical strategy that first samples an abstract action sequence by employing Monte Carlo tree search (MCTS) in conjunction with a symbolic planner (which can be of any off-the-shelf planner), and subsequently samples the corresponding geometric state variables. Moreover, we introduce two methods to improve planning speed. First, the reward of MCTS is generated by evaluating the result of TP (i.e. action sequence) using MP. Second, when performing MP, a goal state candidate is generated in advance based on the given action sequence. We found that prechecking the feasibility of the goal candidate can significantly improve the search speed, particularly in problems with tight goal constraints, such as the *kitchen* domain. To validate the effectiveness of our approach, we compared it with our baselines and PDDLStream [7].

II. RELATED WORK

Unlike the typical MP problems, the overall system configuration of manipulation planning is extended due to the existence of unactuated objects. The topology of this problem

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) No. 2021R1C1C1005232, and No. 2021R1A4A3032834

¹K. Kim, and M. J. Kim are with the School of Electrical Engineering, and ²D. Park is with the School of Computing, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Republic of Korea. E-mail: {kh11kim, daehyung, minjun.kim} at kaist.ac.kr

has already been well-established by introducing a concept of a *mode* that parametrizes varying collision-free space of the robot. This type of problem is called a *multi-modal motion planning* (MMMP) [8]. In addition to MMMP, the TAMP problem requires a more general setting that considers abstracted states and actions [1]. Namely, despite the fact that MMMP and TAMP are both tackling manipulation planning problems, they have slightly different problem setups. Consequently, solution approaches to MMMP and TAMP have tended to be developed separately.

In the field of TAMP, a symbolic planner-based method is often used [7], [9]–[13]. In most cases, the continuous domain is pre-discretized or sampled online, and a symbolic planner using domain-independent heuristics is utilized to search discretized state space. These methods are very general as they can plan even with abstract states and actions. However, if the continuous domain is not discretized densely, a valid trajectory of the plan may not be found. Conversely, if discretization is too dense, the size of the symbolic planning problem becomes very large. These issues may become problematic when geometric constraints are tight. Consequently, the overall planning easily becomes inefficient as the iteration continues.

In the field of MMMP, researchers often utilize a planner that directly handles geometry by retaining a reachability tree that represents the connectivity of sampled states [8], [14]–[17]. One of the most well-known approaches in this area is the RRT-like MMMP method [18], which samples a random state and extends the search tree towards that sample. Previous studies have shown that well-designed heuristics can improve search speed. [18] biases promising transitions using a utility table, which is calculated offline. In [19], [20], the edge weight of the *mode transition graph* is updated using motion planning, and the result of the graph search is utilized to guide the tree extension.

Inspired by the previous works, we propose a reachability-tree-based approach for TAMP problems in this paper.

III. MANIPULATION PLANNING PROBLEM

We formulate a TAMP problem for robotic manipulation by combining MP, MMMP, and TP methods. In the following, we briefly define our problem first, and then explain the notations in detail.

A. Problem Definition

We define a manipulation planning problem as a tuple, (x_I, \mathcal{X}_G) , where $x_I \in \mathcal{X}$ is an initial state, and \mathcal{X}_G is a set of goal states. The objective of this problem is to find a sequence of states $[x_0, x_1, \dots, x_k]$, where $x_0 = x_I$, $x_k \in \mathcal{X}_G$, and k is the number of state transitions. We assume a quasi-static and fully observable system in this work.

It should be noted that a goal set \mathcal{X}_G can be implicitly expressed as a set of values of some state variables, which is denoted as G . A goal set \mathcal{X}_G is derived as a set of states, which can be made by conjunction of values in G , where unmentioned state variables may have arbitrary values.

B. Notations

1) *Object o , Movable object m , Robot r* : In TP, \mathcal{O} is a set of objects that includes a set of movable objects $\mathcal{M} \subset \mathcal{O}$, and the robot $r \in \mathcal{O}$. Non-movable objects are also used for placing movable objects.

2) *Attachment α* : The kinematic relationship between the two objects is defined using an *attachment*. An attachment α of a movable object m can be defined as a tuple consisting of itself, the parent object, and the transformation between the two,

$$\alpha_m = (m, p, {}^mT_p),$$

where $m \in \mathcal{M}$, $p \in \mathcal{O}$, ${}^mT_p \in SE(3)$.

If the kinematic parent p is a robot r , the attachment is called *grasp*; otherwise, it is called *placement*.

