

Reactive Task and Motion Planning under Temporal Logic Specifications

Shen Li^{†*}, Daehyung Park^{†‡*}, Yoonchang Sung^{†*}, Julie A. Shah[†], and Nicholas Roy[†]

Abstract—We present a task-and-motion planning (TAMP) algorithm robust against a human operator’s cooperative or adversarial interventions. Interventions often invalidate the current plan and require replanning on the fly. Replanning can be computationally expensive and often interrupts seamless task execution. We introduce a dynamically reconfigurable planning methodology with behavior tree-based control strategies toward reactive TAMP, which takes the advantage of previous plans and incremental graph search during temporal logic-based reactive synthesis. Our algorithm also shows efficient recovery functionalities that minimize the number of replanning steps. Finally, our algorithm produces a robust, efficient, and complete TAMP solution. Our experimental results show the algorithm results in superior manipulation performance in both simulated and real-world tasks.

I. INTRODUCTION

Consider the problem of collaborative operation between humans and robots [1]. A central challenge is responding in realtime to changes that might render the current robot task incomplete or inefficient. For example, we might like to ask a robot to “pack up all the objects into the box” as shown in Fig. 1. To achieve this complex goal, task-and-motion planning (TAMP) has been widely used for producing a sequence of actions [2]–[5]. During execution, a human partner might help the robot by packing up some of the objects or providing tools (e.g., a tray) to the robot. However, such changes may make the current plan invalid. The robot’s plan might not have accounted for the new tool. Operating alongside human partners may require replanning-and-execution on the fly.

A popular replanning-and-execution approach is reactive synthesis that finds control strategies to react to environmental model changes including object position [6]–[8], proposition [9], and uncertainty [10] while satisfying a given task specification [11], [12]. However, its reasoning process often requires the planner to construct a graph explicitly capturing feasible states (i.e., nodes) and edge connectivity between the states to find solutions. A wide range of model changes due to interventions leads to computationally expensive state-graph (re-)construction and search. Further, this graph reconstruction may lower execution efficiency, pausing ongoing motion until a new plan is reestablished.

We propose a reactive TAMP method that hierarchically reacts to various levels of task interventions (e.g., the




Fig. 1: A capture of our “3-block+tray” transfer task. Given a temporal logic-based goal specification $\mathcal{FG}(all_obj_in_r_2)$, our UR5 robot produces a complete task-and-motion plan to transfer objects from r_1 to r_2 . When a human operator provides a movable tray r_3 , our system reactively synthesizes a new plan that utilizes r_3 to produce efficient (i.e., minimum-distance cost) behaviors, and also shows seamless execution using a behavior tree.

position changes of objects, addition/deletion of objects) by introducing complementary layers of planning and execution. The planning layer includes complete and efficient reactive synthesis using linear temporal logic (LTL) suitable to formally specify complex objectives and interventions [13]–[16]. In detail, the method produces a sequence of action plans combining symbolic and (joint space) motion-cost based geometric searches over the LTL automata. Given environmental changes, our first contribution is to show how to find a new plan with minimal search effort by incrementally expanding the search space and leveraging experience such as previously explored paths and corresponding costs.

For the execution layer, we propose a behavior tree (BT) that conditionally activates sub-control policies (i.e., subtrees) to reactively follow action plans from the planning layer. BTs have been widely used to handle new or partially updated policies, as well as environmental changes [17]–[19]. We use the modularity of BTs, which allows easier composition, addition, and deletion of policies on the fly, to partially update policies without pausing ongoing executions. Our second contribution is to show that the recovery strategies of conditional BTs [7], [20] can minimize expensive replanning steps with LTL. By relaxing the activation conditions of trees, we maximize the recovery range and minimize replanning requests given environment changes such as object relocation.

In this work, we consider a robot that conducts a pick-and-place task of multiple objects where a human operator either cooperatively or adversarially could move the same objects. The robot can recognize the unexpected changes only by observing the location of objects. We assume that the robot does not suffer from partial and uncertain observations in this work. We do not explicitly model the human as another agent but rather consider as a part of the environment. We show our robot is able to adaptively change its plan on the

*These authors contributed equally to this work. [†]S. Li, D. Park, Y. Sung, J. Shah, and N. Roy are with CSAIL, Massachusetts Institute of Technology, USA. {shenli, daehyung, yooncs8, julie_a_shah, nickroy}@csail.mit.edu. [‡]D. Park is also with the School of Computing, Korea Advanced Institute of Science and Technology, Korea. This work was supported by Lockheed Martin Co. with the sponsor I.D. RPP2016-001 and the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. 2021R1C1C1004368).

fly to complete a given task. The key contributions include:

- We present a robust TAMP method that hierarchically reacts to various environment changes or new temporal logic specifications.
- Our planning layer provides quickly replanned solutions based on past experience, online graph extension, and reconfigurable BT structures.
- Our execution layer allows seamless plan execution even with environment changes and computation delays.
- We demonstrate that our method outperforms various baselines in terms of time efficiency and motion cost-effectiveness via physics-based simulations. We also conduct proof-of-concept experiments with a UR5 robot from Universal Robots.

