

An Abstraction-Free Method for Multi-Robot Temporal Logic Optimal Control Synthesis

Xusheng Luo, Yiannis Kantaros, *Member, IEEE*, and
Michael M. Zavlanos, *Senior Member, IEEE*

Abstract—The majority of existing Linear Temporal Logic (LTL) planning methods rely on the construction of a discrete product automaton, that combines a discrete abstraction of robot mobility and a Büchi automaton that captures the LTL specification. Representing this product automaton as a graph and using graph search techniques, optimal plans that satisfy the LTL task can be synthesized. However, constructing expressive discrete abstractions makes the synthesis problem computationally intractable. In this paper, we propose a new sampling-based LTL planning algorithm that does not require any discrete abstraction of robot mobility. Instead, it incrementally builds trees that explore the product state-space, until a maximum number of iterations is reached or a feasible plan is found. The use of trees makes data storage and graph search tractable, which significantly increases the scalability of our algorithm. To accelerate the construction of feasible plans, we introduce bias in the sampling process which is guided by transitions in the Büchi automaton that belong to the shortest path to the accepting states. We show that our planning algorithm, with and without bias, is probabilistically complete and asymptotically optimal. Finally, we present numerical experiments showing that our method outperforms relevant temporal logic planning methods.

Index Terms—Path planning for multiple mobile robots or agents, formal methods in robotics and automation, motion and path planning, optimization and optimal control.

I. INTRODUCTION

MOTION planning traditionally consists of generating robot trajectories that reach a goal region from a starting point while avoiding obstacles [1]. Methods for point-to-point navigation range from using potential fields and navigation functions [2], [3] to sampling-based algorithms [4]–[6]. More recently, new planning approaches have been proposed that can handle a richer class of tasks, than the classical point-to-point navigation, and can capture temporal goals. Such tasks can be, e.g., sequencing or coverage [7], data gathering [8], intermittent communication [9], or persistent surveillance [10], and can be captured using formal languages, such as Linear Temporal Logic (LTL) [11], developed in concurrency theory.

Control synthesis for robots under complex tasks, captured by Linear Temporal Logic (LTL) formulas, builds upon either bottom-up approaches when independent LTL expressions are assigned to robots [12], [13] or top-down approaches when a global LTL formula describing a collaborative task is assigned to a team of robots [14], [15], as in this work. Common in the above works is that they rely on model checking theory [11],

[16], to find plans that satisfy LTL-specified tasks, without optimizing task performance. Optimal control synthesis under local and global LTL specifications has been addressed in [17], [18] and [19]–[21], respectively. In the case of global LTL tasks [19]–[21], optimal discrete plans are derived for every robot using the individual transition systems that satisfy a bisimulation property [7] and are obtained through an abstraction process [22]–[26], and a Non-deterministic Büchi Automaton (NBA) that represents the global LTL specification. Specifically, by taking the synchronous product among the transition systems and the NBA, a synchronous product automaton can be constructed. Then, using graph-search techniques, optimal motion plans can be derived. A major limitation is the high computational cost of constructing expressive discrete abstractions that result in very large transition systems. Combining these large transition systems with many robots and complex tasks increases dramatically the size of the product automaton so that graph-search techniques become intractable. An additional limitation is that the resulting discrete plans are only optimal given the abstraction that was used to compute them. For a different abstraction, a different optimal plan will be computed. Therefore, global optimality can not be ensured.

To improve on the scalability of optimal control synthesis methods such as those discussed above, we recently proposed new more efficient sampling-based algorithms for discrete transition systems and global temporal logic tasks that avoid the construction of the product automaton altogether [27]–[33]. Specifically, in [27], we transformed given transition systems into trace-included transition systems, meaning that all the behaviors they generate can also be generated by the original transition systems, that have smaller state spaces but are still expressive enough to construct feasible motion plans. We experimentally validated the approach in [27] in [34]. In [28], [29] we proposed a more tractable sampling-based method that builds trees incrementally, similar to the approach proposed here, that approximate the product automaton until a motion plan is constructed, while in [30], [31] we extended the method in [28], [29] by introducing bias in the sampling process in directions that make the most progress towards satisfying the assigned task. As a result, the method in [30], [31] can solve planning problems that are orders of magnitude larger than the planning problems that can be solved using [28], [29]. In [32], we improved scalability of the biased sampling-based algorithm in [30], [31] by using experience from solving previous LTL planning problems. Finally, in [33], we provided a distributed implementation of [29] whereby robots collaborate to build subtrees so that the computational

Xusheng Luo, Yiannis Kantaros and Michael M. Zavlanos are with the Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708, USA. {xusheng.luo, yiannis.kantaros, michael.zavlanos}@duke.edu. This work is supported in part by the ONR under agreement #N00014-18-1-2374 and the AFOSR under award #FA9550-19-1-0169.

time is decreased significantly. However, [27]–[33] require a discrete abstraction of the robot dynamics and the environment, which is computationally expensive to construct [24].

Motivated by existing sampling-based algorithms for point-to-point navigation in continuous state spaces [6], we propose a sampling-based temporal logic planning method that, unlike the methods in [27]–[33], does not require any discrete abstraction of robot mobility. Our algorithm builds incrementally trees that explore both the continuous state space of robot positions and the discrete state space of the NBA, simultaneously. Specifically, we first build a tree that connects an initial state to an *accepting* state of the product automaton. Tracing the nodes from the accepting state back to the root returns a plan that corresponds to the *prefix* part of the plan and is executed once. Then, we construct a new tree rooted at an accepting state in a similar way and propose a cycle-detection method to discover a loop around that root, which corresponds to the *suffix* part of the plan and is executed indefinitely. By construction, the continuous execution of the generated plans satisfies the LTL-based task. Moreover, we show that our algorithm is probabilistically complete and asymptotically optimal. Inspired by [30], [31], to accelerate the construction of a feasible solution, we also propose a biased sampling method guided by transitions in the NBA that trace shortest paths to the accepting states. The simulations show that our biased sampling method can significantly accelerate the detection of feasible but also low-cost plans. Our algorithm can be viewed as an extension of RRT* [6] for LTL-based tasks. Nevertheless, completeness and optimality of our algorithm can not be trivially inherited from RRT* that is designed exclusively for continuous state spaces, or from [27]–[33] that rely on discrete abstractions of robot mobility.

Related are also the methods in [35]–[38]. Particularly, [35] formulates the planning problem as an Integer-Linear-Program (ILP) using a discrete abstraction of the environment. To avoid the construction of discrete abstractions and the corresponding product automata, the authors of [36], [37] encode the LTL specifications as mixed-integer linear programming (MILP) over the continuous system variables, and employ off-the-shelf optimization solvers to obtain the optimal solution. However, [36] only considers a subclass of co-safe LTL formulas that require all robots to return to their original places and stay there forever, which can be satisfied by a finite-length robot trajectory, whereas here we consider arbitrary LTL formulas. Moreover, the methods in [36], [37] require the user-specified length of the prefix-suffix plan that satisfies the LTL formula. On the other hand, the algorithm proposed in [38] relies on the satisfiability modulo convex (SMC) approach [39]. By formulating the LTL planning problem as a feasibility problem over Boolean and convex constraints, and leveraging state-of-the-art SAT solvers and convex optimization solvers, this method scales better than the MILP-based method [36], [37] by relying on a coarse abstraction of the workspace. However, since [38] formulates the problem as a feasibility problem, this method does not have any optimality guarantees.

To the best of our knowledge, the most relevant abstraction-free LTL planning methods are presented in [40]–[44]. Specifically, in [40], [41], sampling-based algorithms employ an

RRG-like algorithm to build incrementally a Kripke structure until it is expressive enough to generate a plan that satisfies a task expressed in deterministic μ -calculus. However, building arbitrary structures to represent transition systems compromises scalability since, as the number of samples increases, so does the density of the constructed graph which requires more resources to save and search for optimal plans using graph search methods. Motivated by this limitation, [42] proposes a similar RRG algorithm, but constructs incrementally sparse graphs representing transition systems that are then used to construct a product automaton. Then correct-by-construction discrete plans are synthesized by applying graph search methods on the product automaton. However, similar to [40], [41], as the number of samples increases, the sparsity of the constructed graph is lost. Therefore, the methods in [40]–[42] can only be used for simple planning problems. To the contrary, our sampling-based method builds trees, instead of graphs of arbitrary structure, so that optimal plans can be easily constructed by tracing the sequence of parent nodes starting from desired accepting states. Combined with the proposed biased sampling method, this allows us to handle larger problems compared to, e.g., [42]. Our approach is also closely related to [43], [44], where high-level states selected from a discrete product automaton state space are used to guide the low-level sampling-based planner in the continuous space. However, the biased sampling method in [43], [44] relies on a discrete abstraction of robot mobility, whereas the biased sampling method proposed here is abstraction-free and guided by the NBA, as in our previous work [30], [31] for discrete state spaces. Compared to [43], [44], we show that our method is asymptotically optimal and more scalable.

The main contribution of this paper can be summarized as follows. First, we propose a new abstraction-free temporal logic planning algorithm for multi-robot systems that is highly scalable. Scalability is due to the use of trees to represent the transition system and the proposed bias in the sampling process. Second, we show that the proposed algorithm, with or without bias, is probabilistically complete and asymptotically optimal. Finally, we provide extensive simulation studies that show that the proposed method outperforms relevant state-of-the-art algorithms [38], [42], [43].

The rest of the paper is organized as follows. In Section II and III we present preliminaries and the problem formulation, respectively. In Section IV we describe the sampling-based planning algorithm. We introduce the biased sampling method in Section V. Furthermore, we examine their correctness and optimality in Section VI. Simulation results and conclusive remarks are presented in Sections VII and VIII, respectively.

II. PRELIMINARIES

In this section, we formally describe Linear Temporal Logic (LTL) by presenting its syntax and semantics. Also, we briefly review preliminaries of automata-based LTL model checking.

Linear temporal logic [11] is a type of formal logic whose basic ingredients are a set of atomic propositions \mathcal{AP} , the boolean operators, conjunction \wedge and negation \neg , and temporal operators, next \bigcirc and until \mathcal{U} . LTL formulas over \mathcal{AP} abide by the grammar $\phi ::= \text{true} \mid \pi \mid \phi_1 \wedge \phi_2 \mid \neg \phi \mid \bigcirc \phi \mid \phi_1 \mathcal{U} \phi_2$.

For brevity, we abstain from deriving other Boolean and temporal operators, e.g., *disjunction* \vee , *implication* \Rightarrow , *always* \square , *eventually* \diamond , which can be found in [11].

An infinite word σ over the alphabet $2^{\mathcal{A}^{\mathcal{P}}}$ is defined as an infinite sequence $\sigma = \pi_0\pi_1\dots \in (2^{\mathcal{A}^{\mathcal{P}}})^\omega$, where ω denotes an infinite repetition and $\pi_k \in 2^{\mathcal{A}^{\mathcal{P}}}$, $\forall k \in \mathbb{N}$. The language $\text{Words}(\phi) = \{\sigma \mid \sigma \models \phi\}$ is defined as the set of words that satisfy the LTL formula ϕ , where $\models \subseteq (2^{\mathcal{A}^{\mathcal{P}}})^\omega \times \phi$ is the satisfaction relation. An LTL formula ϕ can be translated into a Nondeterministic Büchi Automaton defined as follows [45]:

Definition 2.1: (NBA) A *Nondeterministic Büchi Automaton* (NBA) B over $2^{\mathcal{A}^{\mathcal{P}}}$ is defined as a tuple $B = (\mathcal{Q}_B, \mathcal{Q}_B^0, \Sigma, \rightarrow_B, \mathcal{Q}_B^F)$, where \mathcal{Q}_B is the set of states, $\mathcal{Q}_B^0 \subseteq \mathcal{Q}_B$ is a set of initial states, $\Sigma = 2^{\mathcal{A}^{\mathcal{P}}}$ is an alphabet, $\rightarrow_B \subseteq \mathcal{Q}_B \times \Sigma \times \mathcal{Q}_B$ is the transition relation, and $\mathcal{Q}_B^F \subseteq \mathcal{Q}_B$ is a set of accepting/final states.

An *infinite run* ρ_B of B over an infinite word $\sigma = \pi_0\pi_1\pi_2\dots$, $\pi_k \in \Sigma = 2^{\mathcal{A}^{\mathcal{P}}}$, $\forall k \in \mathbb{N}$, is a sequence $\rho_B = q_B^0q_B^1q_B^2\dots$ such that $q_B^0 \in \mathcal{Q}_B^0$ and $(q_B^k, \pi_k, q_B^{k+1}) \in \rightarrow_B$, $\forall k \in \mathbb{N}$. An infinite run ρ_B is called *accepting* if $\text{Inf}(\rho_B) \cap \mathcal{Q}_B^F \neq \emptyset$, where $\text{Inf}(\rho_B)$ represents the set of states that appear in ρ_B infinitely often. The words σ that produce an accepting run of B constitute the accepted language of B , denoted by \mathcal{L}_B . Then [11] proves that the accepted language of B is equivalent to the words of ϕ , i.e., $\mathcal{L}_B = \text{Words}(\phi)$.