3) *Mode σ* : In most manipulation problems, each movable object has one attachment (i.e., a kinematic tree). In this setting, a tuple of attachments of all movable objects in a certain state is called a *mode*:

$$\sigma = \{\alpha_m | m \in \mathcal{M}\} \in \Sigma,$$

where the mode space Σ is an infinite set of modes. Once the mode is determined, all unactuated degrees of freedom in the system are fixed. Namely, a mode defines the collision-free configuration space of a robot, which we denote as $\mathcal{C}_{free}^\sigma$.

4) *Abstract state s* : The planning domain definition language (PDDL) [21] can be used to abstract a state of the system, which we denote as s . Among variables that compose an abstract state, there are abstract attachments in the following form,

$$\tilde{\alpha}_m = (\text{attached } ?m \ ?p),$$

that indicates which object is attached to which parent. Note that the abstract version of an attachment simply omits transformation information mT_p between two objects. Likewise, we can also naturally define an abstract mode $\tilde{\sigma}$. Conversely, the set of variables that is not abstract attachment is defined as non-geometric state ψ . Therefore, $s = (\psi, \tilde{\sigma})$.

5) *Abstract action a* : The abstract action space \mathcal{A} is a set of actions a , which can make a transition from an abstract state s when applicable. We denote an applicable transition as,

$$s' = T(s, a).$$

Similar to the abstract state, we call the special type of action $a_g \in \mathcal{A}$, which can change abstract mode $\tilde{\sigma}$ a geometric action, i.e. `Pick` and `Place`. Non-geometric action $a_n \in \mathcal{A}$ can only change non-geometric variables, such as `wash` or `cook` in *kitchen* domain (see V). A set of geometric actions and a set of non-geometric actions are denoted by \mathcal{A}_g and \mathcal{A}_n , respectively.

6) *Hybrid state x* : Now we can represent the full state of the system,

$$x \in \mathcal{X} = \Psi \times \Sigma \times \mathcal{C},$$

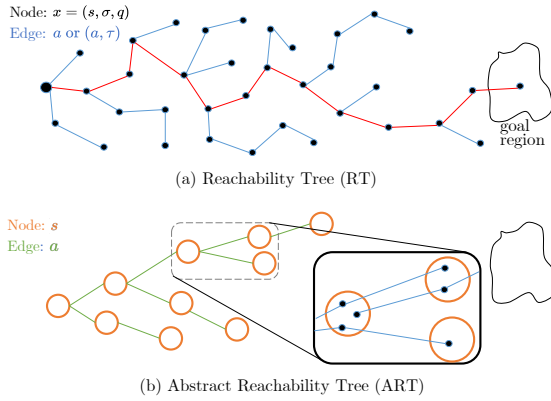


Fig. 2. Our algorithm maintains two trees: (a) The reachability tree (RT) tracks all reachable states from the initial state, and (b) the abstract reachability tree (ART) is an abstraction of the reachability tree, with each node containing multiple RT nodes.

where q is a configuration of the robot, and $\Psi, \Sigma, \mathcal{C}$, and \mathcal{X} are non-geometric state space, mode space, robot configuration space, and hybrid state space, respectively. For the sake of simplicity, we will use the redundant representation as,

$$x = (s, \sigma, q),$$

where the information of the abstract mode is redundantly expressed in both s and σ .

7) *State transition*: There are 3 types of state transitions between two states $x = (s, \sigma, q)$ and $x' = (s', \sigma', q')$, and each transition induces a three-level hierarchical structure in our planner.

- *Configuration transition*: Only q is changed. That is,

$$\|q' - q\| < \epsilon, \quad \sigma = \sigma', s = s',$$

where $\epsilon \in \mathbb{R}$ is a small number which implies motion continuity.

- *Mode transition*: σ can be changed only if q is in the subspace induced by two modes,

$$\sigma \neq \sigma', \quad q = q' \in \mathcal{C}_{free}^\sigma \cap \mathcal{C}_{free}^{\sigma'}, \quad \psi_s = \psi_{s'},$$

where $s' = T(s, a_g)$, $a_g \in \mathcal{A}_g$. ϕ_s and $\phi_{s'}$ represents the nongeometric states of s and s' , respectively.

- *Non-geometric state transition*: Only non-geometric state ψ is changed, therefore,

$$\psi_s \neq \psi_{s'}, \quad q = q', \quad \sigma = \sigma',$$

where $s' = T(s, a_n)$, $a_n \in \mathcal{A}_n$.