II. TEMPORAL-LOGIC PLANNING

We present a brief introduction to the LTL planning framework (more details in [21]) and show how we define variables used in the framework.

The planning problem studied in this work is to move a set of objects from initial locations to desired goal locations while overcoming environmental changes that might occur unpredictably. To enable temporal reasoning and guarantee task completion, we adopt temporal logic-based planning to synthesize a plan for the robot that achieves LTL task specifications, owing to its correct-by-design nature.

We model the robot and its workspace as a transition system (TS), which can be defined as tuples as follows:

Definition 1. (Transition system) $TS = (S, s_0, A, \delta_t, \Pi, \mathcal{L})$, where S is a finite set of states ($s \in S$), s_0 is an initial state, A is a set of actions, $\delta_t : S \times A \rightarrow S$ is a deterministic transition relation, Π is a set of atomic propositions, and $\mathcal{L} : S \rightarrow 2^\Pi$ is a labeling function that maps each state to a subset of atomic propositions that hold true.

We compose an object o_i with a region r_j to denote a state s in TS, *i.e.*, $s = o_i r_j$, implying that the object o_i is located in the region r_j . Let $O = \{o_1, \dots, o_n | o_i \in \mathbb{R}^3\}$ be the set of n prehensile object states exist in the workspace of the robot, where o_i is the 3D positional vector of the i -th object. The objects are placed in m discrete and non-overlapping regions which we achieve by abstractly partitioning a continuous 3D workspace. We denote the set of m region states by $R = \{r_1, \dots, r_m | r_j \in \mathbb{R}^3\}$. The robot can observe which objects belong to which region through its perception. For multiple objects we represent a state s by splitting pairs of an object o_i and region r_j by underscore, *e.g.*, $s = o_1 r_1 _ o_2 r_3$. The cardinality of S is $|R|^{|O|}$. An element of the action set A is represented by $\text{move}_{o_i r_j}$ that is to pick an object o_i from its current location and place o_i in region r_j . An atomic proposition $\pi \in \Pi$ corresponds to a state $s \in S$ that is either true or false statements that must be specified for a given task. For example, $\pi = o_i r_j$ is true if an object o_i is in a region r_j , or false otherwise. \mathcal{L} maps a state $s \in S$ to a set of true atomic propositions $\mathcal{L}(s) \subseteq \Pi$.

An LTL formula φ is a combination of Boolean operators and temporal operators, which is involved with a set of

atomic propositions ($\subseteq \Pi$). The formula φ can express non-Markovian temporal properties, such as temporally-extended goals. The syntax of the formula φ over Π is as follows:

$$\varphi = \pi \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid X\varphi \mid \varphi_1 \mathcal{U}\varphi_2, \quad (1)$$

where we denote $\pi \in \Pi$, “negation” (\neg), “disjunction” (\vee), “next” (X), and “until” (\mathcal{U}). We can additionally derive a Boolean operator “conjunction” (\wedge) and temporal operators “always” (\mathcal{G}) and “eventually” (\mathcal{F}) from the basic syntax. A detailed explanation can be found in [21].

A word is an infinite sequence of atomic propositions in Π defined over the alphabet 2^Π . There exists a union of infinite words that satisfies a given formula φ . Alternatively, the nondeterministic Büchi automaton (BA) constructed according to [22] is generally used to check the satisfiability of φ .

Definition 2. (Büchi automaton) $BA = (Q_b, q_{b,0}, \Sigma, \delta_b, \mathcal{F}_b)$, where Q_b is a finite set of states ($q_b \in Q_b$); $q_{b,0}$ is an initial state; $\Sigma = 2^\Pi$ is an alphabet from the LTL formula; $\delta_b : Q_b \times \Sigma \rightarrow 2^{Q_b}$ is a non-deterministic transition relation; and $\mathcal{F}_b \subseteq Q_b$ is a set of accepting states.

We then generate a product automaton (PA) by producing the TS and BA as follows.

Definition 3. (Product automaton) $PA = TS \times BA = (Q_p, q_{p,0}, \Sigma, \delta_p, \mathcal{F}_p)$, where $Q_p = S \times Q_b$ whose element is $q_p = (s, q_b) \in Q_p$; $q_{p,0} = (s_0, q_{b,0})$; $\Sigma = 2^\Pi$; $\delta_p : Q_p \rightarrow Q_p$; and $\mathcal{F}_p \subseteq Q_p$. Note that δ_p satisfies $(s', q'_b) \in \delta_p(s, q_b)$ iff there exist $a \in A$ and δ_b such that $s' = \delta_t(s, a)$ and $q'_b \in \delta_b(q_b, \mathcal{L}(s))$.