III. PROBLEM FORMULATION

Consider N robots residing in a workspace, represented by a compact subset $\mathcal{W} \subset \mathbb{R}^d$, $d \in \{2, 3\}$. Let \mathcal{O} , an open subset of \mathcal{W} , denote the set of obstacles and $\mathcal{W}_{\text{free}} = \mathcal{W} \setminus \mathcal{O}$ denote the free workspace. Let $\mathcal{R} = \{\ell_j\}_{j=1}^W$ be a set of $W \in \mathbb{N}^+$ closed labeled regions that can have any arbitrary shapes. Given a labeled region of interest ℓ_j , $j \in [W]$, where $[W]$ is the shorthand of $\{1, \dots, W\}$, let $\partial\ell_j$ denote its boundary, i.e., the closure of ℓ_j without its interior. The motion of robot i is captured by a Transition System (TS) defined as follows:

Definition 3.1: (TS) A *Transition System* for robot i , denoted by TS_i , is a tuple $\text{TS}_i = \{\mathcal{W}, \mathbf{x}_i^0, \rightarrow_i, \mathcal{A}\mathcal{P}_i, L_i\}$ where: (a) $\mathcal{W} = \{\mathbf{x}_i\}$ is the set of infinite position states of robot i ; (b) \mathbf{x}_i^0 is the initial position of robot i ; (c) $\rightarrow_i \subseteq \mathcal{W} \times \mathcal{W}$ is the transition relation. We consider robot dynamics subject to holonomic constraints, as in [6]. Let $\overline{\mathbf{x}_i\mathbf{x}'_i}$ denote the straight line connecting the points \mathbf{x}_i and \mathbf{x}'_i . A transition from \mathbf{x}_i to \mathbf{x}'_i , denoted by $(\mathbf{x}_i, \mathbf{x}'_i) \in \rightarrow_i$, exists if (i) $\overline{\mathbf{x}_i\mathbf{x}'_i}$ is inside $\mathcal{W}_{\text{free}}$ and (ii) it crosses any boundary $\partial\ell_j$, $j \in [W]$ at most once, as in [42];¹ (d) $\mathcal{A}\mathcal{P}_i = \bigcup_{j=1}^W \{\pi_i^{\ell_j}\} \cup \{\emptyset\}$ is the set of atomic propositions, where $\pi_i^{\ell_j}$ is true if robot i is inside region ℓ_j (including its boundary) and false otherwise, and \emptyset means robot i is not inside any labeled region; (e) $L_i : \mathcal{W} \rightarrow \mathcal{A}\mathcal{P}_i$ is the observation (labeling) function that returns an atomic proposition satisfied at a location \mathbf{x}_i , i.e., $L_i(\mathbf{x}_i) = \pi_i^{\ell_j}$ or \emptyset .

¹Suppose $\mathbf{x}_i \in \ell_m$, $\mathbf{x}'_i \in \ell_n$ and $\overline{\mathbf{x}_i\mathbf{x}'_i} \subset \mathcal{W}_{\text{free}}$ crosses region ℓ_k twice, further consider a task that requires to always avoid ℓ_k . If the transition depends only on condition (i) whether $\overline{\mathbf{x}_i\mathbf{x}'_i}$ is inside $\mathcal{W}_{\text{free}}$, then we have $(\mathbf{x}_i, \mathbf{x}'_i) \in \rightarrow_i$. However, the path corresponding to $\overline{\mathbf{x}_i\mathbf{x}'_i}$ crosses region ℓ_k thus violating the task. The transition condition (ii) excludes this case where the discrete plan satisfies the task but the underlying continuous path does not, because it passes through an undesired region between two endpoints.

We emphasize that a TS is defined over the continuous state space. Given a TS, we define the Product Transition System (PTS), which captures all possible combinations of robot behaviors captured by their respective TS_i , as follows:

Definition 3.2: (PTS) Given N transition systems $\text{TS}_i = \{\mathcal{W}, \mathbf{x}_i^0, \rightarrow_i, \mathcal{A}\mathcal{P}_i, L_i\}$, the product transition system $\text{PTS} = \text{TS}_1 \times \text{TS}_2 \times \dots \times \text{TS}_N$ is a tuple $\text{PTS} = (\mathcal{W}^N, \mathbf{x}_0, \rightarrow_{\text{PTS}}, C, \mathcal{A}\mathcal{P}, L_{\text{PTS}})$ where (a) $\mathcal{W}^N = \mathcal{W} \times \dots \times \mathcal{W}$ is the set of states and $\mathbf{x} = [\mathbf{x}_1^T, \mathbf{x}_2^T, \dots, \mathbf{x}_N^T]^T \in \mathcal{W}^N$ is a vector that stacks the positions \mathbf{x}_i of all robots $i \in [N]$; (b) $\mathbf{x}_0 = [\mathbf{x}_1^{0,T}, \dots, \mathbf{x}_N^{0,T}]^T \in \mathcal{W}^N$ represents the initial positions of all robots; (c) $\rightarrow_{\text{PTS}} \subseteq \mathcal{W}^N \times \mathcal{W}^N$ is the transition relation defined by the rule $\frac{\bigwedge_{i \in [N]} (\mathbf{x}_i \rightarrow_i \mathbf{x}'_i)}{\mathbf{x} \rightarrow_{\text{PTS}} \mathbf{x}'}$, $\forall \mathbf{x}, \mathbf{x}' \in \mathcal{W}^N$. In words, there exists a transition from \mathbf{x} to \mathbf{x}' if there exists a transition from \mathbf{x}_i to \mathbf{x}'_i for all $i \in [N]$; (d) $C : \mathcal{W}^N \times \mathcal{W}^N \rightarrow \mathbb{R}^+$ is the cost of traveling from one point to another for all robots. The cost function can be defined arbitrarily as long as it is monotonic. Examples include traveled distance or energy consumption. For holonomic planning, we use the Euclidean distance between $\mathbf{x}, \mathbf{x}' \in \mathcal{W}^N$ as the cost function, i.e.,

$$C(\mathbf{x}, \mathbf{x}') = \|\mathbf{x} - \mathbf{x}'\|_2; \quad (1)$$

(e) $\mathcal{A}\mathcal{P} = \bigcup_{i=1}^N \mathcal{A}\mathcal{P}_i$ is the set of atomic propositions; (f) $L = \bigcup_{i=1}^N L_i : \mathcal{W}^N \rightarrow 2^{\mathcal{A}\mathcal{P}}$ is an observation function that returns the set of atomic propositions satisfied at a state $\mathbf{x} \in \mathcal{W}^N$.

Given the PTS and the NBA B , we can define the *Product Büchi Automaton* (PBA) $P = \text{PTS} \times B$ as follows [11]:

Definition 3.3: (PBA) The *Product Büchi Automaton* is defined by the tuple $P = (\mathcal{Q}_P, \mathcal{Q}_P^0, \rightarrow_P, \mathcal{Q}_P^F, C)$, where (a) $\mathcal{Q}_P = \mathcal{W}^N \times \mathcal{Q}_B$ is an infinite set of states; (b) $\mathcal{Q}_P^0 = \mathbf{x}_0 \times \mathcal{Q}_B^0$ is a set of initial states; (c) $\rightarrow_P \subseteq \mathcal{Q}_P \times 2^{\mathcal{A}\mathcal{P}} \times \mathcal{Q}_P$ is the transition relation defined by the rule: $\frac{(\mathbf{x} \rightarrow_{\text{PTS}} \mathbf{x}') \wedge (\mathcal{Q}_B \xrightarrow{L(\mathbf{x})} \mathcal{Q}_B^F)}{q_P = (\mathbf{x}, q_B) \rightarrow_P q'_P = (\mathbf{x}', q'_B)}$. The transition from the state $q_P \in \mathcal{Q}_P$ to $q'_P \in \mathcal{Q}_P$, is denoted by $(q_P, q'_P) \in \rightarrow_P$, or $q_P \rightarrow_P q'_P$; (d) $\mathcal{Q}_P^F = \mathcal{W}^N \times \mathcal{Q}_B^F$ is a set of accepting/final states; (e) With a slight abuse of notation, C is the cost function between two product states, defined as the cost between respective configurations in \mathcal{W}^N .

In this paper, we restrict our attention to LTL formulas that exclude the ‘next’ temporal operator, denoted by $\text{LTL}_{-\circ}$.² $\text{LTL}_{-\circ}$ formulas are satisfied by discrete plans τ that are infinite sequences of locations of N robots in $\mathcal{W}_{\text{free}}^N$, i.e., $\tau = \tau(1)\tau(2)\dots\tau(k)\dots$, where $\tau(k) \in \mathcal{W}_{\text{free}}^N$. If an $\text{LTL}_{-\circ}$ formula is satisfiable, then there exists a plan τ that can be written in a finite representation, called *prefix-suffix* structure, i.e., $\tau = \tau^{\text{pre}}[\tau^{\text{suff}}]^\omega$ where the prefix part $\tau^{\text{pre}} = \tau(1)\tau(2)\dots\tau(K)$ is executed only once followed by the indefinite execution of the suffix part $\tau^{\text{suff}} = \tau(K)\tau(K+1)\dots\tau(K+S)\tau(K+S+1)$, where $\tau(K+S+1) = \tau(K)$ [11]. A discrete plan τ satisfies ϕ if the trace generated by τ , defined as $\text{trace}(\tau) := L(\tau(1))\dots L(\tau(K))[L(\tau(K))\dots L(\tau(K+S+1))]^\omega$, belongs

²The syntax of $\text{LTL}_{-\circ}$ is the same as the syntax of LTL excluding the ‘next’ operator. The choice of $\text{LTL}_{-\circ}$ over LTL is motivated by the fact that we are interested in the continuous time execution of the synthesized plans, in which case the next operator is not meaningful. For example, to satisfy the formula $\diamond(\pi_1^{\ell_1} \wedge \circ\pi_1^{\ell_2})$, robot 1 should visit region ℓ_2 right after visiting ℓ_1 , which is not possible in continuous time since this requires instantaneous motion. This choice of $\text{LTL}_{-\circ}$ is common in relevant works, see, e.g., [46] and the references therein.

to $\text{Words}(\phi)$. We denote $\text{trace}(\tau) \in \text{Words}(\phi)$ by $\tau \models \phi$.

Given a discrete plan τ , either finite or infinite, a continuous path $\tilde{\tau}$ can be generated by joining any two consecutive points in τ using a line segment. In this sense, the continuous path $\tilde{\tau}$ essentially captures the execution of the discrete plan τ . Let $\tilde{\tau} : [0, 1] \rightarrow \mathcal{W}_{\text{free}}^N$ be a parameterized continuous path, where $\tilde{\tau}(0) = \mathbf{x}_0$. We denote its trace by $\text{trace}(\tilde{\tau}) := \{L(\tilde{\tau}(s_i))\}_{i=0}^n$, where $s_{i+1} = \inf\{s \in [s_i, 1] \mid L(\tilde{\tau}(s)) \neq L(\tilde{\tau}(s_i))\}$ and $s_0 = 0$, $L(\tilde{\tau}(s_n)) = L(\tilde{\tau}(1))$. Like the prefix-suffix structure in the discrete plan τ , a continuous path $\tilde{\tau}$ satisfying ϕ can be written as $\tilde{\tau} = \tilde{\tau}_1 | \tilde{\tau}_2 | \tilde{\tau}_3 | \dots$, where $\tilde{\tau}_1$ and $\tilde{\tau}_2$ are prefix and suffix paths, respectively, $\tilde{\tau}_1(1) = \tilde{\tau}_2(0) = \tilde{\tau}_2(1)$, and $|$ stands for the concatenation of two paths. Hence, $\text{trace}(\tilde{\tau}) = \text{trace}(\tilde{\tau}_1) [\text{trace}(\tilde{\tau}_2)]^\omega$. The transition relation defined before between any two configurations \mathbf{x} and \mathbf{x}' ensures that if τ satisfies ϕ , so does $\tilde{\tau}$. Thus, we can focus on finding discrete feasible plans τ for control.

Furthermore, we define the *cost* of a continuous path $\tilde{\tau}$ as [6]

$$J(\tilde{\tau}) = \sup_{\{m \in \mathbb{N}, 0 = s_0 < s_1 < \dots < s_m = 1\}} \sum_{i=1}^m C(\tilde{\tau}(s_i), \tilde{\tau}(s_{i-1})). \quad (2)$$

In essence, this cost is equal to the Euclidean distance traversed by the continuous path $\tilde{\tau}$ in \mathcal{W}^N . The cost of the discrete plan τ corresponding to $\tilde{\tau}$ can be defined equivalently, as the length of a piecewise linear path in \mathcal{W}^N . Specifically, we define the cost of a prefix-suffix plan $\tau = \tau^{\text{pre}} [\tau^{\text{suf}}]^\omega$ as:

$$J(\tau) = wJ(\tau^{\text{pre}}) + (1-w)J(\tau^{\text{suf}}), \quad (3)$$

where $w \in [0, 1]$ is a user-specified parameter that assigns different weights to the cost associated with the prefix plan and recurrent cost of the suffix plan.

The problem addressed in this paper is stated as follows.

Problem 1: Given the product space P , composed of the transition system of robots and the NBA derived from a global LTL $_{\text{O}}$ specification ϕ defined over the set \mathcal{AP} , determine a plan τ that minimizes $J(\tau)$ and satisfies ϕ .

IV. SAMPLING-BASED OPTIMAL CONTROL SYNTHESIS

In this section, we propose a sampling-based algorithm to solve Problem 1, called TL-RRT* for Temporal Logic RRT*, that incrementally constructs trees that explore the product state space $\mathcal{Q}_P = \mathcal{W}^N \times \mathcal{Q}_B$. These trees are used to design a feasible plan $\tau = \tau^{\text{pre}} [\tau^{\text{suf}}]^\omega$, which satisfies the desired LTL formula ϕ and minimizes the cost function (3). Our algorithm, inspired by the RRT* method for point-to-point navigation [6], is described in Alg. 1. In Alg. 1, first the LTL formula is translated to an NBA $B = \{\mathcal{Q}_B, \mathcal{Q}_B^0, \rightarrow_B, \mathcal{Q}_B^F\}$ [line 1, Alg. 1]. Then, in lines 2-6, candidate prefix plans $\tau^{\text{pre},a}$ are constructed, followed by the construction of corresponding suffix plans $\tau^{\text{suf},a}$ in lines 7-17. Finally, the optimal plan $\tau = \tau^{\text{pre},a^*} [\tau^{\text{suf},a^*}]^\omega \models \phi$ is synthesized in lines 18-19.