IV. OUR PLANNING ALGORITHM*

Our algorithm expands a reachability tree (RT) until the goal state is included in the tree (Fig. 2(a)). The key idea is to use sampling of action sequences (i.e. abstract plans) to expand the RT. To accomplish this, we use the *abstract reachability tree* (ART), which is an abstraction of the RT (Fig. 2(b)). Since ART only considers task-level information,

*The corresponding video and code are available on the project website: <https://sites.google.com/view/tree-based-tamp>

*Note that the \cdot (dot) operator is used to indicate a member of a node for algorithm description.

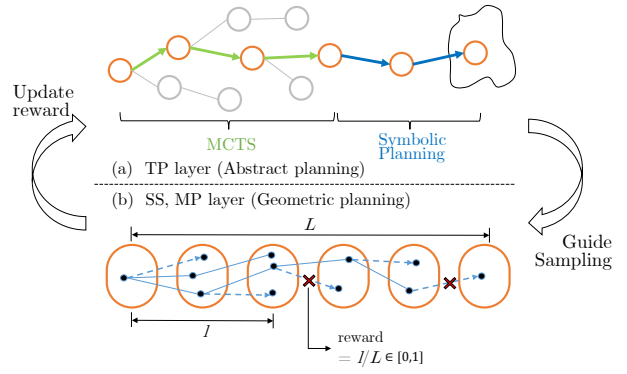


Fig. 3. (a) TP layer utilizes MCTS and a symbolic planner to sample an action sequence. (b) The RT extension is guided by the sampled action sequence. When the extension fails, the SS layer calculates a reward.

one ART node naturally includes multiple RT nodes, as depicted by the black solid box in Fig. 2(b).

Fig. 3 illustrates an overview of the entire algorithm. The tree expansion is performed hierarchically using three layers: (1) action sequence sampling by the TP layer, (2) subgoal sampling (SS) by the SS layer, and (3) trajectory sampling by the MP layer.

First, in the TP layer, an action sequence is sampled using only ART. The action sequence sampler consists of a randomized tree search and a symbolic planner. That is, a partial plan is obtained by the randomized tree search and the remaining plan is completed using a symbolic planner (Fig. 3(a)). Note that using MCTS as a randomized tree search naturally leads to informed sampling in the TP layer. Second, in the SS layer, new RT nodes for exploration are sampled based on the abstract plan, which is the outcome of the TP layer. Note that attachment sampling is required for each abstract action (e.g. `pick` requires a grasp). The SS layer samples the attachments in a batch manner to create a goal state candidate in advance. If the candidate is infeasible, all batch attachments are rejected. Otherwise, the sampled attachments are used to create new RT nodes. Lastly, the MP layer samples the trajectory between a tree node and a new node. If the MP fails, the SS layer calculates the ratio of the successful extension length of an abstract plan as a reward (Fig. 3(b)). The average of rewards is then updated to ART and utilized to bias the sampling of better abstract plans in the next iteration.

A formal description of the two main trees is as follows.

- *Reachability tree (RT)* \mathcal{G} : Each node corresponds to a hybrid state x , and each edge represents the transition between two states. For geometric actions, the edge also includes a trajectory, while for non-geometric actions, the edge only includes the action itself.
- *Abstract reachability tree (ART)* \mathcal{T} : Each node corresponds to an abstract state s , and each edge represents an abstract action a . In addition, each node contains information such as the number of visits n_{visit} , the sum of rewards r_{total} , and a set of RT node pointers V_s . Note that n_{visit} and r_{total} are used in the MCTS in the TP layer, while V_s is used for parent sampling in the SS layer.

Algorithm 1: Task Planning Layer

Input : Initial state x_I , Implicitly defined goal set G ,
SS layer Iteration number k_{ss}

Output: Reachability Tree \mathcal{G}

```
1  $\mathcal{G} \leftarrow \text{ReachabilityTree}()$ 
2  $\mathcal{T} \leftarrow \text{AbstractStateTree}()$ 
3  $\mathcal{G}.\text{AddNode}(x_I)$ 
4  $\mathcal{T}.\text{AddNode}(x_I.s)$ 
5 while within computational budget do
6    $r \leftarrow \text{Queue}()$ 
7    $\pi \leftarrow \text{SampleActionSeq}(\mathcal{T}, G)$  // Alg.2
8    $n_s \leftarrow \text{GetAbsStateNodeSeq}(\mathcal{T}, \pi)$ 
9   for  $i \leftarrow 1$  to  $k_{ss}$  do
10     $rewards \leftarrow \text{SSLayer}(\pi, n_s, G, \mathcal{G})$  // Alg.3
11     $r \leftarrow \text{Concat}(r, rewards)$ 
12   $\text{UpdateTree}(n_s, \text{Average}(r))$ 
```