PA becomes a search graph used for finding a plan. Finding a valid path that starts from an initial state $(s_0, q_{b,0})$ to one of accepting states $(s_F, q_{b,F})$ where $q_{b,F} \in \mathcal{F}_b$ implies computing a feasible plan $\xi = \{(s_i, q_{b,i})\}_{i=\{0,1,\dots,F\}}$ that achieves a given formula φ .

We study a scenario where unexpected human intervention affects the state of objects. A human operator might add new objects, remove existing objects, or move an object to a different region. This is equivalent to keeping the LTL specification φ as it is while dynamically modifying the TS and PA to reflect the changes. Previous works also considered the same reactive approach as ours [9], [10], or relaxed the LTL specification φ instead while keeping the TS [23], [24] for handling the changes. However, they mostly focused on navigation tasks.

The objective in this work is that the robot places all objects O to desired goal regions among regions R . Although human intervention occurs unpredictably, the robot must be able to correct TS and adjust its plan to accomplish a given task. Besides robustness, we would like the proposed replanning algorithm to additionally exhibit time efficiency and completeness.

III. DYNAMICALLY RECONFIGURABLE PLANNING

We address the proposed problem by directly finding a plan over the PA that satisfies an LTL formula φ , *i.e.*, placing all objects O in desired goal regions R . We achieve this by a hierarchical architecture (Fig. 2) that disentangles the high-level planning from low-level motor control commands. We

handle unexpected changes by human intervention in upper layers of the architecture.

The entire proposed hierarchical architecture consists of three layers: (1) The high-level planning layer computes a high-level plan that satisfies LTL specifications; (2) The middle-level execution layer dynamically constructs BT to make high-level plans executable for low-level controller of the robot; and (3) The low-level layer outputs control inputs used in actuators. The inputs to the architecture are a robot model, information about workspace, and LTL specifications φ and the output is low-level control inputs. Additional inputs are environmental changes that can be detected via perception.

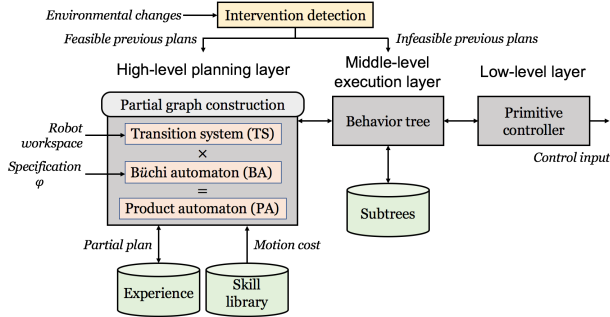


Fig. 2: Overall hierarchical architecture. The high-level layer contains the TS, BA, and PA that check the satisfiability of an LTL formula φ as well as external databases, *i.e.*, experience and skill library. The middle-level layer communicates between a high-level plan and a low-level controller via BT with a subtree database. The low-level layer outputs motor control commands to a robot.

Three desirable properties we aim at are robustness (against intervention), efficiency (in computing a plan), and completeness. We achieve robustness efficiently by (1) handling intervention without involving the expensive task planner and only calling the planner when necessary; (2) reusing partial plans from the past; (3) computing a cost-effective plan by incorporating a motion cost computed by inverse kinematics; and (4) maintaining a partial graph of PA. We achieve completeness by employing an LTL-based planning.

In Section III-A, we present the capability of each component in our system and the pseudo-code of the framework. In Section III-B, we introduce the search algorithm with motion cost to achieve a cost-effective solution. In Section III-C, we present two ways of speeding up the search algorithms: reusing partial plans and maintaining a partial graph of PA. In Section III-D, we explain the BT reconstruction in detail.

A. Robust Hierarchical System Design

In this subsection, we highlight which component of the hierarchical structure deals with which environmental changes. These components are in charge of different replanning queries that arrive at random time, including search graph construction and experience-based graph search in the high-level layer, and BT in the middle-level layer. In particular, we consider the following types of environmental changes: object relocation while previous plans are feasible or infeasible, and the addition and removal of objects and regions. Note again that the change in an LTL formula φ is not considered.

The process of the system handling environmental changes is illustrated in Alg. 1. In line 5, when the object is relocated

to one of existing regions, there are cases where the previous plans can be reused if the plans have been reserved as *subtrees* in the BT (details are explained in Section III-D). In particular, if there are still executable subtrees (line 6), the expensive task planner will not be queried. For example, as shown in Fig. 1, if the robot fails to pick up a block due to perception error, the BT will re-activate specific subtrees to attempt to grasp the block again.