A. Construction of Prefix Plans

In this section, we describe how the prefix plan is constructed [lines 2-6, Alg. 1] by building incrementally a tree, denoted by $\mathcal{T} = \{\mathcal{V}_{\mathcal{T}}, \mathcal{E}_{\mathcal{T}}, \text{Cost}\}$, that explores the state space \mathcal{Q}_P . The set $\mathcal{V}_{\mathcal{T}}$ consists of nodes $q_P = (\mathbf{x}, q_B) \in \mathcal{W}_{\text{free}}^N \times \mathcal{Q}_B$. The set of edges $\mathcal{E}_{\mathcal{T}}$ captures transitions among the nodes in $\mathcal{V}_{\mathcal{T}}$, i.e., $(q_P, q'_P) \in \mathcal{E}_{\mathcal{T}}$, if $q_P \rightarrow_P q'_P$, for $q_P, q'_P \in \mathcal{V}_{\mathcal{T}}$. Moreover, the function $\text{Cost} : \mathcal{V}_{\mathcal{T}} \rightarrow \mathbb{R}^+$ assigns the cost of

Algorithm 1: TL-RRT*

Input: Product transition system PTS, LTL formula ϕ , maximum numbers of iterations $n_{\text{max}}^{\text{pre}}, n_{\text{max}}^{\text{suf}}$

Output: Optimal plans $\tau \models \phi$

- 1 Convert ϕ to an NBA $B = (\mathcal{Q}_B, \mathcal{Q}_B^0, \rightarrow_B, \mathcal{Q}_B^F)$;
 - ; \triangleright Construction of Prefix Plans IV-A
- 2 Define goal set: $\mathcal{Q}_{\text{goal}}^{\text{pre}}$;
- 3 Initial state: $q_P^0 = (\mathbf{x}^0, q_B^0)$;
- 4 $[\mathcal{T}, \mathcal{P}] = \text{ConstructTree}(\mathcal{Q}_{\text{goal}}^{\text{pre}}, \text{PTS}, B, q_P^0, n_{\text{max}}^{\text{pre}})$;
- 5 **for** $a = 1 : |\mathcal{P}|$ **do**
 - 6 | $\tau^{\text{pre},a} = \text{FindPlan}(\mathcal{T}, q_P^0, \mathcal{P}(a))$;
 - ; \triangleright Construction of Suffix Plans IV-B
- 7 **for** $a = 1 : |\mathcal{P}|$ **do**
 - 8 | Initial state: $q_P^0 = \mathcal{P}(a)$;
 - 9 | Define goal set: $\mathcal{Q}_{\text{goal}}^{\text{suf}}(q_P^0)$;
 - 10 | **if** $q_P^0 \in \mathcal{Q}_{\text{goal}}^{\text{suf}}$ **then**
 - 11 | | $\mathcal{T} = (\{q_P^0\}, \{q_P^0, q_P^0\}, 0)$, $\mathcal{S}_a = \{q_P^0\}$;
 - 12 | **else**
 - 13 | | $[\mathcal{T}, \mathcal{S}_a] =$
 - 14 | | | $\text{ConstructTree}(\mathcal{Q}_{\text{goal}}^{\text{suf}}, \text{PTS}, B, q_P^0, n_{\text{max}}^{\text{suf}})$;
 - 15 | **for** $e = 1 : |\mathcal{S}_a|$ **do**
 - 16 | | $\tau_a^{\text{suf},e} = \text{FindPlan}(\mathcal{T}, q_P^0, \mathcal{S}_a(e))$;
 - 17 | | $e^* = \text{argmin}_e (J(\tau_a^{\text{suf},e}))$;
 - 18 | | $\tau_a^{\text{suf},a} = \tau_a^{\text{suf},e^*}$;
 - ; \triangleright Construction of optimal plans IV-C
- 19 $a^* = \text{argmin}_a (J(\tau^{\text{pre},a}) + J(\tau^{\text{suf},a}))$;
- 20 **return** $\tau = \tau^{\text{pre},a^*} [\tau^{\text{suf},a^*}]^\omega$;

reaching node q_P from the root. Let τ denote the plan/path from the root to q_P in the tree, then $\text{Cost}(q_P) = J(\tau)$.³

The tree \mathcal{T} is rooted at the initial state $q_P^0 = (\mathbf{x}^0, q_B^0)$ where $\mathbf{x}^0 \in \mathcal{W}_{\text{free}}^N$ and $q_B^0 \in \mathcal{Q}_B^0$.⁴ We define the goal region as

$$\mathcal{Q}_{\text{goal}}^{\text{pre}} = \{q_P = (\mathbf{x}, q_B) \in \mathcal{W}_{\text{free}}^N \times \mathcal{Q}_B \mid q_B \in \mathcal{Q}_B^F\}, \quad (4)$$

i.e., the goal region collects all final states of the PBA. To construct a prefix plan, the tree \mathcal{T} attempts to find a sequence of states in $\mathcal{W}_{\text{free}}^N \times \mathcal{Q}_B$ that connects an initial state q_P^0 to any *accepting* state in $\mathcal{Q}_{\text{goal}}^{\text{pre}}$. In Alg. 2, the set $\mathcal{V}_{\mathcal{T}}$ initially contains only the initial state q_P^0 . The set of edges is initialized as $\mathcal{E}_{\mathcal{T}} = \emptyset$. By convention, the cost of q_P^0 is zero [line 1, Alg. 2]. Below we describe the incremental construction of the tree $\mathcal{T} = \{\mathcal{V}_{\mathcal{T}}, \mathcal{E}_{\mathcal{T}}, \text{Cost}\}$ [lines 2-12, Alg. 2].

1) *Sampling a state \mathbf{x}^{rand} :* The first step to construct the tree \mathcal{T} is to sample a state $\mathbf{x}^{\text{rand}} \in \mathcal{W}_{\text{free}}^N$. We generate independently one sample from $\mathcal{W}_{\text{free}}$ for each robot and stack these samples to construct \mathbf{x}^{rand} . Any distribution can be used to get \mathbf{x}^{rand} as long as it is bounded away from zero on $\mathcal{W}_{\text{free}}$.

2) *Constructing a state $\mathbf{x}^{\text{new}} \in \mathcal{W}_{\text{free}}^N$:* Given the state $\mathbf{x}^{\text{rand}} \in \mathcal{W}_{\text{free}}^N$, we construct the state $\mathbf{x}^{\text{new}} \in \mathcal{W}_{\text{free}}^N$. To do so, we first find the set $\mathcal{Q}_P^{\text{nearest}} \subseteq \mathcal{V}_{\mathcal{T}}$, which consists of nodes $q_P^{\text{nearest}} = (\mathbf{x}^{\text{nearest}}, q_B^{\text{nearest}}) \in \mathcal{V}_{\mathcal{T}}$, that are the closest to \mathbf{x}^{rand} in terms of a given distance function [line 5, Alg. 2].

³We also refer a path to a sequence of connecting nodes in a graph, which can be distinguished easily from the continuous path of a discrete plan.

⁴In what follows, for the sake of simplicity we assume that the set of initial states \mathcal{Q}_B^0 consists of only one state, i.e., $\mathcal{Q}_B^0 = \{q_B^0\}$. In case there are more than one initial states in B , then the lines 2-6 in Alg. 1 should be executed for each possible initial state q_B^0 .

Algorithm 2: Function $\text{ConstructTree}(Q_{\text{goal}}, \text{PTS}, B, q_P^0, n_{\text{max}})$

```

1  $\mathcal{V}_{\mathcal{T}} = \{q_P^0\}; \mathcal{E}_{\mathcal{T}} = \emptyset; \text{Cost}(q_P^0) = 0; \mathbf{x}^{\text{new}} = \mathbf{x}_0;$ 
2 for  $n = 1 : n_{\text{max}}$  do
3   while  $\mathbf{x}^{\text{new}} \notin \mathcal{W}_{\text{free}}^N$  do
4      $\mathbf{x}^{\text{rand}} = \text{Sample}(\mathcal{W}_{\text{free}});$ 
5      $\{(\mathbf{x}^{\text{nearest}}, q_B^{\text{nearest}})\} = \text{Nearest}(\mathcal{T}, \mathbf{x}^{\text{rand}});$ 
6      $\mathbf{x}^{\text{new}} = \text{Steer}(\mathbf{x}^{\text{nearest}}, \mathbf{x}^{\text{rand}});$ 
7      $Q_P^{\text{near}} = \text{Near}(\mathbf{x}^{\text{new}}, \mathcal{T}) \cup \{(\mathbf{x}^{\text{nearest}}, q_B^{\text{nearest}})\};$ 
8     for  $b = 1 : |Q_B|$  do
9        $q_B^{\text{new}} = Q_B(b), \quad q_P^{\text{new}} = (\mathbf{x}^{\text{new}}, q_B^{\text{new}});$ 
10       $[\mathcal{T}, \text{added}] = \text{Extend}(q_P^{\text{new}}, Q_P^{\text{near}}, \mathcal{T});$ 
11      if  $\text{added} = 1$  then
12         $[\mathcal{T}, \text{Cost}] = \text{Rewire}(q_P^{\text{new}}, Q_P^{\text{near}}, \mathcal{T});$ 
13  $\mathcal{P} = \mathcal{V}_{\mathcal{T}} \cap Q_{\text{goal}};$ 
14 return  $\mathcal{T}, \mathcal{P};$ 

```

Here we employ the Euclidean distance. Note that Q_P^{near} may include multiple nodes with the same $\mathbf{x}^{\text{nearest}}$ but different q_B^{nearest} . We identify the nodes in Q_P^{near} using the function $\text{Nearest} : \mathcal{W}_{\text{free}}^N \rightarrow 2^{\mathcal{V}_{\mathcal{T}}}$ defined as

$$\text{Nearest}(\mathbf{x}^{\text{rand}}) = \underset{q_P \in \mathcal{V}_{\mathcal{T}}}{\text{argmin}} \left\| \mathbf{x}^{\text{rand}} - q_P|_{\mathcal{W}_{\text{free}}^N} \right\|, \quad (5)$$

where $(\cdot)|_{(\cdot)}$ stands for the projection, i.e., $q_P|_{\mathcal{W}_{\text{free}}^N} \in \mathcal{W}_{\text{free}}^N$.

Next, given the state $\mathbf{x}^{\text{nearest}} \in \mathcal{W}_{\text{free}}^N$, we construct the state $\mathbf{x}^{\text{new}} \in \mathcal{W}_{\text{free}}^N$ using the function $\text{Steer} : \mathcal{W}_{\text{free}}^N \times \mathcal{W}_{\text{free}}^N \rightarrow \mathcal{W}^N$ [line 6, Alg. 2]. Specifically, the state \mathbf{x}^{new} is within distance at most $\eta > 0$ from $\mathbf{x}^{\text{nearest}}$, i.e., $\|\mathbf{x}^{\text{nearest}} - \mathbf{x}^{\text{new}}\| \leq \eta$, where η is user-specified, and should also satisfy $\|\mathbf{x}^{\text{new}} - \mathbf{x}^{\text{rand}}\| \leq \|\mathbf{x}^{\text{nearest}} - \mathbf{x}^{\text{rand}}\|$. The step size η should consider the clutter in the environment, the maximum control input, and the duration of the time unit when tracking the path using a robot controller. Furthermore, we check if \mathbf{x}^{new} lies in the obstacle-free space, i.e., $\mathbf{x}^{\text{new}} \in \mathcal{W}_{\text{free}}^N$. If not, we return to the Sample step.

3) *Constructing the node set Q_P^{near} :* Given $\mathbf{x}^{\text{new}} \in \mathcal{W}_{\text{free}}^N$, we construct the set $Q_P^{\text{near}} \subseteq \mathcal{V}_{\mathcal{T}}$ [line 7, Alg. 2] that is the union of the set Q_P^{near} and the set

$$\text{Near}(\mathbf{x}^{\text{new}}, \mathcal{T}) = \{q_P = (\mathbf{x}, q_B) \in \mathcal{V}_{\mathcal{T}} \mid \|\mathbf{x} - \mathbf{x}^{\text{new}}\| \leq r_n(\mathcal{V}_{\mathcal{T}})\}, \quad (6)$$

which collects nodes that lie within a certain distance $r_n(\mathcal{V}_{\mathcal{T}})$ from \mathbf{x}^{new} , where the connection radius $r_n(\mathcal{V}_{\mathcal{T}})$ is selected as

$$r_n(\mathcal{V}_{\mathcal{T}}) = \min \left\{ \gamma_{\text{TL-RRT}^*} \left(\frac{\log |\mathcal{V}_{\mathcal{T}}|}{|\mathcal{V}_{\mathcal{T}}|} \right)^{\frac{1}{\text{dim}+1}}, \eta \right\}. \quad (7)$$

In (7), $\eta > 0$, and $|\mathcal{V}_{\mathcal{T}}|_{\sim}$ is the equivalence class defined as $|\mathcal{V}_{\mathcal{T}}|_{\sim} = \{\mathbf{x} \mid (\mathbf{x}, q_B) \in \mathcal{V}_{\mathcal{T}}, \forall q_B \in Q_B\}$, i.e., $|\mathcal{V}_{\mathcal{T}}|_{\sim}$ collects all \mathbf{x} 's appearing in the tree. Also, $|\cdot|$ is the cardinality of a set, dim is the dimension of the product workspace \mathcal{W}^N , i.e., $\text{dim} = d \cdot N$, and $\gamma_{\text{TL-RRT}^*}$ is a constant as

$$\gamma_{\text{TL-RRT}^*} \geq (2 + \theta) \left(\frac{(1 + \epsilon/4)J(\tau^*)}{(\text{dim} + 1)\theta(1 - \kappa)} \cdot \frac{\mu(\mathcal{W}_{\text{free}}^N)}{\zeta_{\text{dim}}} \right)^{\frac{1}{\text{dim}+1}}, \quad (8)$$

where $\epsilon \in (0, 1)$, $\theta \in (0, 1/4)$ and $\kappa \in (0, 1)$ are three ratio terms, τ^* is the optimal path, $\tau^{*,\text{pre}}$ for the prefix part and $\tau^{*,\text{suf}}$ for the suffix part, $\mu(\mathcal{W}_{\text{free}}^N)$ is the measure of $\mathcal{W}_{\text{free}}^N$, and ζ_{dim} is the volume of the unit sphere in \mathbb{R}^{dim} . Intuitively, as the size of free workspace or the cost of the optimal path