Algorithm 2: Action Sequence Sampler

Input : Abstract reachability tree \mathcal{T} , Implicitly defined
goal set G

Output: Action sequence π

```
1  $n_{last}, \pi_{rand} \leftarrow \text{RandomizedTreeSearch}(\mathcal{T})$ 
2  $\pi_{plan} \leftarrow \text{TaskPlanner}(n_{last}, G)$ 
3  $\text{ExtendTree}(\mathcal{T}, n_{last}, \pi_{plan})$ 
4  $\pi \leftarrow \text{Concat}(\pi_{rand}, \pi_{plan})$ 
```

A. Task Planning Layer (Alg. 1)

The task planning layer is shown in Alg. 1 and Fig. 3(a). First, RT \mathcal{G} and ART \mathcal{T} are initialized (line 1-4). Within the loop, `SampleActionSeq` (Alg. 2) outputs a sequence of actions π , ensuring that the initial state can be successfully transitioned to the abstract goal state set. Note that the set of abstract goal states can be derived by the implicitly defined goal set G (line 7). Moreover, a node sequence of abstract states n_s can also be derived using π (line 8). This sampled sequence π and n_s are evaluated k_{ss} times by the subgoal sampling layer (line 10). The ART \mathcal{T} is updated using the average of rewards obtained by the SS layer.

B. Action Sequence Sampler (Alg. 2)

The Action Sequence Sampler should provide a randomized sequence of abstract actions that can guide RT \mathcal{G} to the goal set \mathcal{X}_G . The key idea is to create a partial action sequence using a randomized tree search, where the remaining plan can be completed using a symbolic planner.

For `RandomizedTreeSearch`, we use a mechanism of Monte Carlo tree search (MCTS), which continues to evaluate the value of each node and outputs a better state sequence in the next iteration (line 1). In such a framework, it is important to balance between exploration for searching for better rewards and exploitation that makes use of current information. For this purpose, we employ a ϵ -greedy strategy in node selection. That is, the child node is selected with probability ϵ , otherwise, the child with the highest value is chosen. Because our lower-level planners are semi-complete, it should be possible to resample the same action sequence. For this reason, each tree node has a probability of terminating the sequence π . This allows the `RandomizedTreeSearch` to resample the previously sampled sequence. As a result, the

Algorithm 3: Subgoal Sampling Layer

Input : Action sequence π , Abstract state node sequence
 n_s , Implicitly defined goal set G , Reachability
tree \mathcal{G} , Iteration number of goal candidate
sampling k_{goal}

Output: A sequence of rewards $rewards$

// Goal candidate generation

```
1  $x_I \leftarrow \mathcal{G}.\text{root}$ 
2 for  $i \leftarrow 1$  to  $k_{goal}$  do
3    $\alpha \leftarrow \text{SampleBatchAttachments}(\pi, G)$  // Alg.4
4    $x_G \leftarrow \text{MakeGoalCandidate}(x_I, \alpha)$ 
5   if there is no collision at  $x_G$  then break
6    $x_G \leftarrow \text{null}$ 
7 if  $x_G \leftarrow \text{null}$  then return
// Batch Extension of RT
8  $rewards \leftarrow \text{Queue}()$ 
9 for  $i \leftarrow 1$  to  $\pi.\text{length}$  do
10   $n \leftarrow n_s[i]; a \leftarrow \pi[i]; \alpha \leftarrow \alpha[i]; n' \leftarrow n_s[i+1]$ 
11   $x \leftarrow \text{SampleRTNode}(n)$  // Select parent
12  if  $x = \text{null}$  then break
13   $s \leftarrow n.s; \sigma \leftarrow x.\sigma; q \leftarrow x.q; s' \leftarrow n'.s$ 
14
15  if  $a \notin \mathcal{A}_g$  then
// Non-geometric transition
16     $x' \leftarrow (s', \sigma, q)$ 
17  else
// Mode transition
18     $\sigma' \leftarrow \text{MakeNextMode}(\sigma, \alpha)$ 
19     $q' \leftarrow \text{SampleTransition}(\sigma, \sigma')$ 
20     $x' \leftarrow (s', \sigma', q')$ 
21   $success \leftarrow \text{MPLayer}(x, x', \mathcal{G})$ 
22  if  $success$  then  $n'.V_s \leftarrow n'.V_s \cup \{x'\}$ 
23  else  $rewards.\text{Push}((i-1)/\pi.\text{length})$  // Reward
24 if the last tree extension was successful then
25   $success \leftarrow \text{MPLayer}(x', x_G, \mathcal{G})$ 
26  if  $success$  then
27     $\mathcal{G}.\text{solution} \leftarrow x_G$  // Solution found
```

