

# Sampling-based path planning under temporal logic constraints with real-time adaptation

Yizhou Chen, Ruoyu Wang, Xinyi Wang, and Ben M. Chen

**Abstract**—Replanning in temporal logic tasks is extremely difficult during the online execution of robots. This study introduces an effective path planner that computes solutions for temporal logic goals and instantly adapts to non-static and partially unknown environments. Given prior knowledge and a task specification, the planner first identifies an initial feasible solution by growing a sampling-based search tree. While carrying out the computed plan, the robot maintains a solution library to continuously enhance the unfinished part of the plan and store backup plans. The planner updates existing plans when meeting unexpected obstacles or recognizing flaws in prior knowledge. Upon a high-level path is obtained, a trajectory generator tracks the path by dividing it into segments of motion primitives. Our planner is integrated into an autonomous mobile robot system, further deployed on a multicopter with limited onboard processing power. In simulation and real-world experiments, our planner is demonstrated to swiftly and effectively adjust to environmental uncertainties.

## I. INTRODUCTION

In classical motion planning problems, a robot is commanded to generate a trajectory to reach a target configuration without collision [1], [2], [3]. Some works extend the motion planning problems to goals expressed in terms of temporal logic [4], [5], [6]. They consider tasks with complex temporal and spatial constraints that can be expressed in a form closely linked to natural language. Linear temporal logic (LTL) is widely used to specify such tasks since it is a mathematically precise language with sufficient expressiveness [7]. Hybrid robot controllers [8] can be computed to satisfy tasks expressed in LTL formulae.

In classical controller synthesis approaches [9], [10], [11], abstraction on the environment is created a priori which simplifies the motion of robots into symbolic transitions. A discrete plan can be retrieved from a product automaton constructed by taking the product of the transition system and the task-related automata. However, most existing works assume the environment is static, and its finite model can be obtained in advance. In the real world, the model of the environment is not always available and is even unpredictable due to shifting behavioral patterns. Directly resynthesis of the whole plan can lead to undesired consequences since the execution history will be lost, and the revised plan may violate the temporal specification. An example of a task is “Pick up a box in the warehouse first, then drop it off at

the office before picking up another box.” If the robot fails to locate the drop zone and decides to resynthesis the plan from scratch, the new plan will have the robot pick up a box again, which is prohibited by the task formula.

In case the previous plan is unachievable, some replanning methods are developed [12], [13], [14] in recent years based on controller synthesis approaches. The work in [12] proposes a plan revising mechanism assuming the environment is partially known and static. Once one transition in the current plan is broken, it seeks ways to find the shortest path bridging up the two components again. Authors in [13] develop a method to locally “patch” the defunct parts caused by the environmental change and identify new solutions. The work [14] develops an iterative repair strategy to resolve unexpected obstacles by combining local patching with refined triangulation. However, the above-mentioned repair operations cannot deal with both propositional change and moving obstacles, and replanning on large product automaton is

Unlike the abstraction-dependent methods above, our approach uses a sampling tree that explores the workspace for solutions. The closest related work to this paper is [15], which presents an abstraction-free method that incrementally builds trees to search for a satisfying solution. Note that it is an offline planner with no consideration of uncertainties. A lot of researchers, either dealing with classical path planning problems [16], [17] or temporal logic tasks [15], put much effort into improving sample efficiency in sampling-based planning. These methods invent various heuristics to speed the convergence of the final solution to optimality. Some work [18], [19] implement receding-horizon sampling trees for interleaved planning and acting. However, the executed edges are abandoned and cannot be used in future planning. In [20], the authors develop a multi-query sampling-based path planner by preserving all nodes sampled. Inspired by the approach above, we reuse the sampled points connected to the historical path for real-time replanning.

We propose a novel sampling-based planning and replanning paradigm for temporal logic tasks. The contributions of this work are threefold:

- A temporal logic constrained path planner that rapidly responds to real-time workspace knowledge updates using a dual-root tree;
- A graph-theoretic approach that reuses sampled nodes and resolves state duplication during replanning;
- Integration of algorithms with an online robotic perception-and-planning system and demonstration of its advantages in simulation and real-world experiments.

The work was supported in part by the Research Grants Council of Hong Kong SAR under Grants 14209020 and 14206821, and in part by the Hong Kong Centre for Logistics Robotics. Authors are with the Chinese University of Hong Kong, Shatin, N.T., Hong Kong 999077. (Email: josephchen@link.cuhk.edu.hk; rywang@link.cuhk.edu.hk; xywang-mae@link.cuhk.edu.hk; bmchen@cuhk.edu.hk).

## II. PRELIMINARIES

### A. Linear Temporal Logic on Finite Traces

We restrict the task specification in the form of Linear Temporal Logic on finite traces ( $LTL_f$ ), which has sufficient expressiveness to command robots with constrained power supply. Let  $AP$  be a set of atomic propositions where  $\alpha \in AP$  is a Boolean variable. We consider  $LTL_f$  formulas whose syntax is as follows:

$$\varphi := \text{true} \mid \alpha \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond\varphi \mid \square\varphi \mid \varphi\mathcal{U}\varphi$$

where  $\neg$  (negation),  $\vee$  (disjunction),  $\wedge$  (conjunction) are boolean operators, and  $\diamond$  (“eventually”),  $\square$  (“always”) and  $\mathcal{U}$  (“until”) are temporal operators. Note that we exclude the “next” operator due to the continuous execution without instantaneous motion in our abstraction-free approach. We refer readers to the work [21] for the semantics of  $LTL_f$ .