Object relocation could also occur where previous plans are no longer feasible. To handle this case, our system will use the same graph (skipping line 13) and query a search algorithm to produce a new plan (line 16). Particularly, the search algorithm (Section III-B) improves its planning efficiency by reusing the planning experiences computed and cached from the previous queries and maintaining a partial graph (Section III-C). The BT will then receive the new plan and add/delete subtrees while keeping reusable subtrees for execution (line 17) (Section III-D).

When the number of objects and that of regions change, our system will reconstruct the TS and PA (line 13) before querying the experience-based graph search for planning (line 16) and send the new plan to the BT for execution (line 17) (Section III-D).

B. Planning for LTL specifications

We use a weighted directed graph $G = (V, E, w)$ to represent the PA and deploy a graph search algorithm for planning. In particular, each $q_p \in Q_p$ corresponds to a node $v \in V$. For each q_p and q'_p where $q'_p \in \delta_p(q_p)$, there is an edge $e \in E$ between the nodes associated with q_p and q'_p .

We incorporate the notion of motion cost into the PA as a weight w on an edge e of G . We compute the motion cost as the incremental displacement of a robot's end effector, which can be obtained from a skill library assumed to be given. For example, when a manipulator approaches some location, the motion cost can be computed by inverse kinematics and saved as one of skills in the library. This allows for cost-effective plans with respect to a cumulative motion cost.

To increase the planning efficiency, we propose two improvements over general graph search algorithms: (1) leveraging the computations from the past to reduce the computation in motion costs for a new replanning query; and (2) constructing G on the fly to avoid the overhead of generating a graph from scratch whenever the world changes.

C. Search Acceleration

We attempt to speed up the search algorithm by reusing previous plans and partially constructing a search graph. When searching for a path in G , computing a motion cost on edge is costly as it requires solving an inverse kinematics problem. We cache the motion cost of edges newly explored so far as well as corresponding node names as a table data structure. If a new plan during replanning visits a node that has been previously visited, we look up the motion cost table to avoid inverse kinematics recomputation. Reusing previous plans expedites the search as environmental changes often affect only a small part of G .

Algorithm 1: LTL-BT System. Efficiently handling various types of interferences and updating the subtrees Ψ in BT.

Input: BA, TS, Ψ

```

1  $s \leftarrow s_0$ ; // Initial state
2  $E \leftarrow \emptyset$ ; // Empty experience
3 while Environmental changes do
4    $s_{new} \leftarrow \text{GetStateFromPerception}()$ ;
5   if Objects relocated then
6      $i \leftarrow$  Find the subtree index  $i$  of
        $(A^i, C_{pre}^i, C_{post}^i) \in T$  such that  $C_{pre}^i \subset s_{new}$ ;
7     if  $i$  is found then
8       // Subtree  $i$  can resolve this
       // change by selecting action  $A^i$ 
       // No replanning is needed
9       continue;
10    end
11     $s \leftarrow s_{new}$ ;
12    if Objects added/removed or regions added then
13      // Search graph reconstruction
14      TS  $\leftarrow$  ReconstructTS(TS,  $s$ );
15    end
16     $q_{p,0} \leftarrow (s, q_{b,0})$ ; // Construct PA state
17    // Replan to find action tuples with
18    // partial graph construction, use
19    // cache experience (Sec. III-B, III-C)
20     $P, E \leftarrow \text{LTLPlanner}(q_{p,0}, \text{BA}, \text{TS}, E)$ ;
21    // Update subtrees in BT (Sec. III-D)
22     $\Psi \leftarrow \text{BTReconfiguration}(P, \Psi)$ ; // Alg. 2
23  end

```

Graph reconstruction is computationally expensive when the number of objects and that of regions vary. The complexity of fully reconstructing the PA grows exponentially as the number of new objects increases. The cardinality of Q_p is $|S| \times |Q_b|$ where $|S| = |R|^{|O|}$. To fully construct the graph G , the program needs to check for valid transitions for all pairs $\forall q_p, q'_p \in Q_p$, which results in $O(|Q_p|^2) = O(|R|^{2|O|} |Q_b|^2)$. Our system constructs G on the fly within the search (Alg. 1 line 16) to reduce this overhead, rather than fully constructing it before the search. With an admissible heuristic, our system can avoid expanding all $v \in V$ and avoid constructing the entire G while finding the same path when heuristics are not used. In particular, within the search algorithm (line 16), when a node associated with q_p is chosen to be expanded, the system will query δ_p to produce the list of next states $\{q'_p\} = \delta_p(q_p)$ and add them to V and E accordingly. In the worst case, if the heuristic is not effective at all, all nodes in V will be expanded, and G will be fully constructed. We call this approach *partial graph construction*. We show the efficiency of partial graph construction against the approach of fully constructing G every time before the search (*full graph construction*) in Section V.