`RandomizedTreeSearch` outputs the randomized plan π_{plan} and the final ART node n_{last} .

`TaskPlanner` can be any off-the-shelf symbolic planner (line 2). Then, the planned sequence π_{plan} is used to expand the ART \mathcal{T} (line 3). This increases the probability of selecting a similar sequence to the current sequence π in the next iterations.

C. Subgoal Sampling Layer (Alg. 3)

In the SS layer, the expansion of the RT is achieved through three steps. Firstly, this layer performs batch sampling of attachments required to realize the sampled action sequence. Secondly, new transition states are generated using the tree nodes and the sampled attachments. Lastly, a reward signal is generated to provide information to the TP layer.

The SS layer first samples all attachments for tree extension in a batched manner, and generates a candidate goal state x_G based on the batch samples (lines 1-7). The attachment samples are rejected if the goal state candidate is not in collision. Detailed explanations of `SampleBatchAtt` will be introduced in IV-D. `MakeGoalCandidate` makes a goal state candidate $x_G = (s_G, \sigma_G, q_G)$, where

- s_G can be derived by applying action sequence π to the initial abstract state $x_I.s$;
- σ_G can be generated by applying mode changes by a sequence of attachments α to the initial mode $x_I.\sigma$;
- q_G can be found in G or randomly sampled.

If x_G is collision-free, attachments and goal states are used for the further tree extension process. Since sampling a goal candidate is computationally cheap, unachievable attempts can be rejected by validating the geometric feasibility of x_G .

After sampling the goal candidate, a batch extension of RT is conducted using π , n_s , and α . Fig. 3(b) shows the batch extension process of RT. First, a parent RT node x for expansion is sampled from the abstract state node n (line 12). If a is a non-geometric action, the new state x' is created by changing the abstract state s to s' , which is a non-geometric state transition (lines 16-17). If a is a geometric action, an adjacent mode σ' is generated by changing an attachment α from the previous mode σ (line 19), and the transition configuration q between σ and σ' is sampled (line 20). Note that this robot configuration sampling can be achieved by projecting a random configuration into a subspace induced by the two modes σ and σ' , i.e. numerical inverse kinematics.

The connection between the new node x' and the parent node x is checked through the MP layer (line 22). If connected, the pointer of x' is stored in the next abstract state node s' (line 23). If the connection fails, a reward signal is calculated (line 24). We define a reward as the total extension length from the root to the current node, which is scaled by the length of the action sequence to have a value between 0 and 1. Note that multiple rewards are received through the batch extension process, and therefore, the TP layer is updated using the average of received rewards.

If the last action of the action sequence π is successful, the SS layer finally checks a connection between the last extended state x' and the goal candidate x_G (lines 25-26). If the extension to the goal candidate x_G succeeds, a solution is found (lines 17-18). The path from the initial state x_I to the goal state x_G can be derived by backtracking the RT \mathcal{G} .

D. Attachment Batch Sampler (Alg. 4)

The attachment batch sampler samples all necessary attachments in the action sequence π at once. It is important to note that to complete the planning, the specific actions must select the corresponding attachments in the implicitly defined goal set G . We denote this action and attachment as *goal-making action* and *goal attachment*. Determining the goal-making actions can be achieved by examining the action sequence in reverse order and identifying the action that leads to the final mode. To accomplish this, a set \mathcal{M}_{goal} is employed to keep track of the movables whose final attachments are not determined.