For every  $LTL_f$  formula  $\varphi$  there is a deterministic finite automaton (DFA) that accepts a sequence of input words to satisfy  $\varphi$ . DFA is formally defined by a five-tuple  $\mathcal{A}_\varphi = (Q, 2^{AP}, \delta, Q_0, \mathcal{F})$ , where  $Q$  is a finite set of automata states;  $Q_0 \subseteq Q$  is a set of initial states;  $2^{AP}$  is the input alphabet;  $\delta: Q \times 2^{AP} \rightarrow Q$  is the transition relation;  $\mathcal{F} \subseteq Q$  is a set of accepting automata states.

A run of a finite length  $n$  on DFA is a sequence of states  $r = q_0q_1 \dots q_n$ , where  $q_0 \in Q_0$  and  $q_{k+1} \in \delta(q_k, \lambda)$  for some  $\lambda \in 2^{AP}$ ,  $k \in [0, n-1]$ , which implies the run starts at an initial state and follows the transition relation. A run  $r$  is called an *accepting run* if  $q_n \in \mathcal{F}$ .

### B. Robot Model

In this paper, we model a mobile robot as a hybrid system with state space  $S = Q \times X$ , where  $Q$  is the set of discrete modes and  $X$  is the set of continuous states. A *hybrid state* is defined as a pair  $s = (q, x) \in S$ , and its projection on  $X$  is denoted as  $s|_X = x$ . Considering no external disturbance caused by the environment, the evolution of a robot can be achieved by applying a motion primitive  $u$  during time  $t$  at a state in the domain  $x \in X$ , which can be written as  $\zeta(x, u, t)$ . A finite *path* is a finite sequence of hybrid states  $\tau = s_0s_1 \dots s_n$ , and its projection  $\tau|_X$  is resulted from a series of motion primitives  $\zeta(x_0, u_0, t_0)\zeta(x_1, u_1, t_1) \dots \zeta(x_n, u_n, t_n)$ . The corresponding *trace* of a path  $\tau$  is  $\sigma = L(x_0)L(x_1) \dots L(x_n)$ , where  $L: X \rightarrow 2^{AP}$  is a labeling function that returns the atomic propositions that are true in the state  $x \in X$ . We define that a path  $\tau$  satisfies an  $LTL_f$  formula  $\varphi$  if and only if the input words of the corresponding trace lead to an acceptance state, formally written as  $\exists q_f = \delta(q_n, L(x_n)), q_f \in \mathcal{F} \Leftrightarrow \tau \models \varphi$ .

Follow [14], an *event-driven trace* of a path is defined as

$$\bar{\sigma} = L(x_0)L(x_{i_0})L(x_{i_0+i_1}) \dots L(x_{i_0+\dots+i_{l-1}}),$$

for  $i_0 + \dots + i_{l-1} < n$  by removing states with repetitive discrete modes in a path. In another word, the *event-driven trace* captures the label change in the trajectory and  $L(x_{I_j}) \neq L(x_{I_{j+1}})$  where  $I_j = \sum_{k=0}^j i_k$ . In practice, there may exist multiple paths leading to the acceptance of a DFA with different *event-driven traces*.

## III. PROBLEM DESCRIPTION AND PROPOSED APPROACH

### A. Problem Description

This paper considers a mobile robot with sensory capability moving in a partially unknown workspace to accomplish high-level goals. The mobile robot can either be a differential drive ground vehicle or a multicopter, which can remain stationary and change its heading. Meanwhile, the robot is not collision resilient and is forced to stay away from obstacles by the planner. Therefore, collision avoidance constraints are excluded from formulae in the following sections. The workspace  $X$  consists of a set of regions of interest  $\Pi$  whose properties are the combinations of atomic propositions  $AP$ , and a set of regions occupied by fixed obstacles  $X_{\text{occ}}$ . The robot is given the workspace model  $X_{\text{init}}$  as prior knowledge. We assume  $X_{\text{init}}$  suffices the robot to compute a trajectory within obstacle-free areas  $X_{\text{free}} = X_{\text{init}} \setminus X_{\text{occ}}$  that satisfies the specification  $\varphi$ . In contrast to the work [14], we relaxed the assumption that the propositional regions are fully-known. Due to the incomplete information given in advance, satisfaction is not guaranteed since the actual region property might not coincide with that in prior knowledge. The workspace could be updated by receiving knowledge from either onboard sensors or external observers.

The problem that this paper focused on is stated as follows:

**Problem 3.1:** Given a partially unknown workspace  $X_{\text{init}}$  with dynamic obstacles, a hybrid robot model whose state space is  $S$ , and a task specification  $\varphi$  expressed in  $LTL_f$ , find a path  $\tau$  that satisfies  $\varphi$  with adaptation to any admissible uncertainties in real-time.

We suppose that the uncertainty in the environment is admissible in the sense that the propositions of each region do not vary dynamically, and there always exists a collision-free finite path that results in the task’s acceptance. Once detected property changes or dynamic obstacles, the robot is supposed to be able to brake and avoid reaching/colliding with them in time while satisfying kinodynamic constraints. We assume no perception error within a sensing area, while actuation uncertainties are taken care of by feedback motion planning at a lower level.

### B. Proposed Approach

To address Problem 3.1, we propose a path planning framework whose diagram can be viewed in Fig. 1. The planning part can be divided into two phases: initial planning and reactive planning.