D. Behavior Tree & Reconfiguration

The BT is a compositional decision-making system that can group a set of behaviors, execute between groups at a

fixed frequency, and change the structure on the fly. The execution starts from the root of the BT. It then gradually traverses the execution permission (the tick) from the root to child nodes. Once a node receives the tick, the node execution returns: 1) SUCCESS if the node execution was successfully completed, 2) RUNNING if the execution is not completed yet, 3) FAILURE if the execution was failed, and 4) INVALID otherwise. Fig. 3 shows the execution flow given a sequence of action plans, where \rightarrow and $?$ represent *sequence* and *selector* behaviors, respectively (see details in [25]).




Fig. 3: Subtree structures. Chooser is a custom selector that does not interrupt the low priority of a child node from the higher priority of nodes.

There are previous works [24], [26], [27] adopting both the LTL and BT for synthesizing controls from high-level plans. However, their methods cannot deal with environmental changes studied in this work.

We now describe how to build the internal structure of the BT given a high-level plan and reconfigure it on the fly.

1) *Automatic tree construction*: We convert a sequence of action plans P to a sequence of subtrees Ψ , which can be attached to the task root of a BT. Similar to the planning domain description language (PDDL [28]), we use preconditions and effects to achieve complex sequences of behaviors similar to [20], [27]. To build the tree, we parse a plan P into a sequence of action tuples $T = \{(A^1, C_{pre}^1, C_{post}^1), (A^2, C_{pre}^2, C_{post}^2), \dots\}$, where A^i , C_{pre}^i , and C_{post}^i represent an action, a pre-condition, and a post-condition, respectively. For example, given two consecutive states $(o_1r_1_o_2r_1, o_1r_2_o_2r_1)$, the action is o_1 's move from r_1 to r_2 , where $C_{pre} = \{o_1r_1, o_2r_1\}$ and $C_{post} = \{o_1r_2, o_2r_1\}$.

We then convert the tuple to a *state condition*-based subtree as shown in Fig. 3a, where C_R represents a set of running conditions (e.g., $o_1r_{hand_o_2r_1}$) from [7] and the *State Update* decorator forces the current state to be updated after the completion of an action. Note that C_{post}^i equals to C_{pre}^{i+1} in this sequential plan so that we can ignore post conditions. This structure can also re-execute the previous actions if the current state satisfies the precondition. However, the state-based condition may be too conservative in that a full state specification cannot ignore unrelated environmental changes.

Alternatively, we introduce a relaxed version of subtree constructor, the *action condition*-based subtree (see Fig. 3b), that is designed to prevent unnecessary replanning given unrelated or minor environmental changes such as object additions, deletions, and repositions. The subtree also allows more recovery behaviors. In particular, the precondition node checks soft conditions including the location of the target object, object source, and object destination (e.g., $C_{pre} = o_1r_1$). We also use a customized *selector*, called *chooser* [25], as a

Algorithm 2: BT Reconfiguration. P_A^i and Ψ_A^i represent the action name of the i -th action tuple and subtree, respectively. Likewise, P and A with the subscript C represent their preconditions. The superscript $-i$ accesses the i th element from last.

Output: Ψ

```

1  $P \leftarrow (P^1, \dots, P^n)$  // A seq. of action tuples
2  $\Psi \leftarrow (\Psi^1, \dots, \Psi^m)$  // A seq. of subtrees
3 for  $i \leftarrow 1$  to  $n$  do
4   if  $P_A^{-i} = \Psi_A^{-i}$  &  $|\Psi| \geq i$  then
5      $\Psi_C^{-i} \leftarrow P_C^{-i}$  // Update conditions
6   else if  $P_A^{-i} \neq \exists \Psi_A^{-i}$  then
7      $j \leftarrow$  find the index of  $P_A^{-i-1}$  at  $\Psi_A^{1:-i}$ 
8     Remove subtrees at  $\Psi^{j:-i}$ 
9   else
10    Add a subtree for  $P^{-i}$  at  $\Psi^{-i}$ 
11  end
12 end
13 if  $|\Psi| \neq n$  then
14   Remove  $\forall s \in \Psi^{1:|\Psi|-n}$ 
15 end

```

root of each subtree that is to prevent interruption from higher priority of sibling subtrees. This enables BT to complete an ongoing action.