This attachment assignment procedure ignores non-geometric actions, which do not require attachments (lines 5-7). If a is a geometric action, then it is necessary to determine its corresponding attachment. First, we obtain the target object m from action a . If there is no attachment assigned for the object m and if the action a has a corresponding goal

Algorithm 4: Attachment Batch Sampler

Input : Action sequence π , Implicitly defined goal set G
Output: A sequence of attachments α

```

1  $\alpha \leftarrow \text{Queue}()$ 
2  $\mathcal{M}_{goal} \leftarrow \mathcal{M}$ 
3 for  $i \leftarrow \pi.length$  to 1 do
4    $a \leftarrow \pi[i]$ 
5   if  $a \notin \mathcal{A}_g$  then
6      $\alpha.Push(null)$ 
7     continue
8    $m \leftarrow \text{TargetMovable}(a)$ 
9   if  $m \in \mathcal{M}_{goal}$  and  $\text{IsGoalAttachment}(a, G)$ 
10    then
11      $\alpha_m \leftarrow \text{GetGoalAttachment}(a, G)$ 
12    else
13      $\alpha_m \leftarrow \text{SampleAttachment}(a)$ 
14     $\alpha.Push(\alpha_m)$ 
15     $\mathcal{M}_{goal} \leftarrow \mathcal{M}_{goal} - \{m\}$ 
16  $\alpha \leftarrow \text{Reverse}(\alpha)$ 

```




Fig. 4. (a) Non-monotonic domain. (b) Blocktower domain.

attachment in the implicitly defined goal set G , then it can be concluded that the action a is the goal-making action. In this case, instead of sampling, we simply choose the goal attachment from the set G (lines 9-10). If the action is not directly related to the goal, the attachment is sampled based on the action (lines 11-12).

V. EMPIRICAL EVALUATION

This section validates the proposed method using three TAMP benchmark problems:

- *Kitchen m domain* [22]: We provide m number of movable food blocks and three placeable objects (i.e., dish, sink, and stove) as shown in Fig. 1. The task is to cook every food block, which is initially placed on the dish. A food block must be cleaned before being cooked. Therefore, there are two non-geometric actions `cook` and `wash` that can be executed when a food block is placed on the `stove` and `sink`, respectively. This domain has a tight geometric constraint, i.e., the area of the `sink` and `stove` is small, so the blocks can only be placed densely.
- *Non-monotonic m domain* [23]: We provide m blocks of different colors and m blockers with a higher height. The robot task is to place colored blocks on the floor which must have the same color as the blocks. As shown in Figure 4(a), a difficulty arises since the robot cannot reach the color blocks at once. Therefore, grey blockers

must be removed before approaching the color blocks. However, since the task planner is not aware of such geometry, so interaction between the TP layer and the lower-level layers is essential.

- *The blocktower m domain [24]:* We provide m blocks and 3 placeable plates. Initially, we stacked blocks in random order on the left or right plates, as shown in Fig. 4(b). The robot task is to stack them on the center plate in a certain order. To account for the task constraint where a movable object can be placed on another movable object, two abstract actions, namely, *stack* and *unstack*, are introduced.

For comparison, we implemented PDDLStream, which is the most well-known and readily available TAMP solver online. In addition, two baseline planners are designed:

- *No-reward planner:* Our planner uses the reward as a bias for better action sequence sampling. To verify the effectiveness of this, the reward is not used in this baseline.
- *No-rejection planner:* Our planner rejects sample batches by making a goal candidate and checking its geometrical feasibility. To verify the effectiveness of this process, the feasibility of the goal state is not checked in this baseline.

All planners including PDDLStream were written in Python. For the proposed method and baseline planners, the abstract domain is modeled in PDDL, and the task plan is solved by Pyperplan symbolic planner [25]. For the MP layer, our implementation of RRTCConnect [26] is used. Collision checking and inverse kinematics are implemented using Pybullet physics library [27]. All algorithms were executed on an Intel Core i7-10700 CPU, 16Gb RAM PC. The following hyperparameters were used in the proposed method including the baselines: $k_{ss} = 2$, $k_{goal} = 10$, and $\epsilon = 0.5$. For PDDLStream, *adaptive*, which is known as the best-performing implementation, was used with default parameters. All experiments were performed 30 times, and the timeout was set as 100 seconds.

A. Discussion

Simulation results are presented in Fig. 5 using a cumulative distribution function. In the kitchen domain, it is a better strategy to choose a simple task plan. For example, it is more advantageous to complete the wash-cook task in series, because *sink* has a tight goal region. Our planner showed a success rate of 97%, while the no-rejection baseline showed a success rate of 90%, and the PDDLStream and no-reward baseline showed a success rate of around 70%.