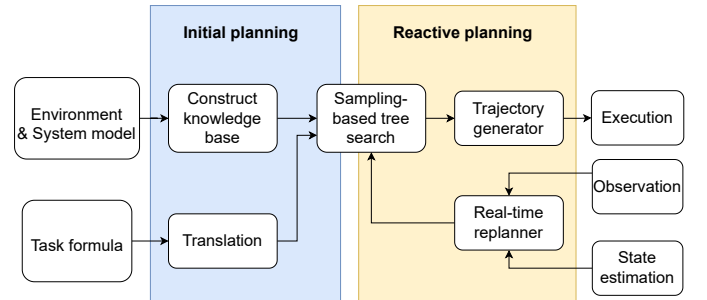


Fig. 1: The framework of real-time temporal logic planning

Prior to planning, the system is provided with an environment model, a robot model, and a task formula. The robot model helps grow a sampling-based search tree  $\mathcal{T}$  in  $X_{\text{free}}$  progressively with kinetically compatible edges. The task specified in  $LTL_f$  as  $\varphi$  is translated into a DFA and is further converted to a state monitor that guides the sampling-tree expansion. A knowledge base of the workspace is constructed based on the environment model with labeled regions  $\Pi$  and obstacles  $X_{\text{occ}}$ . The knowledge base serves as an initial map  $X_{\text{init}}$  in the initial planning to obtain an initial solution whose trace satisfies the task formula  $\varphi$ , as to be explained in Section IV-B.

In the online reactive planning part, once a solution is successfully obtained, the robot can track the nominal path safely with the help of a trajectory generator. During execution, the sampling-based planner continuously searches for a distinctive solution and improves the existing solutions, as will be described in Section IV-C. The state estimation helps the planner keep track of the execution progress and updates the robot's continuous state  $x$ . In the meantime, the robot receives observations from both external observers and its onboard sensors and revises the corresponding part of the knowledge base. The evolution of both the environment state and the robot state can trigger the replanning mechanism to deal with the non-determinism in the domain, as will be discussed in Section IV-D and Section IV-E.

#### IV. SOLUTION OF THE PROBLEM

##### A. Automata preprocess and temporally forbidden zones

To obtain a minimized automaton without ambiguity, we use MONA [22] and  $LTL_f2DFA$  [23] to translate an  $LTL_f$  formula to a DFA  $\mathcal{A}_\varphi$ . Chances are that, given a prefix of a run  $q_0q_1 \dots q_i$  on DFA, there does not exist a sequence  $q_iq_{i+1} \dots q_n$  where  $q_{k+1} \in \delta(q_k, \lambda)$  for any  $\lambda \in 2^{AP}$ ,  $k \in [i, n-1]$  and  $q_n \in \mathcal{F}$ , i.e., the accepting states cannot be reached from the automata state  $q_i$  on the DFA. To overcome this issue, we apply a preprocessing routine on DFA before planning. In detail, we construct an auxiliary automaton that reverses all transitions of the original DFA and finds the sets of automata states  $\mathcal{B}$  that are not reachable from any state in  $\mathcal{F}$ . During planning, the planner disables any transition relation  $\delta(q_k, \lambda)$  that can have the motion tree to include a node with an automata state  $q_{k+1} \in \mathcal{B}$ .

During execution, to ensure the states in  $\mathcal{B}$  are never reached, temporally forbidden zones  $\Pi_{\text{fbd}}$  are extracted and serve as time-variant hard constraints in path tracking. When the robot's state is  $s_{\text{cur}} = (q_{\text{cur}}, x_{\text{cur}})$ , the planner obtains the next discrete mode as  $q_{\text{next}} = \delta(q_{\text{cur}}, L(x_{\text{cur}}))$  due to the determinism of DFA. The planner examines each labeled region  $\pi \in \Pi$ , and tries to intake the corresponding label  $L(\pi)$ . If the input word transitions  $q_{\text{next}}$  toward a state in  $\mathcal{B}$ , the corresponding region will be added into  $\Pi_{\text{fbd}}$ .

##### B. Initial plan computation

After constructing the initial map  $X_{\text{init}}$  and translating the  $LTL_f$  specification, the robot builds a motion tree toward the temporal goal. Similar to anytime RRT\* [18], the robot does

not wait till a near-optimal solution is calculated to execute. Instead, once a feasible solution is found, the robot begins to follow the nominal path with a trajectory generator.

Like the work [15], we employ a sampling-based tree search method monitored by an automaton. The search tree  $\mathcal{T}$  is composed of node set  $\mathcal{V}_{\mathcal{T}}$  and edge set  $\mathcal{E}_{\mathcal{T}}$ . Additionally, we maintain a graph  $\mathcal{G} = (\mathcal{T}, \mathcal{V}_{\text{iso}})$  containing not only the search tree but also a set of isolated vertices.

Every time a sample is drawn on a state  $x_{\text{rand}} \in X_{\text{free}}$ , the nearest node in the tree  $s_{\text{nearest}} = (q_{\text{nearest}}, x_{\text{nearest}})$  steers to  $x_{\text{rand}}$  with robot dynamics in time  $\Delta t$  and ends at  $x_{\text{new}}$ . By pairing with possible automata states, a set of hybrid states  $S_{\text{new}} = \{(q, x_{\text{new}}) \mid q \in Q \setminus \mathcal{B}\}$  whose continuous states are the same but automata states are different is generated.  $S_{\text{new}}$  is recorded in each new vertex  $s_{\text{new}} \in S_{\text{new}}$  for resolving state duplication and will be referred to as  $S^{\text{red}}$  in Section IV-D. Afterward, the planner tries to add the hybrid states as nodes to the tree. The set of vertices that is near  $x_{\text{new}}$  is  $S_{\text{near}} = \{s = (q, x) \in \mathcal{V}_{\mathcal{T}} \cup \mathcal{V}_{\text{iso}} \mid \|x - x_{\text{new}}\| \leq r_n(\mathcal{V}_{\mathcal{T}})\}$ , where  $r_n$  is a radius calculation function as defined in [15]. For each  $s_{\text{new}} = (q_{\text{new}}, x_{\text{new}})$ , we find the node  $s_{\text{best}} = (q_{\text{best}}, x_{\text{best}}) \in S_{\text{near}}$  with the least cost-to-come value that satisfies the transition relation  $q_{\text{new}} = \delta(q_{\text{best}}, L(x_{\text{best}}))$ , and then add an edge  $\varepsilon = \langle s_{\text{best}}, s_{\text{new}} \rangle$  to the tree  $\mathcal{T}$ . The *cost-to-come values*—the cumulative distance from the tree root to the states—are allocated to the newly added nodes  $S_{\text{added}}$  in the tree. The remaining nodes  $S_{\text{new}} \setminus S_{\text{added}}$  are included in graph  $\mathcal{G}$ , and they can be added to the tree in the future by rewiring. The rewiring process is similar to that in [20] except for the hybrid state space and constraints on transition relations. Note that not only tree nodes  $\mathcal{V}_{\mathcal{T}}$  but also isolated vertices  $\mathcal{V}_{\text{iso}}$  can be rewired. The planner checks if the newly sampled node  $s_{\text{new}} = (q_{\text{new}}, x_{\text{new}})$  can lead to the acceptance one hop away: If  $\delta(q_{\text{new}}, L(x_{\text{new}})) \in \mathcal{F}$ , then a new solution is identified immediately.