2) *Online reconfiguration:* We introduce how we can reconfigure a BT without interfering with running tasks (*i.e.*, running subtrees) on the fly. Given environmental changes, the planner will produce a new plan P' , which may not be completely different from the previous plan P so it is not necessary to reconstruct a sequence of subtrees S from scratch. Alg. 2 shows how we can find sharable subtrees between P and P' from the back. If there is a new action in P' , we add a new subtree to S . Otherwise, we remove the unmatched subtrees. Note that if there are any running and non-interruptible subtrees, our framework postpones the change of the subtree until the return of the subtree is changed.

IV. EXPERIMENTAL SETUP

A. Statistical Evaluation

We evaluated our method with baselines on two environment settings in a physics-based simulator, GAZEBO (Ver. 7) [29]. The manipulation settings were composed of colored blocks, trays, and/or baskets (see Fig. 5). Given an LTL specification, $\varphi = \mathcal{FG}(all_obj_in_r_2)$, our method generated a plan for transferring all blocks from r_1 to r_2 . Note that other LTL specifications can also be handled in our framework. Per each method and setting, we created 30 test scenarios by randomly placing objects in the beginning and then introduced an additional environmental change, such as repositioning, adding, or deleting a randomly selected object at a random time. Descriptions of each scenario and setting are as follow: 1) In the *3-block* scenario, as shown in Fig. 5(a), there are three blocks initially located at r_1 . The desired plan is to move each block one after another from r_1 to r_2 . 2) In the *3-block + tray* scenario, as shown in Fig. 5(b), there are three blocks initially located at r_1 and a movable tray r_3 near r_1 . The optimal plan is to transfer all the blocks to r_3




Fig. 4: Comparison of on/offline BT with *action* and *state* conditions. first, move r_3 near r_2 , and then relocate blocks from r_3 to r_2 . By leveraging r_3 , the robot can reduce the total motion costs of transferring blocks from r_1 to r_2 .


B. Demonstration Setup

We demonstrated the proposed method on two manipulation scenarios of a real UR5 robot (see Fig. 5). The robot observed objects using an Intel D435 RGB-D camera mounted on the wrist. It first recognized object segments by using *Detectron2* [30] with *Mask R-CNN* [31] from Facebook, where we trained *Detectron2* images collected from similar settings after labeling it via *COCO Annotator* [32]. The segments were then used for estimating the 3D pose and grasping point of each object from the depth image. Here, we assumed there are unique colors of objects to assign unique ID. An operator randomly relocated, added, or removed objects including a tray during the experiment.


V. RESULTS

We benchmarked our method against a set of baselines and demonstrated its robustness and efficiency via simulations. To show the performance of the execution layer, we compared three variations of the BT: *online* reconfiguration with *action condition*, *online* reconfiguration with *state condition*, and *offline* reconstruction with *action condition* over two scenarios. Note that the offline reconstruction stops a running node to remove all subtrees. It then constructs a new tree from scratch in an idle mode. Fig. 4 shows the *online* reconfiguration with *action condition* resulted in a smaller number of BT changes and took less completion time compared with that of *state condition*. Due to the short replanning time of A* with experience, the online method with *action* and *state* conditions resulted in a similar completion time in the 3-Block scenario. On the other hand, the *offline* reconstruction incurred longer average completion time than *online* BT. Overall, the results demonstrate that the *online* BT reconfiguration with *action condition* is the most efficient among the variations.

We benchmarked three search algorithms, A* with experience, A*, and Dijkstra, in the 3-block + tray scenario, each with three types of environmental changes (repositioning, adding, or deleting a block). We split all planning processes in each trial into (1) initial planning in the very beginning and (2) replannings, used for intervention handling. Table I highlights the robustness and efficiency of our hierarchical system. The column “# replan” presents that both means and medians of the number of replannings for object repositions are lower than those for additions and deletions across all algorithms. This result demonstrates that the system’s



(a) A “3-block” transfer task. An operator added and then repositioned objects during the task.



(b) A “3-block + tray” task. An operator provided a movable tray during the task.

Fig. 5: Demonstrations on the 3-block and 3-block+tray scenarios, where a UR5 robot is commended to place all objects on the right place. A human operator artificially generated environmental changes during the object transfer tasks. The robot with the proposed TAMP method successfully completed the task by producing a reactive and cost-effective solution. Please see the attached video: <https://youtu.be/lPpMVfBzZH0>.

efficiency in handling repositions by using the BT alone without the expensive replanning unless necessary. The mean of the number of replannings for object reposition is non-zero because sometimes it is necessary to replan. For example, if o_i is repositioned to r_1 when r_3 is near r_2 , since the state $o_i r_1 r_3 r_2$ is not in the optimal initial plan, no subtrees in BT can handle this case alone without replanning. Table I highlights the efficiency of the experience-based search. A* with experience leverages the similarity between the initial plan and the plans for intervention handling. In Table I, A* with experience has a significantly lower total replanning time than A* and Dijkstra among all types of environmental changes. The initial planning time is almost the same across three algorithms due to a lack of experience in the very beginning. To demonstrate the overall efficiency, we report the completion time of each trial. Among all types of environmental changes, A* with experience achieves significantly shorter completion time due to faster replanning. The success rate is always 100% in Table I as our system guarantees task completion.