Algorithm 3: Function $\text{Extend}(q_P^{\text{new}}, Q_P^{\text{near}}, \mathcal{T})$

```

1  $\text{added} = 0; c = \infty$  for all
    $q_P^{\text{near}} = (\mathbf{x}^{\text{near}}, q_B^{\text{near}}) \in Q_P^{\text{near}}$  do
2   if  $\mathbf{x}^{\text{near}} \rightarrow \mathbf{x}^{\text{new}} \wedge (q_B^{\text{near}}, L(\mathbf{x}^{\text{near}}), q_B^{\text{new}}) \in \rightarrow_B$  then
3      $c' = \text{Cost}(q_P^{\text{near}}) + C(\mathbf{x}^{\text{near}}, \mathbf{x}^{\text{new}});$ 
4     if  $c' \leq c$  then
5        $\text{added} = 1; q_P^{\text{min}} = q_P^{\text{near}};$ 
6        $c = c'; \text{Cost}(q_P^{\text{new}}) = c';$ 
7 if  $\text{added} = 1$  then
8    $\mathcal{V}_{\mathcal{T}} = \mathcal{V}_{\mathcal{T}} \cup \{q_P^{\text{new}}\}, \mathcal{E}_{\mathcal{T}} = \mathcal{E}_{\mathcal{T}} \cup \{(q_P^{\text{min}}, q_P^{\text{new}})\};$ 
9 return  $\mathcal{T}, \text{added};$ 

```

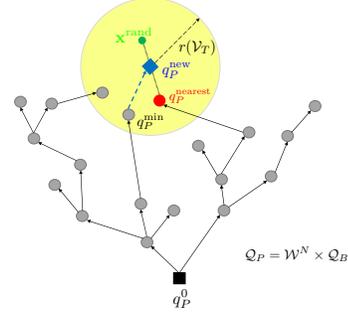


Fig. 1. Graphical depiction of Alg. 3. The black square is the root of the tree and the gray and red disks represent nodes in the set $\mathcal{V}_{\mathcal{T}}$. The green dot is the state \mathbf{x}^{rand} generated by function Sample . The blue diamond and the red dot stand for states q_P^{new} and q_P^{min} , respectively. q_P^{new} is connected to the tree \mathcal{T} through the node q_P^{min} that incurs the smallest cost from the root (see the blue dashed edge), which may not be the node q_P^{nearest} .

become larger, $\gamma_{\text{TL-RRT}^*}$ also becomes larger, which leads to a larger connection radius so that more nodes are considered when adding the new state to the tree and further improving the tree, thus covering more space; see the following steps 4) and 5). The lower bound of $\gamma_{\text{TL-RRT}^*}$ is the same as that in [47] and will be discussed in Theorem 6.5.

After obtaining \mathbf{x}^{new} and Q_P^{near} , we pair a Büchi state q_B^{new} with \mathbf{x}^{new} to construct a state $q_P^{\text{new}} = (\mathbf{x}^{\text{new}}, q_B^{\text{new}})$ which will be examined if it can be added to the tree \mathcal{T} through nodes in Q_P^{near} . The addition is accomplished by the function Extend described in Alg. 3. This procedure is repeated for all states $(\mathbf{x}^{\text{new}}, Q_B(b))$ [line 8, Alg. 2], where $Q_B(b)$ denotes the b -th state in the set Q_B for $b \in [|Q_B|]$, assuming an arbitrary enumeration of the elements in Q_B [line 9, Alg. 2]. Appending $q_B^{\text{new}} = Q_B(b)$ to \mathbf{x}^{new} , we construct the state $q_P^{\text{new}} = (\mathbf{x}^{\text{new}}, q_B^{\text{new}})$. In what follows, we describe the function Extend for a given state q_P^{new} , which is illustrated in Fig. 1.

4) *Extending \mathcal{T} towards q_P^{new} :* In lines 1-6 of Alg. 3, we select the parent of the state q_P^{new} among Q_P^{near} . Specifically, for each state $q_P^{\text{near}} = (\mathbf{x}^{\text{near}}, q_B^{\text{near}}) \in Q_P^{\text{near}}$, we check (i) if $\mathbf{x}^{\text{near}} \rightarrow \mathbf{x}^{\text{new}}$, and (ii) if $(q_B^{\text{near}}, L(\mathbf{x}^{\text{near}}), q_B^{\text{new}}) \in \rightarrow_B$ [line 2, Alg. 3]. In words, we check whether q_P^{near} is a candidate parent of q_P^{new} . If so, the cost of the state q_P^{new} is $\text{Cost}(q_P^{\text{near}}) + C(\mathbf{x}^{\text{near}}, \mathbf{x}^{\text{new}})$ [line 3, Alg. 3], where $\text{Cost}(q_P^{\text{near}})$ is the cost of reaching q_P^{near} from q_P^0 , and $C(\mathbf{x}^{\text{near}}, \mathbf{x}^{\text{new}})$ from (1) is the cost of reaching q_P^{new} from q_P^{near} . Among all candidate parents q_P^{near} of q_P^{new} , we pick the one that incurs the minimum cost $\text{Cost}(q_P^{\text{new}})$, denoted by q_P^{min} [lines 4-6, Alg. 3]. If the set of candidate parents is not empty, then the sets $\mathcal{V}_{\mathcal{T}}, \mathcal{E}_{\mathcal{T}}$ are updated as:

Algorithm 4: Function $\text{Rewire}(q_P^{\text{new}}, Q_P^{\text{near}}, \mathcal{T})$

```

1 for all  $q_P^{\text{near}} = (\mathbf{x}_P^{\text{near}}, q_B^{\text{near}}) \in Q_P^{\text{near}}$  do
2   if  $\mathbf{x}^{\text{new}} \rightarrow \mathbf{x}^{\text{near}} \wedge (q_B^{\text{new}}, L(\mathbf{x}^{\text{new}}), q_B^{\text{near}}) \in \rightarrow_B$  then
3     if  $\text{Cost}(q_P^{\text{near}}) > \text{Cost}(q_P^{\text{new}}) + C(\mathbf{x}^{\text{new}}, \mathbf{x}^{\text{near}})$ 
4       then
5          $\text{Cost}(q_P^{\text{near}}) = \text{Cost}(q_P^{\text{new}}) + C(\mathbf{x}^{\text{new}}, \mathbf{x}^{\text{near}})$ ;
6          $\mathcal{E}_{\mathcal{T}} = \mathcal{E}_{\mathcal{T}} \setminus \{(\text{Parent}(q_P^{\text{near}}), q_P^{\text{near}})\}$ ;
7          $\mathcal{E}_{\mathcal{T}} = \mathcal{E}_{\mathcal{T}} \cup \{(q_P^{\text{new}}, q_P^{\text{near}})\}$ ;
8   return  $\mathcal{T}, \text{Cost}$ ;
```

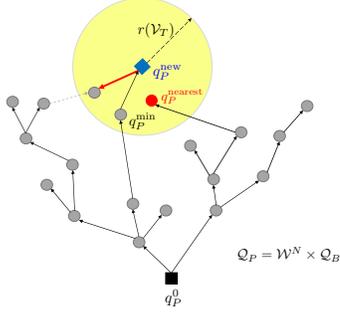


Fig. 2. Graphical depiction of Alg. 4. The dashed gray arrow stands for the edge that will be deleted from the set $\mathcal{E}_{\mathcal{T}}$ and the thick red arrow stands for the new edge that will be added to $\mathcal{E}_{\mathcal{T}}$ during the execution of Alg. 4.

$\mathcal{V}_{\mathcal{T}} = \mathcal{V}_{\mathcal{T}} \cup \{q_P^{\text{new}}\}$ and $\mathcal{E}_{\mathcal{T}} = \mathcal{E}_{\mathcal{T}} \cup \{(q_P^{\text{min}}, q_P^{\text{new}})\}$ [lines 8, Alg. 3]. If q_P^{new} is added to $\mathcal{V}_{\mathcal{T}}$, the *rewiring* process follows.

5) *Rewiring through q_P^{new}* : Once a new state $q_P^{\text{new}} = (\mathbf{x}^{\text{new}}, q_B^{\text{new}})$ has been added to the tree, we *rewire* the nodes in $q_P^{\text{near}} \in Q_P^{\text{near}}$ through the node q_P^{new} [line 12, Alg. 2] if this will decrease the cost $\text{Cost}(q_P^{\text{near}})$, i.e., if this will decrease the cost of reaching the node q_P^{near} starting from the root q_P^0 . The rewiring process is described in Alg. 4; see also Fig. 2.

Specifically, for each $q_P^{\text{near}} \in Q_P^{\text{near}}$, we first check if a transition from q_P^{new} to q_P^{near} is *feasible*, i.e., if $\mathbf{x}^{\text{new}} \rightarrow \mathbf{x}^{\text{near}}$ and if $(q_B^{\text{new}}, L(\mathbf{x}^{\text{new}}), q_B^{\text{near}}) \in \rightarrow_B$. If so, then we check if the cost of q_P^{near} will decrease due to rewiring through q_P^{new} , i.e., if $\text{Cost}(q_P^{\text{near}}) > \text{Cost}(q_P^{\text{new}}) + C(\mathbf{x}^{\text{new}}, \mathbf{x}^{\text{near}})$ [line 3, Alg. 4]. If so, then the cost of q_P^{near} is updated as $\text{Cost}(q_P^{\text{new}}) + C(\mathbf{x}^{\text{new}}, \mathbf{x}^{\text{near}})$ [line 4, Alg. 4], and the parent of q_P^{near} becomes the node q_P^{new} [line 5-6, Alg. 4], where the function $\text{Parent} : \mathcal{V}_{\mathcal{T}} \rightarrow \mathcal{V}_{\mathcal{T}}$ maps a node to the unique parent node. By convention, we assume that $\text{Parent}(q_P^0) = q_P^0$, where q_P^0 is the root of the tree \mathcal{T} .

6) *Computing paths*: The construction of the tree \mathcal{T} ends after $n_{\text{max}}^{\text{pre}}$ iterations, where $n_{\text{max}}^{\text{pre}}$ is user-specified [line 2, Alg. 2]. Then we construct the set $\mathcal{P} = \mathcal{V}_{\mathcal{T}} \cap Q_{\text{goal}}^{\text{pre}}$ [line 13, Alg. 2] that collects all states $q_P \in \mathcal{V}_{\mathcal{T}}$ that belong to the goal region $Q_{\text{goal}}^{\text{pre}}$. Given the tree \mathcal{T} and the set \mathcal{P} [line 4, Alg. 1], we can compute the prefix plans [lines 5-6, Alg. 1]. In particular, the path that connects the a -th state in the set \mathcal{P} , denoted by $\mathcal{P}(a)$, to the root q_P^0 constitutes the a -th prefix plan and is denoted by $\tau^{\text{pre},a}$ [function FindPlan in line 6, Alg. 1]. To compute $\tau^{\text{pre},a}$, only the parent of each node in \mathcal{T} is required, due to the tree structure of \mathcal{T} . Notice that the computational complexity of computing the prefix plan $\tau^{\text{pre},a}$ in the tree \mathcal{T} is $O(|\mathcal{V}_{\mathcal{T}}|)$. On the other hand, if the product state-space $\mathcal{W}^N \times Q_B$ was searched by a graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ of arbitrary structure, as in [40]–[42], then the

computational complexity of the Dijkstra algorithm to find the optimal plan $\tau^{\text{pre},a}$ that connects the state $\mathcal{P}(a)$ to the root would be $O(|\mathcal{E}| + |\mathcal{V}| \log(|\mathcal{V}|))$, where $|\mathcal{E}| + |\mathcal{V}| \log(|\mathcal{V}|) > |\mathcal{V}|$.

B. Construction of Suffix Plans

Once the prefix plans $\tau^{\text{pre},a}$ are constructed, the construction of the respective suffix plans $\tau^{\text{suf},a}$ follows [lines 7-17, Alg. 1]. The suffix plan $\tau^{\text{suf},a}$ is a sequence of states in $\mathcal{W}_{\text{free}}^N \times Q_B$ that starts from the state $\mathcal{P}(a)$, $a \in [|\mathcal{P}|]$, and ends at the same state $\mathcal{P}(a)$, i.e., a cycle around state $\mathcal{P}(a)$. Projection of this sequence on $\mathcal{W}_{\text{free}}^N$ gives the suffix plan $\tau^{\text{suf},a}$.