##### C. Real-time solution improvement

Since the workspace may be very large, the initial feasible plan can be sub-optimal and has a large accumulative cost value. We develop a solution update scheme that continuously improves the solution and steers the robot following the least-cost path toward the temporal goal.

Every time an accepting path  $\tau$  is found, the planner sets  $\tau$  as the current solution if no solutions have been found so far. Otherwise, the planner inspects the solutions obtained before and only updates those solutions that are similar to the new solution while costing more. The intuition of the procedure above is that different from traditional path planning whose goal is a point in cartesian space, a temporal logic goal can be achieved as long as the corresponding run is accepting. An example is “go to site A or site B,” where distinct paths driving the robot to either site A or site B can both fulfill the given task. To exclude duplicated solutions, the planner compares two solutions on their *event-based traces* as well as path topology. We use the concept of homotopy class [24] to decide whether paths are similar in topology, i.e.,

whether they are with the same start and end positions and can be deformed into each other without crossing obstacles or labeled regions. The set of paths that satisfies the temporal goal and are distinct from each other are stored in a solution library. The best solution is elected from all path candidates with the least cost.

#### D. Robot state update

In existing works on anytime path planning [18], [25], the planning only focuses on the pending portion of the current plan while the nodes traveled are not useful anymore. However, considering the potential incidents in the real world, the robot may have to turn around and travel through the visited nodes in order to satisfy temporal logic tasks. During the execution of the plan, we keep the traveled edges alive by moving the root along with the robot motion. Instead of setting the next un-visited node as root like in [18], [20], we transform the classical tree into a dual-root tree during the robot movement. Fig. 2 uses the task “go to site A or site B” as an illustration. Denoting the executed edge connecting the last root  $R_{old}$  and the current motion target  $s_{tgtOld}$  as  $\varepsilon_{old} = \langle R_{old}, s_{tgtOld} \rangle$ , we reverse this edge once the robot arrives at  $s_{tgtOld}$ . Then  $s_{tgtOld}$  is set as the new root of the tree  $R_{new}$ . We further denote the current edge to execute as  $\varepsilon_{new} = \langle R_{new}, s_{tgtNew} \rangle$ , the subtree stemming from  $s_{tgtNew}$  as  $\mathcal{T}_{future}$ , and the part  $\mathcal{T} \setminus \mathcal{T}_{future} \setminus \varepsilon_{new}$  as  $\mathcal{T}_{past}$ .  $\mathcal{T}_{past}$  will not be visited as the execution progresses on the current solution unless the replanning is triggered. A rewire process starts from  $R_{new}$  at every timestamp to minimize the cost-to-come values of each node in  $\mathcal{T}_{past}$ . Additionally,  $s_{tgtNew}$  is treated as an auxiliary root  $R_{aux}$ , and the cost-to-come value of every node in  $\mathcal{T}_{future}$  is calculated by retracing to  $R_{aux}$ . This is an essential step to prevent nodes in  $\mathcal{T}_{future}$  be rewired to  $\mathcal{T}_{past}$  during execution.




Fig. 2: Illustration of a dual-root tree during execution. The motion tree  $\mathcal{T}$  is composed by  $\mathcal{T}_{past} \cup \varepsilon_{new} \cup \mathcal{T}_{future}$ . Both  $R_{new}$  and  $R_{aux}$  have zero cost-to-come values.

One can notice that when the edge  $\varepsilon_{old}$  is reverted, the transition relation  $R_{old}|_Q = \delta(R_{new}|_Q, R_{new}|_X)$  might not hold. Hence Algorithm 1 is applied on  $\mathcal{T}_{past}$  in a depth-first-search manner. At the beginning of the propagation,  $R_{new}$  is the parent node, and  $R_{old}$  is the child node. In Line 3, the function  $\text{CheckATran}(\cdot)$  checks if the transition relation from  $q_m$  to  $q_n$  is valid given the label of the parent node. If the transition is broken, routines described in Line 5-10 will be applied. The algorithm searches the set of the successor states of  $q_m$  and finds the state  $q_k$ , which can be

evolved from  $q_m$  by taking the word  $L(x_i)$ . The automata state of the corresponding node is therefore updated. This process ensures that all nodes in  $\mathcal{T}_{past}$  are reachable in future replanning. Recall the sampling process in Section IV-B,  $s_c$  records a node set  $S_c^{rcd}$  with all possible combinations of  $x_c$  and  $q \in Q \setminus \mathcal{B}$ . If there is no node in  $S_c^{rcd}$  with the same hybrid state as  $s_c$  (Line 7), the algorithm adds a copy of  $s_c$  to the isolated vertices set  $\mathcal{V}_{iso}$  (Line 8) before modifying  $s_c$ .