Comparing the performance of the two baselines, the no-reward baseline, which exploits the goal candidate sampling, performed better than the no-rejection baseline in the kitchen 3 and 4 domains. This is intuitively acceptable, as the sparsity of blocks in these domains makes it likely that a randomly sampled abstract plan will be feasible. In contrast, in kitchen 5, the no-rejection baseline, which is the informed planner by a previous reward, performed better. The proposed planner

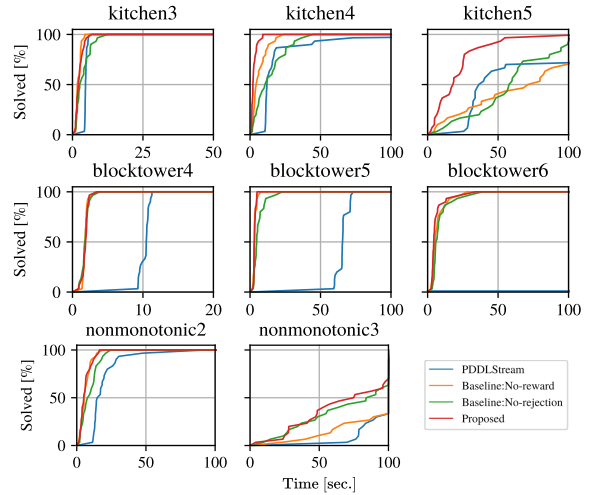


Fig. 5. The comparison result of the proposed planner, two baseline planners (no-reward, no-rejection), and PDDLStream. The results are expressed as a cumulative distribution function.

showed the highest performance for all scenarios by utilizing both features.

The effectiveness of rewards can also be seen in non-monotonic domains, in which the blocker blocks must be removed first. In the proposed method, the task planner learns through rewards that the action sequence for removing the blocker blocks is more promising. As a result, in non-monotonic domains, our proposed method showed the best performance with 66%. The rejection of using goal candidates did not lead to a significant improvement in performance in this domain.

Finally, in the blocktower domain, the proposed method including the baselines outperformed PDDLStream. This is because the proposed method solves a much simpler symbolic planning problem compared to the PDDLStream. In fact, [24] has already reported that the blocktower domain with more than 4 objects is challenging for PDDLStream. In the case of blocktower 6, PDDLStream could not find any solutions, while the proposed method showed a 90% success rate within 13 seconds. Compared to PDDLStream, our planner stores geometrical planning instances, such as grasp, placement, configuration, and trajectories in a reachability tree instead of including them in the task planning problem. In other words, our approach keeps the task planning problem manageable and prevents it from becoming more complicated by sampled instances over time.

VI. CONCLUSION

This paper presents a tree-based TAMP solver, inspired by reachability tree-based MMMP solvers. The proposed method samples the abstract action sequence in such a way that the tree expansion does not violate the task constraints defined in the abstract domain. The evaluation of the sampled sequence by the geometric planners allows us to bias the search towards more promising abstract state regions. In addition, a sampling rejection scheme that pre-generates a goal candidate can be used to resolve the tight goal constraints effectively. The comparative study using two baselines and PDDLStream showed the effectiveness of our approach.