3-block + tray	Success	Init plan time (s)	Total replan time (s)	# replan	Complete time (s)
Rel, A*, exp	30/30	68.39 ± 1.00	0.69 ± 0.43	0.80, 0	466.50 ± 34.80
Rel, A*	30/30	67.19 ± 0.04	33.50 ± 15.26	0.63, 0	530.17 ± 19.04
Rel, Dijkstra	30/30	67.36 ± 0.04	30.40 ± 13.06	0.47, 0	561.57 ± 27.29
Rem, A*, exp	30/30	67.62 ± 0.22	0.10 ± 0.02	1.93, 1	394.48 ± 13.14
Rem, A*	30/30	67.34 ± 0.07	52.57 ± 8.18	1.30, 1	536.25 ± 31.49
Rem, Dijkstra	30/30	67.57 ± 0.08	60.53 ± 9.16	1.40, 1	478.43 ± 17.54
Add, A*, exp	30/30	67.78 ± 0.32	20.61 ± 0.86	1.37, 1	647.00 ± 24.67
Add, A*	30/30	69.16 ± 1.29	89.60 ± 8.73	1.13, 1	734.74 ± 33.32
Add, Dijkstra	30/30	67.60 ± 0.10	116.16 ± 16.35	1.33, 1	720.81 ± 25.03

TABLE I: Results of statistical evaluation regarding 3 search algorithms (A* with planning experiences, A*, and Dijkstra), with partial graph construction, in handling 3 types of environmental changes (relocating 1 object, removing 1 object, and adding 1 object). The metrics used here are number of successes/number of trials, initial planning time (95% confidence interval (CI)), total replanning time (95% CI), number of replannings (mean, median), and completion time (95% CI).

We validated the efficiency of the partial graph construction over the full graph construction through simulations. We measured the time taken for both constructing the PA and replanning when an object was added as environmental changes for efficiency comparison. We fixed the number of regions to be five while varying the number of states in the PA to vary from 50 to 31,250. We randomly generated locations of regions ten times. Fig. 6 shows the mean and the standard deviation of a total time measured over ten random

environments for both graph construction methods. The full graph construction performed comparably with the partial graph construction up to 1,250 states but started degrading exponentially afterward.




Fig. 6: Comparison of time taken for constructing the PA and replanning between the full graph construction and the partial graph construction.

We finally demonstrated the proposed TAMP method with a real UR5 robot. Fig. 5 shows captures of two representative intervention scenarios. Given the same LTL specification, our robot successfully initiated a complete task and motion plan to transfer all objects to another storage by visually recognizing the environment. When an operator added and relocated objects, the online BT enabled the robot to continue seamless task execution while replanning the new setup. The proposed graph construction and experience-based planning enabled the robot quickly to transition to a new plan. Given a new transfer tool (see Fig. 5b), our system could also show cost-effective behaviors by utilizing the tray as we humans do. These demonstrations show our framework enables the robot to achieve robust and efficient high-level planning and middle-level execution in human-robot operation scenarios.

VI. DISCUSSION & CONCLUSION

We introduce a hierarchically reactive TAMP method, which generates a complete and semi-optimal task plan by using LTL and approximated motion costs. Our framework provides an efficient replanning approach leveraging experience-based graph search and node modification algorithms. In addition, to minimize the replanning and also maximize the recovery control capability, we combine it with a dynamically reconfigurable BT with the *action condition* subtree structure. Our statistical evaluation shows the superior efficiency compared to other baseline methods. We also demonstrate the robustness in real world with a UR5 robot.