---

#### Algorithm 1: PropagateState

---

**Input:** Parent node  $s_p$ , child node  $s_c$ , DFA  $\mathcal{A}_\varphi$ , isolated vertices set  $\mathcal{V}_{iso}$   
**Output:** Successfulness of the propagation

```

1  $s_p = (q_m, x_i)$ 
2  $s_c = (q_n, x_j)$ 
3 if  $\text{CheckATran}(q_m, L(x_i), q_n, \mathcal{A}_\varphi)$  then
4   return  $\top$ 
5 for  $q_k \in \{q_v \mid \exists \lambda \in 2^{AP}, q_v = \delta(q_m, \lambda)\}$  do
6   if  $\text{CheckATran}(q_m, L(x_i), q_k, \mathcal{A}_\varphi)$  then
7     if  $|\text{SameStateNodes}(s_c)| = 0$  then
8        $\text{AddNodeCopy}(s_c, \mathcal{V}_{iso})$ 
9      $s_c \leftarrow (q_k, x_j)$ 
10    return  $\top$ 
11 return  $\perp$ 

```

---

However, after line 9 in Algorithm 1 modifies the automata state of  $s_c$ , duplication may occur between  $s_c$  and the node with the same hybrid states in  $S_c^{rcd}$ . Similar problems occur after the propagation in dealing with propositional updates, as to be mentioned in Section IV-E. It will incur problems in extending and rewiring nodes as several nodes with identical hybrid states and costs exist simultaneously. Besides, duplicated nodes make it difficult to maintain tree sparsity. To prevent the tree from becoming entangled and excessively dense, we require an algorithm to visit all the duplicated nodes and only keep the one that has the least cost-to-come value.

We propose Algorithm 2 to deduplicate the nodes in linear time. When handling robot state updates, the algorithm is recursively applied on the  $\mathcal{T}_{past}$  in a preorder fashion and eliminates one duplicated node at a time. In Line 1, denoting the visiting node as  $s_m$ , all the nodes with the same hybrid state are identified as  $S_{dup}$  given the node set  $S_m^{rcd} = \{s \in \mathcal{G} \mid s|_X = s_m|_X\}$ . In Line 2-Line 4, the duplicated and isolated states  $\mathcal{V}_{isoDup}$  are detected and be removed from the isolated vertices set  $\mathcal{V}_{iso}$  as well as duplicated states set  $S_{dup}$  permanently. In Line 7, one element  $s_d$  in the set  $S_{dup}$  is selected, while the rest of the nodes are to be processed later. If the pair of states  $s_d, s_m$  is directly connected, i.e., are in parent-child relationship, the child node is marked as  $s_{worse}$ . Otherwise, the algorithm evaluates the cost of the two nodes  $s_m, s_d$  and marks the node that costs more as  $s_{worse}$ . After that, in Line 16-Line 20, the inward and outward edges of  $s_{worse}$  is re-connected to the better node  $s_{better}$ , and  $s_{worse}$  node is safely deleted. The offspring of the  $s_{better}$  is visited recursively by Algorithm 2 until every node in the subtree is marked as visited.

As a consequence of Algorithm 1 and Algorithm 2 to-

---

**Algorithm 2:** MergeNode

---

**Input:** Node  $s_m$ , tree vertices set  $\mathcal{V}_{\mathcal{T}}$ , tree edge set  $\mathcal{E}_{\mathcal{T}}$ , isolated vertices set  $\mathcal{V}_{\text{iso}}$   
**Output:** Successfulness of the merging

```
1  $S_{\text{dup}} \leftarrow \text{SameStateNodes}(s_m)$ 
2  $\mathcal{V}_{\text{isoDup}} \leftarrow S_{\text{dup}} \cap \mathcal{V}_{\text{iso}}$ 
3  $\mathcal{V}_{\text{iso}} \leftarrow \mathcal{V}_{\text{iso}} \setminus \{\mathcal{V}_{\text{isoDup}}\}$ 
4  $S_{\text{dup}} \leftarrow S_{\text{dup}} \setminus \{\mathcal{V}_{\text{isoDup}}\}$ 
5 if  $|S_{\text{dup}}| = 0$  then
6   return  $\top$ 
7  $s_d \leftarrow S_{\text{dup}}.\text{pop}()$ 
8 if  $s_d = \text{Parent}(s_m) \vee \text{Cost}(s_d) \leq \text{Cost}(s_m)$  then
9    $s_{\text{better}} \leftarrow s_d$ 
10   $s_{\text{worse}} \leftarrow s_m$ 
11 else
12    $s_{\text{better}} \leftarrow s_m$ 
13    $s_{\text{worse}} \leftarrow s_d$ 
14  $s_{\text{better}}.\text{visited} \leftarrow \top$ 
15  $s_{\text{worse}}.\text{visited} \leftarrow \top$ 
16 for  $s_{\text{child}} \in \text{Children}(s_{\text{worse}})$  do
17    $\mathcal{E}_{\mathcal{T}} \leftarrow \mathcal{E}_{\mathcal{T}} \setminus \{ \langle s_{\text{worse}}, s_{\text{child}} \rangle \}$ 
18    $\mathcal{E}_{\mathcal{T}} \leftarrow \mathcal{E}_{\mathcal{T}} \cup \{ \langle s_{\text{better}}, s_{\text{child}} \rangle \}$ 
19  $\mathcal{E}_{\mathcal{T}} \leftarrow \mathcal{E}_{\mathcal{T}} \setminus \{ \langle \text{Parent}(s_{\text{worse}}), s_{\text{worse}} \rangle \}$ 
20  $\mathcal{V}_{\mathcal{T}} \leftarrow \mathcal{V}_{\mathcal{T}} \setminus \{s_{\text{worse}}\}$ 
21 return  $\top$ 
```

---

gether, all nodes in  $\mathcal{T}$  are updated with future-reachable automata states, and the cost-to-come value of each node is less than or equal to the original value.