REFERENCES

- [1] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez, "Integrated task and motion planning," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 4, pp. 265–293, 2021.
- [2] M. J. Kim, A. Werner, F. Loeffl, and C. Ott, "Passive Impedance Control of Robots With Viscoelastic Joints Via Inner-Loop Torque Control," *IEEE Trans. Robot.*, vol. 38, no. 1, pp. 584–598, 2021.
- [3] J. Jeong, H. Mishra, C. Ott, and M. J. Kim, "A Memory-based SO(3) Parameterization: Theory and Application to 6D Impedance Control with Radially Unbounded Potential Function," in *Proc. IEEE Int. Conf. Robot. & Automat.*, 2022, pp. 8338–8344.
- [4] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.
- [5] M. Ghallab, D. Nau, and P. Traverso, *Automated Planning: theory and practice*. Elsevier, 2004.
- [6] S. M. LaValle, "Rapidly-exploring random trees : A new tool for path planning," *Annu. Res. Rep.*, 1998.
- [7] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling, "PDDLStream: Integrating symbolic planners and blackbox samplers via optimistic adaptive planning," in *Proc. Int. Conf. Autom. Plan. Sched.*, vol. 30, 2020, pp. 440–448.
- [8] K. Hauser and J.-C. Latombe, "Multi-modal motion planning in non-expansive spaces," *Int. J. Robot. Res.*, vol. 29, no. 7, pp. 897–915, 2010.
- [9] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling, "Sampling-based methods for factored task and motion planning," *Int. J. Robot. Res.*, vol. 37, no. 13-14, pp. 1796–1825, 2018.
- [10] S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, and P. Abbeel, "Combined task and motion planning through an extensible planner-independent interface layer," in *Proc. IEEE Int. Conf. Robot. & Automat.*, 2014, pp. 639–646.
- [11] T. Lozano-Pérez and L. P. Kaelbling, "A constraint-based method for solving sequential manipulation planning problems," in *Proc. IEEE/RSJ Int. Conf. on Intell. Robot. Syst.*, 2014, pp. 3684–3691.
- [12] L. P. Kaelbling and T. Lozano-Pérez, "Hierarchical task and motion planning in the now," in *Proc. IEEE Int. Conf. Robot. & Automat.*, 2011, pp. 1470–1477.
- [13] N. T. Dantam, Z. K. Kingston, S. Chaudhuri, and L. E. Kavraki, "An incremental constraint-based framework for task and motion planning," *Int. J. Robot. Res.*, vol. 37, no. 10, pp. 1134–1151, 2018.
- [14] J. Mirabel and F. Lamiroux, "Manipulation planning: addressing the crossed foliation issue," in *Proc. IEEE Int. Conf. Robot. & Automat.*, 2017, pp. 4032–4037.
- [15] T. Siméon, J.-P. Laumond, J. Cortés, and A. Sahbani, "Manipulation planning with probabilistic roadmaps," *Int. J. Robot. Res.*, vol. 23, no. 7-8, pp. 729–746, 2004.
- [16] K. Hauser, T. Bretl, J.-C. Latombe, K. Harada, and B. Wilcox, "Motion planning for legged robots on varied terrain," *Int. J. Robot. Res.*, vol. 27, no. 11-12, pp. 1325–1349, 2008.
- [17] J. Barry, L. P. Kaelbling, and T. Lozano-Pérez, "A hierarchical approach to manipulation with diverse actions," in *Proc. IEEE Int. Conf. Robot. & Automat.*, 2013, pp. 1799–1806.
- [18] K. Hauser and V. Ng-Thow-Hing, "Randomized multi-modal motion planning for a humanoid robot manipulation task," *Int. J. Robot. Res.*, vol. 30, no. 6, pp. 678–698, 2011.
- [19] Z. Kingston and L. E. Kavraki, "Scaling Multimodal Planning: Using Experience and Informing Discrete Search," *IEEE Trans. Robot.*, 2022.
- [20] Z. Kingston, A. M. Wells, M. Moll, and L. E. Kavraki, "Informing multi-modal planning with synergistic discrete leads," in *Proc. IEEE Int. Conf. Robot. & Automat.*, 2020, pp. 3199–3205.
- [21] D. McDermott, M. Ghallab, A. Howe, C. Knoblock, A. Ram, M. Veloso, D. Weld, and D. Wilkins, "PDDL-the planning domain definition language," 1998.
- [22] T. Ren, G. Chaltatzaki, and J. Peters, "Extended tree search for robot task and motion planning," *arXiv preprint arXiv:2103.05456*, 2021.
- [23] F. Lagriffoul, N. T. Dantam, C. Garrett, A. Akbari, S. Srivastava, and L. E. Kavraki, "Platform-independent benchmarks for task and motion planning," *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3765–3772, 2018.
- [24] M. Khodeir, B. Agro, and F. Shkurti, "Learning to Search in Task and Motion Planning with Streams," *arXiv preprint arXiv:2111.13144*, 2021.
- [25] Y. Alkhazraji, M. Frorath, M. Gritznern, M. Helmert, T. Liebetraut, R. Matmüller, M. Ortlieb, J. Seipp, T. Springenberg, P. Stahl *et al.*, "Pyperplan," 2020, URL <https://doi.org/10.5281/zenodo.3700819>.
- [26] J. J. Kuffner and S. M. LaValle, "RRT-connect: An efficient approach to single-query path planning," in *Proc. IEEE Int. Conf. Robot. & Automat.*, vol. 2, 2000, pp. 995–1001.
- [27] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," 2016, URL <http://pybullet.org>.