For this purpose, we build a tree $\mathcal{T} = \{\mathcal{V}_{\mathcal{T}}, \mathcal{E}_{\mathcal{T}}, \text{Cost}\}$ similarly as in Section IV-A. The only differences are that: (i) the root of the tree is now $q_P^0 = \mathcal{P}(a)$, i.e., the *accepting* state [line 8, Alg. 1] detected during the construction of the prefix plans, (ii) the goal region, given the root q_P^0 , is defined as

$$Q_{\text{goal}}^{\text{suf}}(q_P^0) = \{q_P = (\mathbf{x}, q_B) \in \mathcal{W}_{\text{free}}^N \times Q_B \mid (q_B, L(\mathbf{x}), q_P^0|_{Q_B}) \in \rightarrow_B \wedge \mathbf{x} \rightarrow q_P^0|_{\mathcal{W}_{\text{free}}^N}\}, \quad (9)$$

i.e., it collects all states q_P from which a transition to the root q_P^0 is feasible, but this transition will not be included in $\mathcal{E}_{\mathcal{T}}$. Note that before constructing trees to compute suffix plans, we first check if $q_P^0 \in Q_{\text{goal}}^{\text{suf}}$, i.e., if $(q_P^0|_{Q_B}, L(q_P^0|_{\mathcal{W}_{\text{free}}^N}), q_P^0|_{Q_B}) \in \rightarrow_B$ [line 10, Alg. 1]. If so, the tree \mathcal{T} is trivial, as it consists of only the root, and a loop around it with zero cost [line 11, Alg. 1]. If $q_P^0 \notin Q_{\text{goal}}^{\text{suf}}$, then the tree \mathcal{T} is constructed by Alg. 2 [line 13, Alg. 1].

Once a tree rooted at $q_P^0 = \mathcal{P}(a)$ is constructed, a set $\mathcal{S}_a \subseteq \mathcal{V}_{\mathcal{T}}$ is constructed that collects all states $q_P \in \mathcal{V}_{\mathcal{T}} \cap Q_{\text{goal}}^{\text{suf}}(q_P^0)$ [lines 14-15, Alg. 1]. Then for each state $q_P \in \mathcal{S}_a$, there exists a suffix plan, denoted by $\tau_a^{\text{suf},e}$, $\forall e \in [|\mathcal{S}_a|]$, and we compute the cost $J(\tau_a^{\text{suf},e})$ using $J(\tau_a^{\text{suf},e}) = \text{Cost}(\mathcal{S}_a(e)) + C(\mathcal{S}_a(e)|_{\mathcal{W}_{\text{free}}^N}, q_P^0|_{\mathcal{W}_{\text{free}}^N})$, where $\mathcal{S}_a(e)$ denotes the e -th state in \mathcal{S}_a . Among all detected suffix plans $\tau_a^{\text{suf},e}$ associated with the *accepting* state $\mathcal{P}(a)$, we pick the one with the minimum cost, which constitutes the suffix plan $\tau^{\text{suf},a}$ [lines 16-17, Alg. 1]. This process is repeated for $a \in [|\mathcal{P}|]$ [line 7, Alg. 1].

Remark 4.1: The execution of the suffix paths requires that the robots visit the suffix waypoints infinitely often. However, in practice, it may be hard to visit the exact locations of these waypoints infinitely often due to possible uncertainties or disturbances in the dynamics or environment. We note that our algorithm TL-RRT* is not sensitive to such disturbances as the robots can visit a neighborhood of each waypoint wherein the same observation is made.

C. Construction of the Optimal Discrete Plan

By construction, any plan $\tau^a = \tau^{\text{pre},a}[\tau^{\text{suf},a}]^{\omega}$, with $\mathcal{S}_a \neq \emptyset$, $a \in \{1, \dots, |\mathcal{P}|\}$ satisfies the global LTL specification ϕ . The cost $J(\tau^a)$ of each plan τ^a is defined in (3). Among all plans $\tau^a \models \phi$, we pick the one with the smallest cost $J(\tau^a)$ denoted by $\tau = \tau^{a^*}$, where $a^* = \text{argmin}_a J(\tau^a)$ [lines 18-19, Alg. 1].

Remark 4.2: The plan τ satisfies the LTL formula if all robots move synchronously by reaching next waypoints at the same time. Future research includes extension of TL-RRT* so that the plans can be executed asynchronously while satisfying the LTL specification, as in [20].

V. BIASED SAMPLING METHOD

In this section, we propose a biased sampling method that biases the construction of the tree towards shortest paths, in

Algorithm 5: Changes made to Alg. 2

```

1 while  $\mathbf{x}^{\text{new}} \notin \mathcal{W}_{\text{free}}^N$  do
2    $[\mathbf{x}^{\text{rand}}, q_P^{\text{closest}}] = \text{BiasedSample}(\mathcal{W}, \mathcal{T})$ ;
3    $\mathbf{x}^{\text{new}} = \text{Steer}(\mathbf{x}^{\text{closest}}, \mathbf{x}^{\text{rand}})$ ;
4    $\mathcal{Q}_P^{\text{near}} = \text{Near}(\mathbf{x}^{\text{new}}, \mathcal{T}) \cup \{q_P^{\text{closest}}\}$ ;

```

terms of the number of transitions, to the final states in the NBA to accelerate the construction of low-cost feasible plans. We call this method biased TL-RRT*. Intuitively, the biased TL-RRT* extends more often nodes in the tree that are closer to accepting states and the new nodes ($\mathbf{x}^{\text{new}}, q_B^{\text{new}}$) that are added to the tree are such that any subsequent nodes added to the tree via them are even closer to the accepting states.

To this end, similar to [30], [31], we first prune the NBA by removing infeasible transitions. Specifically, a transition from a state q_B to q'_B is infeasible if it is enabled by a propositional formula, e.g., $\pi_i^{\ell_j} \wedge \pi_i^{\ell_k}$, that requires a robot to be present at more than one disjoint regions simultaneously.

Moreover, we define a distance function $\rho : \mathcal{Q}_B \times \mathcal{Q}_B \rightarrow \mathbb{N}$ between any two Büchi states in the NBA, which captures the minimum number of transitions in the pruned NBA

$$\rho(q_B, q'_B) = \begin{cases} |SP_{q_B, q'_B}|, & \text{if } SP_{q_B, q'_B} \text{ exists,} \\ \infty, & \text{otherwise,} \end{cases} \quad (10)$$

where SP_{q_B, q'_B} denotes the shortest path in the pruned NBA from q_B to q'_B and $|SP_{q_B, q'_B}|$ is the number of hops/transitions in this shortest path. Using this metric, we can define a feasible accepting Büchi state $q_B^F \in \mathcal{Q}_B^F$ as (i) $\rho(q_B^0, q_B^F) \neq \infty$ and (ii) $\rho(q_B^F, q_B^F) \neq \infty$. If there are multiple feasible accepting states, we randomly select one, denoted by $q_B^{F, \text{feas}}$, and use it throughout Alg. 2. Moreover, given the tree \mathcal{T} , we define the set $\mathcal{D}_{\min} \subseteq \mathcal{V}_{\mathcal{T}}$ that collects the nodes $q_P \in \mathcal{V}_{\mathcal{T}}$ which have the minimum distance $\rho(q_B, q_B^{F, \text{feas}})$. Intuitively, these nodes are the “closest” to the accepting states that have Büchi state $q_B^{F, \text{feas}}$. More details can be found in [30], [31]. Alternatively, a distance metric over automata can also be defined in terms of the number of sets of atomic propositions that still need to be satisfied until reaching an accepting state [48]. Note that any distance metric over automata can be used with our biased sampling method.

A. Construction of Prefix Plans

The biased sampling-based algorithm for the prefix plans is similar to Alg. 2 except that the selection of \mathbf{x}^{rand} and the construction of the set $\mathcal{Q}_P^{\text{near}}$ [lines 3-7, Alg. 2] are now determined by Alg. 5. To sample a state \mathbf{x}^{rand} , we employ the function `BiasedSample` in Alg. 6.

1) *Selection of $q_P^{\text{closest}} = (\mathbf{x}^{\text{closest}}, q_B^{\text{closest}}) \in \mathcal{V}_{\mathcal{T}}$* [line 1, Alg. 6]: We select a node q_P^{closest} from the tree to expand. Specifically, selection of q_P^{closest} is biased towards nodes that are closer to the accepting states, i.e., nodes in the set \mathcal{D}_{\min} . First, we sort the nodes in the sets \mathcal{D}_{\min} and $\mathcal{V}_{\mathcal{T}} \setminus \mathcal{D}_{\min}$ in the opposite order that they were added to the tree. Then the point q_P^{closest} is sampled from a discrete probability function $f_{\text{closest}}(q_P | \mathcal{V}_{\mathcal{T}}, \mathcal{D}_{\min}) : \mathcal{V}_{\mathcal{T}} \rightarrow [0, 1]$, defined as:

$$f_{\text{closest}}(q_P = i | \mathcal{V}_{\mathcal{T}}, \mathcal{D}_{\min}) = \begin{cases} p_{\text{closest}} \mathbb{P}_{\text{UG}}(q_P = i; \frac{1}{|\mathcal{D}_{\min}|}), & \text{if } i \in \mathcal{D}_{\min} \\ (1 - p_{\text{closest}}) \mathbb{P}_{\text{UG}}(q_P = i; \frac{1}{|\mathcal{V}_{\mathcal{T}} \setminus \mathcal{D}_{\min}|}), & \text{if } i \in \mathcal{V}_{\mathcal{T}} \setminus \mathcal{D}_{\min}. \end{cases} \quad (11)$$

Algorithm 6: Function BiasedSample(\mathcal{W}, \mathcal{T})

```

1 Sample  $q_P^{\text{closest}} = (\mathbf{x}^{\text{closest}}, q_B^{\text{closest}}) \in \mathcal{V}_{\mathcal{T}}$  by (11);
2 Sample  $q_B^{\text{succ},1}, q_B^{\text{succ},2} \in q_B^{\text{closest}}$ ;
3 if  $q_B^{\text{succ},1}$  and  $q_B^{\text{succ},2}$  exist then
4   Select  $\sigma^*$  enabling  $q_B^{\text{succ},1} \rightarrow_B q_B^{\text{succ},2}$ ;
5   Sample  $\mathbf{x}^{\text{rand}} = [\mathbf{x}_1^{\text{rand},T}, \dots, \mathbf{x}_N^{\text{rand},T}]^T$  by (14);
6 else
7   return  $\mathbf{x}^{\text{rand}} = \emptyset, q_P^{\text{closest}} = \emptyset$ ;
8 return  $\mathbf{x}^{\text{rand}}, q_P^{\text{closest}}$ ;

```

where $p_{\text{closest}} \in (0.5, 1)$ is a user-specified parameter denoting the probability of selecting any node in \mathcal{D}_{\min} to be q_P^{closest} . \mathbb{P}_{UG} denotes the UG distribution [49], which compounds the uniform and geometric distributions. Given a countably infinite set \mathbb{N}^+ and a parameter p such that $0 < p < 1$, the probability mass function of a UG distribution is defined as

$$\mathbb{P}_{\text{UG}}(q_P = i; p) = \sum_{n=i}^{\infty} \frac{1}{n} p(1-p)^{n-1}, \quad i \in \mathbb{N}^+. \quad (12)$$

To apply the UG distribution to \mathcal{D}_{\min} , we compute the probability for $q_P \in [|\mathcal{D}_{\min}|]$ and then normalize so that the probabilities add to 1. The UG distribution has the property:

$$\mathbb{P}_{\text{UG}}(q_P = i + 1) = \mathbb{P}_{\text{UG}}(q_P = i) - \frac{p(1-p)^{i-1}}{i}, \quad (13)$$

where $\mathbb{P}_{\text{UG}}(q_P = 1) = -\frac{p \log p}{1-p}$. We have $\mathbb{P}_{\text{UG}}(q_P = i + 1) < \mathbb{P}_{\text{UG}}(q_P = i)$ and $\mathbb{P}_{\text{UG}}(q_P = i + 1)$ approaches $\mathbb{P}_{\text{UG}}(q_P = i)$, $\forall i \in \mathbb{N}^+$, as p goes to 0. Thus, the UG distribution in (12) approximates a uniform distribution in the limit. Since the UG distribution is unimodal with a mode of 1, sorting the nodes in the opposite order makes the probability that a newly-added node is selected as q_P^{closest} slightly larger compared to nodes that were added before, accelerating the expansion of the tree. The reason we prefer UG distribution to the uniform distribution is that the latter is ill-defined over an infinite set, which occurs when the node size goes to infinity. Note that mathematically, the UG distribution is defined over a countably infinite set. To sample from a finite node set of the tree with size n , we first compute the cumulative probability function of the event $\{i \leq n\}$ assuming the node set is infinite, then we renormalize, so that the probability that a node is selected as q_P^{closest} is proportional to that when the node set is infinite.

2) *Selection of $q_B^{\text{succ},1}, q_B^{\text{succ},2} \in \mathcal{Q}_B$* : Next, given the node $q_P^{\text{closest}} = (\mathbf{x}^{\text{closest}}, q_B^{\text{closest}})$, we select two other successive Büchi states $q_B^{\text{succ},1}$ and $q_B^{\text{succ},2}$ in the pruned NBA, such that $q_B^{\text{closest}} \xrightarrow{L(\mathbf{x}^{\text{closest}})}_B q_B^{\text{succ},1} \rightarrow_B q_B^{\text{succ},2}$ and this two-hop transition proceeds towards the accepting state $q_B^{F, \text{feas}}$. That is, $q_B^{\text{succ},1}$ is closer to $q_B^{F, \text{feas}}$ than q_B^{closest} and $q_B^{\text{succ},2}$ is closer than $q_B^{\text{succ},1}$ in terms of the distance function ρ . If the new state q_P^{new} is added to the tree through the selected node q_P^{closest} , the first hop implies that the Büchi state $q_B^{\text{succ},1}$ in q_P^{new} can be q_B^{new} no matter what \mathbf{x}^{new} is. The resulting q_P^{new} is more likely to be selected as q_P^{closest} in the next iterations. If so, the selection of \mathbf{x}^{new} should enable the second hop, i.e., $q_B^{\text{succ},1} \xrightarrow{L(\mathbf{x}^{\text{new}})}_B q_B^{\text{succ},2}$, so that another product state including $q_B^{\text{succ},2}$ can be added to the tree, growing the tree further towards the accepting state $q_B^{F, \text{feas}}$. We refer the reader to [31] for more details.