### E. Environment state update

In real-world applications, prior knowledge may not be precise, making the acceptance condition  $\tau \models \varphi$  invalid. To tackle this problem, we enable the robot to sense the environment and receive observations from other agents. We consider a mobile robot moving in a scenario where the properties of regions can change, and unexpected obstacles may block the path. The robot continuously analyzes sensor data and extracts *knowledge* in two forms:  $\kappa_{\text{reg}} = (\pi, L(\pi))$  or  $\kappa_{\text{obs}} = (v, x_v)$ , where  $\pi \in \Pi$  stands for the region covered by the sensor,  $L(\pi)$  is the set of propositions satisfied in  $\forall x \in \pi$ ,  $v$  is an unexpected obstacle instance, and  $x_v$  is the position of  $v$ . If knowledge stored in memory is found inconsistent with the newly obtained knowledge, the robot will update the knowledge base and perform plan revision.

Every time  $\kappa_{\text{reg}}$  is obtained, labels of the region  $\pi$  are updated. The planner first calls Algorithm 1 to correct the states of nodes in the whole tree  $\mathcal{T}$  given the updated workspace, and then uses Algorithm 2 to resolve state duplication. Afterward, each possible solution is examined to see if its trace satisfies  $\varphi$ . If there still exist feasible solutions, the solution with the least accumulative cost is set as the current plan. Otherwise, if the knowledge change is critical and all previous solutions are infeasible, the robot will come to a full stop at  $x_{\text{stop}}$ . The new root is set as  $s_{\text{stop}} = (q_{\text{stop}}, x_{\text{stop}})$ , where  $q_{\text{stop}} \in Q$  is calculated by the previous root. The robot waits at  $x_{\text{stop}}$  until the planner obtains a new solution as described in Section IV-B.

Given knowledge  $\kappa_{\text{obs}}$ , if the obstacle  $v$  is not observed before, or the traveled distance  $\Delta x_v$  exceeds a threshold, the

planner updates  $X_{\text{free}}$  and triggers the replanning procedure. The planner first blocks the edges that cross  $v$  by setting their weights of them as positive infinity. Afterward, like the situation in dealing with  $\kappa_{\text{reg}}$ , the robot will carry out the best executable solution or await a feasible solution generated by the sampling-based search.

After coping with environmental state change, the search tree is corrected so that every edge in  $\mathcal{T}$  satisfies its transition relation in  $\mathcal{A}_{\varphi}$ . If no solutions have been found so far, the planning is converted to an LTL-constrained sampling-based tree search described in the work [15], which has been proved to be probabilistically complete. Hence by continually growing  $\mathcal{T}$  during replanning, a solution that satisfies  $\varphi$  is bound to be found if one exists in the domain  $X$ . Therefore, we can conclude that Problem 3.1 is solved.

## V. IMPLEMENTATION AND EXPERIMENTAL RESULTS

For comparison and illustration purposes, we present case studies on fire extinguishing tasks. The target scenario contains three regions of interest, three known obstacles, and one unknown obstacle, as can be seen in Fig. 4(a) and Fig. 4(c). Specifically,  $l_1$  and  $l_3$  are labeled as grasslands and  $l_2$  is a pond in the robot's prior knowledge, and the robot can validate a label or detect an obstacle within its sensing area ( $3\text{m} \times 3\text{m}$ ). To ensure that the sensing distance is enough for braking, we set the maximum velocity of the robot and the obstacle as 0.5 m/s and 0.2 m/s, respectively.

### A. Comparative results in simulation

We conduct simulations in three workspaces with different types of uncertainties.  $X_A$  is a workspace where the initial knowledge of either one of *grassland* regions is inaccurate.  $X_B$  is a workspace where one obstacle moves from one side of the workspace to the other side.  $X_C$  combines the property uncertainty and occupancy uncertainty of  $X_A$  and  $X_B$ . A vehicle is given a task formula  $\varphi_1 = \diamond(\text{pond} \wedge \diamond \text{grassland})$ , which instructs the vehicle to fetch water from the pond and then put out the fire in the grasslands. We select methods from Guo *et al.* [12] and Luo *et al.* [15] to compare the replanning speed and solution quality, and all simulations are conducted in Python with an Intel i7 CPU. The abstraction-dependent approach [12] is used to compare the performance only in  $X_A$  since the planner cannot deal with dynamic behaviors in  $X_B$  and  $X_C$ . The abstraction-free approach [15] is implemented by us in a reactive way that, once the current path is found invalid, the vehicle saves the execution progress and builds the search tree again. Results can be viewed in Table I, where each entry is obtained by averaging 30 simulation rounds with randomized initial robot configurations. Our planner outperformed Luo *et al.* [15] by adapting to uncertainties faster and planning shorter paths to satisfy temporal goals. Furthermore, our approach has broader applicability than Guo *et al.* [12].