REFERENCES

- [1] A. Ajoudani, A. M. Zanchettin, S. Ivaldi, A. Albu-Schäffer, K. Kose, and O. Khatib, "Progress and prospects of the human-robot collaboration," *Autonomous Robots*, vol. 42, no. 5, pp. 957–975, 2018.
- [2] N. T. Dantam, Z. K. Kingston, S. Chaudhuri, and L. E. Kavraki, "Incremental task and motion planning: A constraint-based approach," in *Proceedings of Robotics: Science and Systems (RSS)*, vol. 12, p. 00052, Ann Arbor, MI, USA, 2016.
- [3] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling, "Sampling-based methods for factored task and motion planning," *International Journal of Robotics Research*, vol. 37, no. 13-14, pp. 1796–1825, 2018.
- [4] W. Vega-Brown and N. Roy, "Asymptotically optimal planning under piecewise-analytic constraints," in *Algorithmic Foundations of Robotics XII*, pp. 528–543, Springer, 2020.
- [5] 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 *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 639–646, IEEE, 2014.
- [6] K. He, A. M. Wells, L. E. Kavraki, and M. Y. Vardi, "Efficient symbolic reactive synthesis for finite-horizon tasks," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8993–8999, IEEE, 2019.
- [7] C. Paxton, N. Ratliff, C. Eppner, and D. Fox, "Representing robot task plans as robust logical-dynamical systems," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5588–5595, IEEE, 2019.
- [8] T. Migimatsu and J. Bohg, "Object-centric task and motion planning in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 844–851, 2020.
- [9] M. Guo, K. H. Johansson, and D. V. Dimarogonas, "Revising motion planning under linear temporal logic specifications in partially known workspaces," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5025–5032, IEEE, 2013.
- [10] S. C. Livingston, R. M. Murray, and J. W. Burdick, "Backtracking temporal logic synthesis for uncertain environments," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5163–5170, IEEE, 2012.
- [11] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.
- [12] E. M. Wolff, U. Topcu, and R. M. Murray, "Efficient reactive controller synthesis for a fragment of linear temporal logic," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5033–5040, IEEE, 2013.
- [13] S. Chinchali, S. C. Livingston, U. Topcu, J. W. Burdick, and R. M. Murray, "Towards formal synthesis of reactive controllers for dexterous robotic manipulation," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5183–5189, IEEE, 2012.
- [14] K. He, M. Lahijanian, L. E. Kavraki, and M. Y. Vardi, "Towards manipulation planning with temporal logic specifications," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 346–352, IEEE, 2015.
- [15] E. Plaku and S. Karaman, "Motion planning with temporal-logic specifications: Progress and challenges," *AI communications*, vol. 29, no. 1, pp. 151–162, 2016.
- [16] C. K. Verginis and D. V. Dimarogonas, "Timed abstractions for distributed cooperative manipulation," *Autonomous Robots*, vol. 42, no. 4, pp. 781–799, 2018.
- [17] I. Millington and J. Funge, *Artificial intelligence for games*. CRC Press, 2009.
- [18] A. Marzinotto, M. Colledanchise, C. Smith, and P. Ögren, "Towards a unified behavior trees framework for robot control," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5420–5427, IEEE, 2014.
- [19] M. Colledanchise and P. Ögren, *Behavior trees in robotics and AI: An introduction*. CRC Press, 2018.
- [20] E. Giunchiglia, M. Colledanchise, L. Natale, and A. Tacchella, "Conditional behavior trees: Definition, executability, and applications," in *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC)*, pp. 1899–1906, IEEE, 2019.
- [21] H. Kress-Gazit, M. Lahijanian, and V. Raman, "Synthesis for robots: Guarantees and feedback for robot behavior," *Annual Review of Control, Robotics, and Autonomous Systems*, 2018.
- [22] C. Baier and J.-P. Katoen, *Principles of model checking*. MIT press, 2008.
- [23] M. Guo and D. V. Dimarogonas, "Reconfiguration in motion planning of single- and multi-agent systems under infeasible local ltl specifications," in *Proceedings of the IEEE Conference on Decision and Control*, pp. 2758–2763, IEEE, 2013.
- [24] J. Tumova, A. Marzinotto, D. V. Dimarogonas, and D. Kragic, "Maximally satisfying ltl action planning," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1503–1510, IEEE, 2014.
- [25] "Py trees." <http://py-trees.readthedocs.io>. Accessed: 2020-07-10.
- [26] M. Colledanchise, R. M. Murray, and P. Ögren, "Synthesis of correct-by-construction behavior trees," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6039–6046, IEEE, 2017.
- [27] M. Lan, S. Lai, T. H. Lee, and B. M. Chen, "Autonomous task planning and acting for micro aerial vehicles," in *Proceedings of the IEEE International Conference on Control and Automation (ICCA)*, pp. 738–745, IEEE, 2019.
- [28] M. Fox and D. Long, "Pddl2. 1: An extension to pddl for expressing temporal planning domains," *Journal of artificial intelligence research*, vol. 20, pp. 61–124, 2003.
- [29] N. Koenig and A. Howard, "Design and use paradigms for gazebo, an open-source multi-robot simulator," in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, vol. 3, pp. 2149–2154, IEEE, 2004.
- [30] Y. Wu, A. Kirillov, F. Massa, W.-Y. Lo, and R. Girshick, "Detectron2." <https://github.com/facebookresearch/detectron2>, 2019.
- [31] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proceedings of the International Conference on Computer Vision (ICCV)*, pp. 2961–2969, 2017.
- [32] J. Brooks, "COCO Annotator." <https://github.com/jsbroks/coco-annotator/>, 2019.