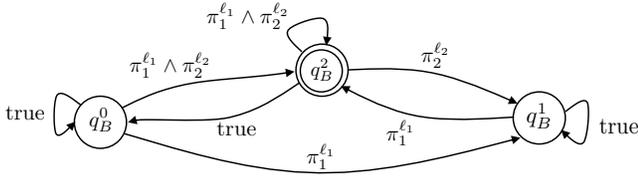


Fig. 3. The NBA for the specification $\square\Diamond\pi_1^{\ell_1} \wedge \square\Diamond\pi_2^{\ell_2}$.

Example 1. Consider the specification $\square\Diamond\pi_1^{\ell_1} \wedge \square\Diamond\pi_2^{\ell_2}$, whose NBA is shown in Fig. 3. When finding the prefix plan, we have $q_B^2 = q_B^{F,feas}$ and the distance $\rho(q_B^0, q_B^2) = 1, \rho(q_B^1, q_B^2) = 1, \rho(q_B^2, q_B^2) = 0$. Initially, the tree \mathcal{T} only contains the root $q_P^0 = (\mathbf{x}_0, q_B^0)$. Assume also that the atomic propositions $\pi_1^{\ell_1}, \pi_2^{\ell_2} \notin L(\mathbf{x}_0)$. According to steps 1) and 2), a possible case of selected vertices is $q_B^{closest} = q_B^0$ (because the tree only has the node (\mathbf{x}^0, q_B^0)), $q_B^{succ,1} = q_B^0$ (because $q_B^0 \xrightarrow{true} q_B^0$ and $q_B^0 \xrightarrow{L(\mathbf{x}_0)} q_B^1, q_B^2$) and $q_B^{succ,2} = q_B^2$ (because $q_B^0 \rightarrow q_B^2$ and $\rho(q_B^0, q_B^2) < \rho(q_B^1, q_B^2)$). That is, when the NBA is at q_B^0 , it should more frequently follow the shortest path q_B^0, q_B^2 . ■

3) *Sampling of $\mathbf{x}^{rand} \in \mathcal{W}$:* Given states $q_B^{succ,1}$ and $q_B^{succ,2}$, we select a propositional formula σ^* such that $q_B^{succ,1} \xrightarrow{\sigma^*} q_B^{succ,2}$ in the NBA [line 4, Alg. 6]; see Fig. 4. For this, we convert the condition enabling the transition from $q_B^{succ,1}$ to $q_B^{succ,2}$ in the disjunctive normal form of $\bigvee \bigwedge \pi_i^{\ell_j}$, that is, $q_B^{succ,1} \xrightarrow{\bigvee \bigwedge \pi_i^{\ell_j}} q_B^{succ,2}$, and we select one clause $\bigwedge \pi_i^{\ell_j}$ randomly or the one with the minimum length to be σ^* . Next, we construct a set \mathcal{L} whose i -th element $\mathcal{L}(i)$ denotes the labeled region that robot i should visit according to the symbol $\pi_i^{\ell_j}$ in σ^* . Note that $\mathcal{L}(i)$ may not appear in σ^* . If $\mathcal{L}(i)$ is an empty symbol, denoted by ϵ^* , then robot i can be anywhere. In this case, with probability p_{idle} , we allow robot i to stay at its current location to incur zero cost; otherwise, robot i moves to any other reachable location. If $\pi_i^{\ell_j}$ appears in σ^* , then with high probability we draw one sample from \mathcal{W}_{free} that is closer to the location $\mathbf{x}_i^{\mathcal{L}(i)}$ than $\mathbf{x}_i^{closest}$ is, where $\mathbf{x}_i^{\mathcal{L}(i)}$ is a point inside region $\mathcal{L}(i) = \ell_j$, such as its centroid. For polygonal environments we employ the geodesic path defined as the shortest path between two points, that is, a sequence of line segments that connect two points and pass through reflex vertices of the polygonal boundary [50]. In order for robot i to reach $\mathbf{x}_i^{\mathcal{L}(i)}$ fast, it should head towards the second vertex in this shortest path, denoted by $\mathbf{x}_i^{target} = SP_{\mathbf{x}_i^{closest}, \mathbf{x}_i^{\mathcal{L}(i)}}(2)$, where $SP_{\mathbf{x}_i^{closest}, \mathbf{x}_i^{\mathcal{L}(i)}}$ denotes the shortest geodesic path from $\mathbf{x}_i^{closest}$ to $\mathbf{x}_i^{\mathcal{L}(i)}$; see Fig. 4. Given \mathbf{x}_i^{target} , we select \mathbf{x}_i^{rand} as

$$\mathbf{x}_i^{rand} = \mathbb{1}_{[Y \leq y_{rand}]} \mathbf{x}_i^{rand,1} + \mathbb{1}_{[Y > y_{rand}]} \mathbf{x}_i^{rand,2}, \quad (14)$$

where $Y \in [0, 1]$ is a random variable drawn from a uniform distribution, $y_{rand} \in (0.5, 1)$ is a weighting factor, $\mathbb{1}_{[Y \leq y_{rand}]}$ is an indicator variable which is 1 if the event $\{Y \leq y_{rand}\}$ occurs, otherwise 0. Also in (14), $\mathbf{x}_i^{rand,1}$ is a point following a normal distribution centered at \mathbf{x}_i^{target} and $\mathbf{x}_i^{rand,2}$ is a point following a uniform distribution that is bounded away from zero on \mathcal{W}_{free} . The fact that y_{rand} is greater than 0.5 ensures that \mathbf{x}_i^{rand} is closer to \mathbf{x}_i^{target} , i.e., robot i moves to \mathbf{x}_i^{target} with high probability. Specifically, the relative position of $\mathbf{x}_i^{rand,1}$ with respect to $\mathbf{x}_i^{closest}$ can be determined by two parameters

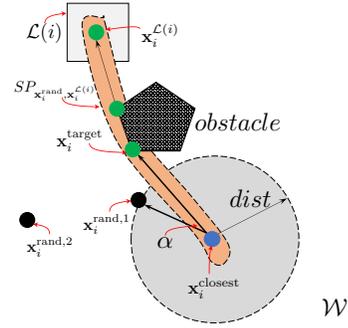


Fig. 4. Graphical depiction of the proposed biased sampling for robot i . Given $q_B^{succ,1}$ and $q_B^{succ,2}$, the symbol σ^* that enables this transition is used to sample \mathbf{x}^{rand} . $\mathcal{L}(i)$ is the labeled region where robot i should be to satisfy σ^* . The path surrounded by light orange area is the shortest path from $\mathbf{x}_i^{closest}$ to $\mathbf{x}_i^{\mathcal{L}(i)}$ generated by visibility graph. $\mathbf{x}_i^{rand,1}$ is the point generated by a normal distribution, with distance d from $\mathbf{x}_i^{closest}$ and angle α between two thick arrow lines. $\mathbf{x}_i^{rand,2}$ is the point generated by a uniform distribution.

in the 2D workspace.⁵ One is the distance $dist$ between $\mathbf{x}_i^{rand,1}$ and $\mathbf{x}_i^{closest}$, and the other one is the angle α formed by two line segments connecting $\mathbf{x}_i^{closest}$ with $\mathbf{x}_i^{rand,1}$ and \mathbf{x}_i^{target} , respectively; see also Fig. 4. We use a 2-dimensional normal distribution to sample $dist$ and α , with mean μ_d, μ_α and standard deviation σ_d, σ_α , i.e.,

$$f(dist, \alpha) = \frac{1}{2\pi\sigma_{dist}\sigma_\alpha} \exp\left(-\frac{1}{2}\left[\frac{dist^2}{\sigma_{dist}^2} + \frac{\alpha^2}{\sigma_\alpha^2}\right]\right).$$

Since the distance is non-negative, we use the absolute value $|dist|$ and α to obtain $\mathbf{x}_i^{rand,1}$. To obtain $\mathbf{x}_i^{rand,2}$, we draw a uniform sample from \mathcal{W}_{free} . If the dimension of the workspace is 3, we need three parameters, including one distance parameter and two angle parameters. After obtaining \mathbf{x}_i^{rand} by (14), we construct $\mathbf{x}^{rand} = [\mathbf{x}_1^{rand,T}, \dots, \mathbf{x}_N^{rand,T}]^T$ [line 5, Alg. 6].

Finally, given \mathbf{x}^{rand} and $q_P^{closest}$ returned by Alg. 6, we generate \mathbf{x}^{new} using the `Steer` function that returns a location that is closer to \mathbf{x}^{rand} than $\mathbf{x}^{closest}$ is. Then we add $q_P^{closest}$ to set \mathcal{Q}_P^{near} [lines 3-4, Alg. 5]. Next, we construct states q_P^{new} by pairing a Büchi state q_B to \mathbf{x}^{new} , exactly as in Alg. 2.

B. Construction of Suffix Plans

The algorithm to design the suffix part differs from the one proposed for the biased prefix part in Section V-A only in that a cycle around the root needs to be computed. Specifically, once a node q_P^{new} is constructed, we check whether its Büchi component is the same as that of the root, i.e., whether it is the same accepting Büchi state; if so, we store this node in a set \mathcal{P} , otherwise the tree is built exactly as for the biased prefix part. Together with the construction of the tree, for each node $q_P^{new} = (\mathbf{x}^{new}, q_B^{new})$ in \mathcal{P} , we find a path from $\mathbf{x}^{new} = [\mathbf{x}_1^{new,T}, \dots, \mathbf{x}_N^{new,T}]^T$ to \mathbf{x}^0 that forms a cycle. To find this path, we apply RRT* [6] to find a path for each robot $i \in [N]$, that connects \mathbf{x}_i^{new} to \mathbf{x}_i^0 while treating all labeled regions as obstacles. This ensures that during the execution of the plan no other observation will be generated and, therefore, the robots will maintain the desired accepting Büchi state.

Remark 5.1: While in this paper we focus on the multirobot planning problem, the proposed method, similar to RRT*, can address optimal control problems subject to temporal

⁵If the dimension of the workspace is 3, we need three parameters, including one distance parameter and two angle parameters.

logic specifications too, such as manipulation tasks [51] or estimation tasks [52]. For this, it is important that infeasible transitions in the NBA can be identified and pruned since, otherwise, it is possible that TL-RRT* can bias search towards those infeasible transitions, which can significantly deteriorate performance. In motion planning applications, it is simple to identify such infeasible transitions as discussed in Section V; see also [30], [31]. However, in general, identifying infeasible NBA transitions may be problem specific and not easy to do.

VI. CORRECTNESS AND OPTIMALITY

In this section we show that TL-RRT* is probabilistically complete and asymptotically optimal. Note that TL-RRT* does not trivially inherit the completeness and optimality properties of RRT*, since TL-RRT* explores a combined continuous and discrete state space while RRT* is designed to explore only continuous state spaces. The resulting technical differences with RRT* are discussed in the proof of TL-RRT* in Appendices. First, we make the following assumptions.

Assumption 6.1: (Nonpoint regions) Every atomic proposition in the LTL formula ϕ is satisfied over a nonpoint region. More precisely, $\mu(\ell_j) > 0$ where μ is the Lebesgue measure.

Assumption 6.1 ensures that any point within a labeled region can be sampled with nonzero probability. Otherwise, it is impossible to generate a feasible plan. In what follows, we denote by $\mathcal{B}_r((\mathbf{x}, q_B))$ a ball of radius of r centered at (\mathbf{x}, q_B) in the product state space,

$\mathcal{B}_r((\mathbf{x}, q_B)) = \{(\mathbf{x}', q'_B) \in \mathcal{Q}_P \mid \text{dist}((\mathbf{x}, q_B), (\mathbf{x}', q'_B)) \leq r\}$, where, with slight notational abuse, $\text{dist}((\mathbf{x}, q_B), (\mathbf{x}', q'_B)) = \|\mathbf{x} - \mathbf{x}'\|$ is the distance between two product states. In words, a product state $(\mathbf{x}', q'_B) \in \mathcal{Q}_P$ lies in the ball $\mathcal{B}_r((\mathbf{x}, q_B))$ if \mathbf{x}' is at distance less than r from \mathbf{x} . We denote by $\text{int}(\mathcal{B}_r((\mathbf{x}, q_B)))$ the interior of the ball $\mathcal{B}_r((\mathbf{x}, q_B))$. By definition of the distance function dist , a point \mathbf{x}' lies in the ball $\mathcal{B}_r((\mathbf{x}, \cdot))$ regardless of its Büchi state component. This definition of a ball is necessary since the product space consists of the continuous space $\mathcal{W}_{\text{free}}^N$ and the discrete space \mathcal{Q}_B , and there is not a physical notion of distance in \mathcal{Q}_B .

Assumption 6.2: (Convergence space) Let Assumption 6.1 hold. For any reachable product state $(\mathbf{x}, q_B) \in \mathcal{W}_{\text{free}}^N \times \mathcal{Q}_B$ from the root, there exists a constant $\delta_{\mathbf{x}} > 0$ that depends on \mathbf{x} , such that any point \mathbf{x}' in $\mathcal{W}_{\text{free}}^N$ for which (\mathbf{x}', \cdot) lies in the interior of the ball $\mathcal{B}_{\delta_{\mathbf{x}}}((\mathbf{x}, q_B))$, can be paired with the same Büchi state q_B as the center \mathbf{x} . Therefore, the product state (\mathbf{x}', q_B) can also be reached from the root.

Assumption 6.2 ensures that a homotopy class exists around any feasible path such that any path in this class can be continuously deformed into another.