### B. System integration and real-world experiment

The proposed planner produces a path  $\tau = s_0 s_1 \dots s_n$  satisfying the task formula. To drive the robot ahead following the piece-wise linear path  $\tau|_X$ , a trajectory generator is

TABLE I: Performance of different planners

Metrics	Methods	$X_A$	$X_B$	$X_C$
Total Completion Time (s)	Guo <i>et al.</i>	23.15	—	—
	Luo <i>et al.</i>	30.19	34.78	46.94
	Ours	<b>22.90</b>	<b>30.14</b>	<b>29.56</b>
Average Replanning Time (s)	Guo <i>et al.</i>	0.44	—	—
	Luo <i>et al.</i>	5.87	6.82	6.92
	Ours	<b>0.39</b>	<b>1.63</b>	<b>1.24</b>
Total Travel Distance (m)	Guo <i>et al.</i>	10.0	—	—
	Luo <i>et al.</i>	8.49	8.54	9.83
	Ours	<b>8.01</b>	<b>7.03</b>	<b>9.26</b>

required to track the path without deviating from the task formula. Note the trajectory generator is system-dependent due to different robot dynamics. Following the assumption of robot dynamics in Section III, we utilize a motion-primitive-based local motion planner [26], which can be deployed on both multicopter and differential-drive vehicles. In detail, given start configuration  $x_i$  and end configuration  $x_{i+1}$ , the motion planner solves a boundary value problem from  $x_i$  to  $x_{i+1}$  and generates a motion primitive  $\zeta(x_i, u_i, t_i)$  in a model predictive control manner. In the meantime, a volumetric mapper [27] builds a Euclidean distance transform (EDT) map for the motion planner by combining data from prior knowledge and online observation.

In the real-world experiment, we use a multicopter to follow the nominal path in real time. A motion capture system provides position estimation for the multicopter as well as the observation of the obstacle if within the sensing range. The multicopter is asked to accomplish the task: “Don’t fly over the grasslands until you’ve got water from the pond, and eventually put out the fire in a grassland without colliding with obstacles.” The formula is written as  $\varphi_2 = (\neg \text{grassland} \cup \text{pond}) \wedge \diamond \text{grassland}$ , and is translated into a DFA, as in Fig. 3.




Fig. 3: The DFA of task formula  $\varphi_2$

In DFA preprocess (Section IV-A), the algorithm identifies that the automata state **3** belongs to  $\mathcal{B}$ . The resulting temporally forbidden zones  $\Pi_{\text{fbd}}$  together with unexpected obstacles are encapsulated into pointcloud and broadcast to ROS. The volumetric mapper clusters the pointcloud and yields an EDT map pushing the robot away from  $\Pi_{\text{fbd}}$  and obstacles. With the EDT map updated in every planning iteration, the motion planner generates a series of safety-guaranteed motion primitives. All computations are running at 10 Hz on an onboard computer, Nvidia Xavier NX.

One typical execution can be viewed in Fig. 4. The multicopter takes off at the center of the map and carries out the

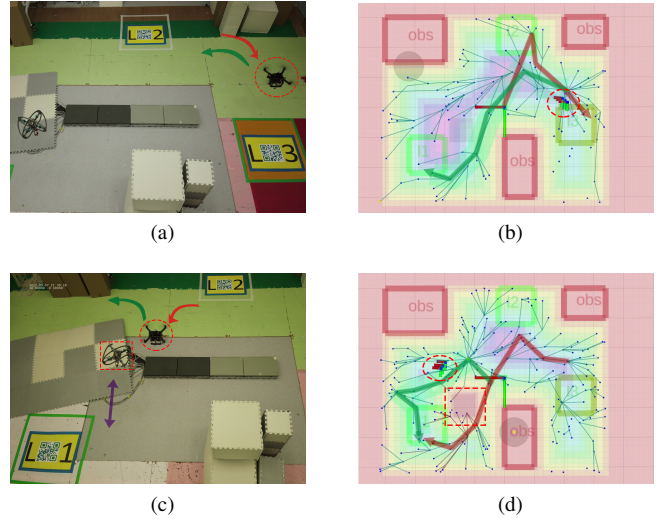


Fig. 4: An indoor experiment of the task  $\varphi_2$ . (a) The multicopter (in the red dashed circle) identifies the labeling error in  $l_3$  (highlighted in yellow). (b) The old path (dark red) is no longer valid, and a new path (green) takes its place almost instantly. The background is colored according to the distance value in EDT. (c) The multicopter detects an unexpected obstacle (in the red dashed rectangle) which moves back and forth. (d) A new accepting path is calculated without intersecting the moving obstacle.

initial plan  $l_2 \rightarrow l_3 \rightarrow l_2$ . In Fig. 4(a), after getting water, the multicopter flies towards  $l_3$  and receives knowledge  $\kappa_{\text{reg}} = (l_3, \emptyset)$ , indicating that the property of region  $l_3$  is incorrect in  $X_{\text{init}}$ . Benefiting from the solution library maintained, the multicopter rapidly adopts an alternative solution, turning around and heading to  $l_1$  (Fig. 4(b)). In Fig. 4(c), the multicopter gets knowledge  $\kappa_{\text{obs}} = (v_1, (1.36, 0.98, 1.02))$  during flight, and potential collision is detected. Thanks to the sample reusing strategy, the multicopter quickly alters its course and flies around the obstacle to  $l_3$  (Fig. 4(d)). The multicopter eventually completes the task by reaching  $l_3$  and putting out the fire. Our planner is shown to successfully accommodate property and occupancy uncertainties through online plan revision and refinement mechanisms. The full video can be found at <https://youtu.be/UH-4KcCUixw>.

## VI. CONCLUSION

In this paper, we have proposed a high-level path planner for mobile robots to complete a temporal logic task regardless of a limited understanding of the workspace. The planner quickly obtains a feasible plan and improves it over time during execution. We develop two algorithms for node reusing and deduplication to overcome issues that may impair the correctness and effectiveness of our plan-revising mechanism when responding to incidents. Results in simulation and real-world missions show that the proposed planner can facilitate online robot tasking in partially-unknown environments. Future work includes developing a user-friendly interface and dedicated motion planning algorithms.

## REFERENCES

- [1] H. Liu, J. Chen, J. Feng, and H. Zhao, "An improved RRT\* UAV formation path planning algorithm based on goal bias and node rejection strategy," *Unmanned Systems*, 2022.
- [2] Y. Dong, Y. Zhang, and J. Ai, "Experimental test of unmanned ground vehicle delivering goods using rrt path planning algorithm," *Unmanned Systems*, vol. 5, no. 01, pp. 45–57, 2017.
- [3] J. A. Douthwaite, S. Zhao, and L. S. Mihaylova, "Velocity obstacle approaches for multi-agent collision avoidance," *Unmanned Systems*, vol. 7, no. 01, pp. 55–64, 2019.
- [4] T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Synthesis of control protocols for autonomous systems," *Unmanned Systems*, vol. 1, no. 01, pp. 21–39, 2013.
- [5] M. Lan, S. Lai, T. H. Lee, and B. M. Chen, "A survey of motion and task planning techniques for unmanned multicopter systems," *Unmanned Systems*, vol. 9, no. 02, pp. 165–198, 2021.
- [6] Y. Kantaros and M. M. Zavlanos, "Stylus\*: A temporal logic optimal control synthesis algorithm for large-scale multi-robot systems," *The International Journal of Robotics Research*, vol. 39, no. 7, pp. 812–836, 2020.
- [7] E. M. Clarke, "Model checking," in *International Conference on Foundations of Software Technology and Theoretical Computer Science*. Springer, 1997, pp. 54–56.
- [8] H. Lin and P. J. Antsaklis, *Hybrid Dynamical Systems: Fundamentals and Methods*. Springer, 2021.
- [9] G. E. Fainekos, A. Girard, H. Kress-Gazit, and G. J. Pappas, "Temporal logic motion planning for dynamic robots," *Automatica*, vol. 45, no. 2, pp. 343–352, 2009.
- [10] D. Tian, H. Fang, Q. Yang, and Y. Wei, "Decentralized motion planning for multiagent collaboration under coupled LTL task specifications," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 2021.
- [11] P. Schillinger, M. Bürger, and D. V. Dimarogonas, "Simultaneous task allocation and planning for temporal logic goals in heterogeneous multi-robot systems," *The international journal of robotics research*, vol. 37, no. 7, pp. 818–838, 2018.
- [12] M. Guo, K. H. Johansson, and D. V. Dimarogonas, "Revising motion planning under linear temporal logic specifications in partially known workspaces," in *2013 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2013, pp. 5025–5032.
- [13] S. C. Livingston, R. M. Murray, and J. W. Burdick, "Backtracking temporal logic synthesis for uncertain environments," in *2012 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2012, pp. 5163–5170.
- [14] M. Lahijanian, M. R. Maly, D. Fried, L. E. Kavraki, H. Kress-Gazit, and M. Y. Vardi, "Iterative temporal planning in uncertain environments with partial satisfaction guarantees," *IEEE Transactions on Robotics*, vol. 32, no. 3, pp. 583–599, 2016.
- [15] X. Luo, Y. Kantaros, and M. M. Zavlanos, "An abstraction-free method for multirobot temporal logic optimal control synthesis," *IEEE Transactions on Robotics*, vol. 37, no. 5, pp. 1487–1507, 2021.
- [16] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Informed RRT\*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2014, pp. 2997–3004.
- [17] M. P. Strub and J. D. Gammell, "Adaptively informed trees (AIT\*) and effort informed trees (EIT\*): Asymmetric bidirectional sampling-based path planning," *The International Journal of Robotics Research*, vol. 41, no. 4, pp. 390–417, 2022.
- [18] S. Karaman, M. R. Walter, A. Perez, E. Frazzoli, and S. Teller, "Anytime motion planning using the RRT," in *2011 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2011, pp. 1478–1483.
- [19] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon" next-best-view" planner for 3d exploration," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 1462–1468.
- [20] K. Naderi, J. Rajamäki, and P. Hämäläinen, "RT-RRT\* a real-time path planning algorithm based on RRT," in *Proceedings of the 8th ACM SIGGRAPH Conference on Motion in Games*, 2015, pp. 113–118.
- [21] G. De Giacomo and M. Y. Vardi, "Linear temporal logic and linear dynamic logic on finite traces," in *IJCAI'13 Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. Association for Computing Machinery, 2013, pp. 854–860.
- [22] S. Zhu, L. M. Tabajara, J. Li, G. Pu, and M. Y. Vardi, "Symbolic LTLf synthesis," in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 2017, pp. 1362–1369. [Online]. Available: <https://doi.org/10.24963/ijcai.2017/189>
- [23] G. De Giacomo and S. Rubin, "Automata-theoretic foundations of FOND planning for LTLf and LDLf goals," in *IJCAI*, 2018, pp. 4729–4735.
- [24] S. Bhattacharya, M. Likhachev, and V. Kumar, "Topological constraints in search-based robot path planning," *Autonomous Robots*, vol. 33, no. 3, pp. 273–290, 2012.
- [25] M. G. Westbrook and W. Ruml, "Anytime kinodynamic motion planning using region-guided search," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 6789–6796.
- [26] S. Lai, M. Lan, and B. M. Chen, "Model predictive local motion planning with boundary state constrained primitives," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3577–3584, 2019.
- [27] Y. Chen, S. Lai, J. Cui, B. Wang, and B. M. Chen, "GPU-accelerated incremental Euclidean distance transform for online motion planning of mobile robots," *IEEE Robotics and Automation Letters*, 2022.