A. Probabilistic Completeness

In this section, we show the probabilistic completeness of TL-RRT* by skipping `Extend` and `Rewire`, and instead connecting q_P^{new} only with nodes in $\mathcal{Q}_P^{\text{nearest}}$, similar to RRT [4]. The reason is that, if there exists a candidate parent q_P^{nearest} of q_P^{new} in $\mathcal{Q}_P^{\text{nearest}}$, then q_P^{new} will be added to the tree \mathcal{T} regardless of which node in $\mathcal{V}_{\mathcal{T}}$ is selected by `Extend` to be its parent. Furthermore, `Rewire` updates the set of edges of the tree and does not play any role in adding new nodes. Therefore,

it can not affect the completeness property. Thus, we focus on finding a candidate parent in $\mathcal{Q}_P^{\text{nearest}}$. If connecting q_P^{new} with nodes in $\mathcal{Q}_P^{\text{nearest}}$ can ensure probabilistic completeness, based on the fact that $\mathcal{Q}_P^{\text{nearest}}$ is a subset of $\mathcal{Q}_P^{\text{near}}$ among which TL-RRT* attempts to find the parent, the probabilistic completeness of TL-RRT* follows directly. Note that when the number of nodes goes to infinity, the distance between the sampled point and the nearest nodes is much smaller than η . Therefore, \mathbf{x}^{new} will be identical to \mathbf{x}^{rand} . This argument can be found in Lemma 1 in [47]. Our proofs are based upon the observation that $\mathbf{x}^{\text{new}} = \mathbf{x}^{\text{rand}}$. The following theorem shows that unbiased TL-RRT* is probabilistically complete.

Theorem 6.3: (Probabilistic Completeness of TL-RRT with Unbiased Sampling)* Let Assumptions 6.1 and 6.2 hold and further assume that sampling in the free workspace is unbiased. Then, TL-RRT* is probabilistically complete, i.e., if there exists a feasible plan that satisfies a given LTL formula ϕ , then TL-RRT* will find it with probability 1.

Proof: We provide a proof sketch and the details can be found in Appendix A. The proof proceeds in two steps. First, given any product state in \mathcal{Q}_P that is one-hop-reachable from the root, we prove that the tree will have a node that is arbitrarily close to it in terms of Euclidean distance, and based on Assumption 6.2, this node will have the same Büchi state as the given product state. The key idea is to show that the expected distance between the tree and the given product state converges to 0 as the tree grows. Next, we extend the above result to multi-hop reachable states from the root. Therefore, given an accepting state that is reachable from the root, a node that is arbitrarily close to it and has the same accepting Büchi state will be added to the tree with probability 1 as the number of iterations of Alg 1 goes to infinity. ■

Theorem 6.3 can also be extended for the case of the biased sampling introduced in Section V. Specifically, we have the following result, whose proof can be found in Appendix B.

Corollary 6.4: (Probabilistic Completeness of TL-RRT with Biased Sampling)* Let Assumptions 6.1 and 6.2 hold. Then, the biased TL-RRT* is probabilistically complete, i.e., if there exists a feasible plan that satisfies a given LTL formula ϕ , then the biased TL-RRT* will find it with probability 1.

B. Asymptotic Optimality

In this section, we show the asymptotic optimality of TL-RRT*. We first define a product plan p given a discrete plan $\tau = \tau(1), \dots, \tau(k), \dots$ satisfying ϕ . Taking $\text{trace}(\tau)$ as the input to the NBA, a run $q_B^1, \dots, q_B^k, \dots$ will be produced. Given the one-to-one correspondence between states in τ and states in this run, we can construct a product state plan p by pairing each position component with a Büchi state, i.e., $p = (\tau(1), q_B^1), \dots, (\tau(2), q_B^2), \dots$. In this case, τ is the projection of p onto $\mathcal{W}_{\text{free}}^N$. Moreover, let $\tau^* = \tau^{*,\text{pre}}[\tau^{*,\text{suf}}]^\omega$ be the optimal plan that satisfies ϕ and incurs the optimal cost defined in (3). We use $\tau^{*,\text{pre}}|\tau^{*,\text{suf}}$ to represent τ^* since it suffices to characterize the optimal plan.

Theorem 6.5: (Asymptotic Optimality of TL-RRT with Unbiased Sampling)* Let Assumptions 6.1 and 6.2 hold and further assume that sampling in the free workspace is unbiased.

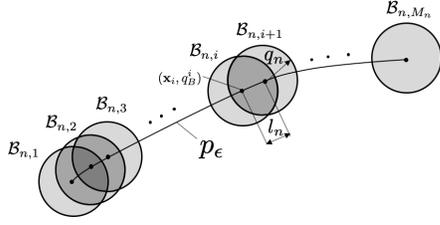


Fig. 5. An illustration of the construction of balls covering the whole product path p_ϵ , modified from [6]. All balls have the same radius q_n . The spacing between centers of two consecutive balls is l_n ; see Appendix C-2 for details.

Consider also the parameter $r_n(\mathcal{V}_\tau)$ defined in (7). Then, TL-RRT* is asymptotically optimal, i.e., the discrete plan $\tau_{n_{\max}^{\text{pre}}, n_{\max}^{\text{suf}}}$ that is generated by this algorithm satisfies

$$\lim_{n_{\max}^{\text{pre}} \rightarrow \infty, n_{\max}^{\text{suf}} \rightarrow \infty} \mathbb{P} \left(\left\{ J(\tau_{n_{\max}^{\text{pre}}, n_{\max}^{\text{suf}}}) \leq (1 + \epsilon) J(\tau^*) \right\} \right) = 1. \quad (15)$$

where $\epsilon \in (0, 1)$, n_{\max}^{pre} and n_{\max}^{suf} are the maximum numbers of iterations used in Alg. 2, $\tau_{n_{\max}^{\text{pre}}, n_{\max}^{\text{suf}}} = \tau^{\text{pre}, n_{\max}^{\text{pre}} | \tau^{\text{suf}, n_{\max}^{\text{suf}}}$, $\tau^* = \tau^* \cdot \text{pre} | \tau^* \cdot \text{suf}$ and J is the cost function defined in (3).

Proof: We provide a proof sketch based on [47] and the details can be found in Appendix C. For any $\delta \in \mathbb{R}^+$, the δ -interior of the space $\mathcal{W}_{\text{free}}^N$, denoted by $\text{int}_\delta(\mathcal{W}_{\text{free}}^N)$, is a subset of $\mathcal{W}_{\text{free}}^N$ containing states that are at least δ distance away from any obstacle. Then, we call a feasible path τ *robust* if there exists $\delta \in \mathbb{R}^+$ such that $\tau \subset \text{int}_\delta(\mathcal{W}_{\text{free}}^N)$. By Assumption 6.2, Problem 1 is robustly feasible. Following [47], there exists a path τ_ϵ such that $J(\tau_\epsilon) \leq (1 + \epsilon/4)J(\tau^*)$, where $\epsilon \in (0, 1)$. We show that the cost of the plan returned by TL-RRT* is at most $(1 + \epsilon)J(\tau^*)$ when the number of iterations goes to infinity. To this end, we augment the path τ_ϵ using Büchi states to obtain a product path p_ϵ . Then, at each iteration, we construct a sequence of equally-spaced balls of same radii centered along the path p_ϵ so that the sequence of balls covers it; see Fig. 5. Furthermore, the radii of these balls are proportional to the connection radius $r_n(\mathcal{V}_\tau)$, which converges to 0 as the iteration n grows, so the number of balls grows to infinity. Consider the event that each ball contains a sample and, for any two adjacent balls, the sample inside the second ball is sampled after that inside the first ball. Furthermore, assume that if a ball is intersected by a boundary of any region, this boundary can be locally approximated by a hyperplane when the ball is extremely small. Then, we require that samples should lie within those parts of the balls that contain their centers. This is because the transition relation between two points requires that the straight line connecting them crosses any boundary at most once. When this event occurs, TL-RRT* will connect together this sequence of samples, augmented by Büchi states, in an ascending order. In this way, the tree contains a path that arbitrarily approximates the path p_ϵ . We prove that, when the connection radius $r_n(\mathcal{V}_\tau)$ satisfies (7), the probability of the above event converges to 1. Finally, the approximation of p_ϵ induces a path whose cost is no more than $(1 + \epsilon)J(\tau^*)$. ■

The next result shows the optimality of the biased TL-RRT* algorithm. The proof can be found in Appendix D.

Corollary 6.6: (Asymptotic Optimality of TL-RRT with Biased Sampling)* Let Assumptions 6.1 and 6.2 hold. Assume

also that the parameter $\gamma_{\text{TL-RRT}^*}$ in (7) satisfies

$$\gamma_{\text{TL-RRT}^*} \geq (2 + \theta) \left(\frac{(1 + \epsilon/4)J(\tau^*)\mu(\mathcal{W}_{\text{free}}^N)}{(1 - y_{\text{rand}})^N (\text{dim} + 1)\theta(1 - \kappa)\zeta_{\text{dim}}} \right)^{\frac{1}{\text{dim} + 1}},$$

where $(1 - y_{\text{rand}})^N$ is the probability that \mathbf{x}^{rand} is obtained following the uniform distribution at each iteration. Then, the biased TL-RRT* is asymptotically optimal in the sense of (15).

VII. SIMULATION RESULTS

In this section, we present three case studies, implemented using Python 3.6.3 on a computer with 2.3 GHz Intel Core i5 and 8G RAM, that illustrate the efficiency and scalability of the proposed algorithm. We first show the correctness and optimality of the unbiased TL-RRT*. Second, we compare our sampling-based TL-RRT* algorithm, with and without bias, with the synergistic method in [43], the RRG method in [42] and the SMC method in [38], with respect to the size of regions. Finally we test the scalability of biased TL-RRT* with respect to the complexity of tasks. It shows that biased TL-RRT* outperforms the synergistic method, the RRG method and the SMC method in terms of optimality and scalability. The implementation is accessible from https://github.com/XushengLuo/TLRRT_star.

In all the following case studies, the LTL formula takes the general form of $\phi = \phi_{\text{spec}} \wedge \phi_{\text{col}}$, where ϕ_{spec} is the specified task and ϕ_{col} means collision avoidance among robots. Specifically, ϕ_{col} requires that, at the same timestamp and in each dimension, the distance between any two robots is larger than R [34], [38]. We assume that robots are equipped with motion primitives that allow them to move safely between consecutive waypoints in the discrete plan, as in [20]. This can be done during the execution phase of the plans using the methods proposed, e.g., in [53], [54]. Specifically, given the possibly intersecting paths returned by TL-RRT*, the methods in [53], [54] can design policies that stop and resume robots to avoid collisions between them.

We consider planning problems for robots that lie in a 1×1 workspace with $W = 6$ isosceles right triangular regions of interest with side length s , and two rectangular obstacles; see Fig. 6. The parameters are set as follows: the step-size η of the function `Steer` is $0.25N$, where N is the number of robots, and w in the cost function (3) is 0.2. In the biased sampling method, $p_{\text{idle}} = 1$, $p_{\text{closest}} = 0.9$, $y_{\text{rand}} = 0.99$, $\mu_d = \mu_\alpha = 0$, $\sigma_d = 1/3$, and $\sigma_\alpha = \pi/108$. The safe distance is $R = 0.005$. Considering that the obstacles are polygonal, the geodesic paths can be constructed using the visibility graph [55].

A. Correctness and optimality using unbiased sampling

In the first simulation, we test the correctness and optimality of the proposed algorithm with unbiased sampling. The side length s is 0.15. We consider a single robot that is initially located at $\mathbf{x}^0 = (0.8, 0.1)$. The assigned task is:

$$\phi_{\text{spec}} = \diamond(\pi_1^{\ell_1} \wedge \diamond\pi_1^{\ell_3}) \wedge (\neg\pi_1^{\ell_1} \mathcal{U}\pi_1^{\ell_2}) \wedge \diamond(\pi_1^{\ell_5} \wedge \diamond(\pi_1^{\ell_6} \wedge \diamond\pi_1^{\ell_4})) \wedge (\neg\pi_1^{\ell_4} \mathcal{U}\pi_1^{\ell_5}) \wedge \square\neg\mathcal{O}. \quad (16)$$

In words, ϕ_{spec} requires robot 1: (a) to eventually visit region ℓ_1 and then ℓ_3 , (b) not visit region ℓ_1 before visiting ℓ_2 , (c) to eventually visit region ℓ_5 and next ℓ_6 and then ℓ_4 , (d) not visit region ℓ_4 before visiting ℓ_5 and (e) avoid obstacles. The LTL formula ϕ_{spec} corresponds to an NBA with $|\mathcal{Q}_B| = 28$ states, $|\mathcal{Q}_B^0| = 1$, and $|\mathcal{Q}_B^F| = 2$. which was constructed using the

TABLE I
PERFORMANCE VERSUS THE NUMBER OF ITERATIONS

n_{\max}^{pre}	$T(\text{s})$	$J(\tau)$	$ \mathcal{P} $
581.9±274.1 412.4±186.0	47.0±37.0 38.7±34.4	0.619±0.046 0.611±0.050	1±0 1±0
600	47.7±7.6 (6) 115.2±27.3 (2)	0.572±0.034 0.543±0.027	3.9±2.6 4.5±2.0
800	95.5±10.9 (6) 282.6±68.7	0.541±0.023 0.514±0.020	5.8±3.1 8.6±3.5
1000	163.2±30.7 580.6±101.6	0.525±0.030 0.507±0.020	7.9±3.6 10.8±3.0

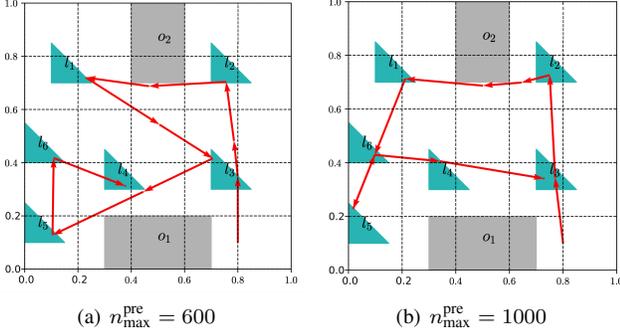


Fig. 6. Simulation Study I: Sample paths.

tool in [56]. The term $\diamond(\pi_1^{\ell_1} \wedge \diamond\pi_1^{\ell_3})$ can capture sequential surveillance and data gathering tasks by visiting specific regions of interest, the term $-\pi_1^{\ell_1} \mathcal{U}\pi_1^{\ell_2}$ can assign certain priority or designate the order between different subtasks.

To determine the connection radius $r(\mathcal{V}_{\mathcal{T}})$ in (7), we set $\gamma_{\text{TL-RRT}^*}$ equal to the lower bound in (8). However, note that the connection radius in (7) can not be computed because it requires knowledge of the optimal plan that is needed in (8). Therefore, we also consider the alternative connection radius

$$r_n(\mathcal{V}_{\mathcal{T}}) = \min \left\{ \gamma_{\text{TL-RRT}^*} \left(\frac{\log |\mathcal{V}_{\mathcal{T}}|}{|\mathcal{V}_{\mathcal{T}}|} \right)^{\frac{1}{\text{dim}}}, \eta \right\}, \quad (17)$$

where $\gamma_{\text{TL-RRT}^*} = \lceil 4 \left(\frac{\mu(\mathcal{W}_{\text{free}}^N)}{\zeta_{\text{dim}}} \right) \rceil$, that was proposed in [6] for RRT* and has been tested extensively in practice. The main difference between (7) and (17) is in the exponent, which is $1/(\text{dim}+1)$ in (7) versus $1/\text{dim}$ in (17). In what follows, to determine the lower bound on $\gamma_{\text{TL-RRT}^*}$ needed to compute the connection radius in (7), we first use the connection radius in (17) to find a feasible path and then use the cost of this path as an approximation of the optimal cost in (8). Moreover, we set $\theta = 0.24$ and $\epsilon = \kappa = 0.01$ that appear in (8). Table I shows statistical results, i.e., mean and standard deviation, on the total runtime, the number $|\mathcal{P}|$ of detected final states of the prefix plan, and the cost $J(\tau)$ of the path for different numbers of iterations n_{\max}^{pre} , averaged over 20 trials. In the first case, the unbiased TL-RRT* was terminated when the first feasible path was detected and in the remaining cases the algorithm was terminated after a fixed number of iterations n_{\max}^{pre} . Note that this task is satisfiable by only the prefix plan. The number inside the parentheses is the number of failed trials when no plan was generated. The first row for each value of n_{\max}^{pre} corresponds to the connection radius in (17), whereas the second row corresponds to the connection radius in (7). The time needed to find a feasible path using the connection radius in (7) is not included. Observe that for the

same connection radius, as the number of iterations increases, the cost of the computed path decreases, as expected due to Theorem 6.5. Observe also in Table I that the connection radius in (7) returns paths with lower cost compared to the connection radius in (17). This is because the connection radius in (7) can discover more accepting states in the set \mathcal{P} and, therefore, more paths, compared to using the one in (17). In our simulations, we observed that the ratio of the connection radius in (7) to that in (17) ranges from 0.8 to 6 as the tree grows. This means that the connection radius in (7) becomes increasingly larger than that in (17) as the tree grows, and a larger connection radius facilitates the discovery of accepting states. Nevertheless, the difference in the cost of the paths returned by TL-RRT* using the two connection radii is not very large, meaning that the connection radius in (17) can still give a good approximation of the optimal path. This is important since (17) is easy to compute as discussed before. Two paths for cases $n_{\max}^{\text{pre}} = 600, 1000$ are depicted in Fig. 6. In practice, we can inflate obstacles and shrink regions to make paths more robust to perturbations, as in [57].

B. Comparison with other methods w.r.t. the size of regions

In this section, by decreasing the size of regions, we compare TL-RRT* to the synergistic (SYN) method [43], the RRG method [42] and the SMC method [38]. Specifically, a team of 2 robots are initially located around (0.8, 0.1) with no collisions between them. The task is:

$$\phi_{\text{spec}} = \square \diamond \pi_1^{\ell_1} \wedge \square \diamond \pi_2^{\ell_2} \wedge \square \diamond (\pi_1^{\ell_4} \wedge \diamond \pi_2^{\ell_4}), \quad (18)$$

which requires: (a) robot 1 to visit region ℓ_1 infinitely often, (b) robot 2 to visit region ℓ_2 infinitely often, (c) robot 1 and 2 to visit region ℓ_4 infinitely often and robot 2 will eventually visit region ℓ_4 after each time robot 1 visits this region. The term $\square \diamond (\pi_1^{\ell_4} \wedge \diamond \pi_2^{\ell_4})$ can capture intermittent connectivity tasks that require robots to reach predetermined communication points infinitely often to exchange information but not necessarily concurrently [8], [9]. The considered LTL formula corresponds to an NBA with $|\mathcal{Q}_B| = 8$ states, $|\mathcal{Q}_B^0| = 1$, and $|\mathcal{Q}_B^F| = 2$. We adopt the connection radius in (17) since it is easy to compute and its use does not significantly affect the cost of the path returned by TL-RRT*, as shown in Section VII-A.

The synergistic planning method in [43] consists of a high-level planner that operates in the product space of N discrete transition systems, corresponding to the N robots, and the Büchi state space, and a low-level sampling-based planner that builds trees in the continuous space guided by the high-level states. In the simulation, we use the same abstraction as that used by the SMC method and new states are sampled uniformly around the selected high-level state.

The RRG method in [42] maintains a sparse approximation of the workspace such that states are “far” away from each other. To this end, it only connects new sampled states to nodes in the graph if their pairwise distance is smaller than $\eta_2(k)$ and larger than $\eta_1(k)$, where k is the number of different position states in the graph. To provide a fair comparison between TL-RRT* and the RRG method, we set the parameter $\eta_1(k)$ in [42] to be $\frac{1}{\sqrt{\pi}} \sqrt{\frac{\mu(\mathcal{W}^N) \Gamma(dN/2+1)}{k}}$ for all $k \geq 1$, where $\mu(\mathcal{W}^N)$ is the total measure (volume) of the configuration space, dN is the dimension of \mathcal{W}^N , Γ is the gamma function, and $\eta_2(k) =$

TABLE II
COMPARISON OF RUNTIMES AND COST FOR THE FIRST FEASIBLE PATH FOR THE LTL FORMULA (18) INVOLVING $N = 2$ ROBOTS

s	runtime					cost				
	$T_{\text{TL-RRT}^*}^{\text{unbias}}(s)$	$T_{\text{TL-RRT}^*}^{\text{bias}}(s)$	$T_{\text{SYN}}(s)$	$T_{\text{RRG}}(s)$	$T_{\text{SMC}}(s)$	$J_{\text{TL-RRT}^*}^{\text{unbias}}$	$J_{\text{TL-RRT}^*}^{\text{bias}}$	J_{SYN}	J_{RRG}	J_{SMC}
0.25	19.0±34.4	0.4±0.3	124.1±41.1	37.9±59.2	15.6±3.9	2.42±0.41	1.75±0.21	3.22±0.46	3.34±0.81	3.19±0.71
0.20	21.4±21.2	0.4±0.1	123.9± 63.0	116.4±133.9	12.9±2.1	2.41±0.32	1.71±0.20	3.40±0.43	3.23±0.73	3.55±0.20
0.15	41.0±28.8	0.6±0.3	446.9±147.1	315.4±288.4	14.6±2.0	2.43±0.42	1.78±0.11	4.46±1.03	3.83±0.78	3.14±0.04
0.10	197.8±140.3	1.0±0.4	378.1±279.6	—	11.6±1.0	2.65±0.51	1.76±0.12	6.00±1.25	—	2.70±0.00

TABLE III
COMPARISON OF RUNTIMES AND COST FOR THE FIRST 5 FEASIBLE PATHS FOR THE LTL FORMULA (18) INVOLVING $N = 2$ ROBOTS

s	$T_{\text{TL-RRT}^*}^{\text{unbias}}(s)$	$T_{\text{TL-RRT}^*}^{\text{bias}}(s)$	$T_{\text{SYN}}(s)$	$T_{\text{RRG}}(s)$	$T_{\text{SMC}}(s)$			$J_{\text{TL-RRT}^*}^{\text{unbias}}$	$J_{\text{TL-RRT}^*}^{\text{bias}}$	J_{SYN}	J_{RRG}	J_{SMC}		
					10	20	*					10	20	*
					0.25	62.3±54.4	1.3±0.3					227.2±54.6	59.0±88.6	17.1±0.7
0.20	78.7±39.4	1.6±0.7	250.3±52.1	167.7±178.1	18.1±0.8	4.6±0.2	4.9±0.4	2.17±0.28	1.56±0.11	2.66±0.18	2.76±0.80	2.98	3.81	2.73
0.15	291.8±280.6	1.8±0.5	575.6±214.9	—	19.5±2.9	5.4±0.9	3.5±0.1	2.18±0.18	1.57±0.07	3.11±0.29	—	3.22	2.99	3.00
0.10	522.3±314.3	3.4±1.5	485.1±189.9	—	19.8±2.2	5.5±1.4	4.5±0.3	2.27±0.27	1.67±0.10	5.34±0.38	—	2.58	3.30	2.70

$2\eta_1(k)$ satisfies (i) $\eta_1(k) < \eta_2(k)$ for all $k \geq 1$, (ii) $c\eta_1(k) > \eta_2(k)$ for some finite $c > 1$ and all $k \geq 0$.

Finally, in our implementation of the SMC method, we use Z3 [58] as the SAT solver and CPLEX as the optimization solver. The workspace is triangulated to create a coarse abstraction and the LTL formula is encoded as Boolean constraints using Bounded Model Checking (BMC) [59], which involves a predetermined parameter specifying the horizon of the plan. However, a feasible initial horizon is not known a priori. Therefore, when the SMC method fails to return a solution, we re-run it after increasing the horizon by 1.

Table II compares the runtimes of the different algorithms until the first feasible plan is discovered and the costs of these plans calculated as in (3), averaged over 20 trials, for various choices of the side length s of the labeled regions, where the symbol “—” means the runtime is larger than 1000s. In the case of the SMC method, the initial horizon is 15 and the runtime of the algorithm includes the time of all failed attempts preceding the first successful attempt. Observe first in Table II that the runtime of the RRG method in [42] is large. This is because, as the graph built by the RRG method grows, the parameter η_1 goes to zero and the constructed graph loses its sparsity. Moreover, as the side length s decreases, it becomes increasingly difficult for the RRG method to sample states that belong to the labeled regions, which further increases the runtime of the algorithm. A similar observation also applies to the synergistic method in [43]. Specifically, since the high-level state space is large, being the product of two discrete transition systems one for each robot, high-level planning is expensive which increases the runtime of the algorithm. Moreover, as with the RRG method, as the side length s decreases, it becomes increasingly difficult also for the synergistic method to sample new states that belong to the labeled regions, which further increases runtime. On the other hand, TL-RRT*, with or without bias, grows trees so that sparsity of the graph is not an issue as with the RRG method. Additionally, biased TL-RRT* guides the sampling process based only on the Büchi states and not on the number of robots or the size of a discrete abstraction of the environment as the synergistic method in [43]. As a result, biased TL-RRT*

is much faster compared to both the RRG method and the synergistic method for any side length s of the labeled regions. Note that, while the runtime of biased TL-RRT* also increases as the side length s decreases, this increase is small since sampling is strongly biased towards labeled regions no matter how small they are. Unlike the other methods, the runtime of the SMC method does not change much with the side length s of the labeled regions since, in this example, the partition of the workspace does not change much as the shapes and locations of the triangular regions remain mostly unchanged. Even so, biased TL-RRT* also outperforms the SMC method in terms of runtime in all cases. As for the cost $J(\tau)$, the plan found by the biased TL-RRT* outperforms all other methods. This is because biasing sampling along shortest paths in the Büchi automaton has the effect that detours in the workspace that increase cost are avoided.

Table III shows comparative results until 5 feasible paths are found using the sampling-based methods. For the SMC method, results correspond to the runtime and cost of the first feasible plan for initial horizons equal to 10 and 20, where the 0 standard deviation is dropped for simplicity. The last column in the SMC case, denoted by “*”, shows the results with “perfect” initial horizons. A “perfect” initial horizon is the smallest horizon needed to obtain a feasible plan. Here, the first successful trial was obtained for a horizon 17 for all side lengths of labeled regions. The results shown in Table III, further validate those discussed above in Table II. For all sampling-based methods, a noticeable reduction in cost can be obtained at the expense of a growing runtime. As for the SMC method, if initialized with a smaller horizon, such as 10, it suffers more failures until a feasible plan is found, thus increasing its runtime. Since the first feasible plan is found for horizon 17, the runtime for finding a solution with a larger initial horizon 20 varies slightly. Noticeably, biased TL-RRT* can reduce cost by slightly increasing runtime. Also, here too, it significantly outperforms all other methods.

C. Scalability w.r.t the complexity of tasks and the number of robots using biased sampling

Below we demonstrate the scalability of the biased TL-RRT* at the expense of optimality. As in Section VII-A,

we use two connection radii in the functions `Extend` and `Rewire`. The first connection radius is determined by (17), and the second one is 0, which amounts to connecting q_P^{new} with q_P^{closest} directly. Note that the connection radius 0 suffices to find feasible plans and, thus, can fully demonstrate the scalability of the biased TL-RRT*. In both cases, we set the step-size η in the function `Steer` to be large enough so that \mathbf{x}^{new} can be directly reached from $\mathbf{x}^{\text{closest}}$.

We consider a team of up to $N = 56$ robots accomplishing a set of tasks. We adopt the representation in [30] where the formula ϕ_{spec} is written in the following compact form:

$$\begin{aligned} \phi_{\text{spec}} = & \square\Diamond\xi_1 \wedge \square\Diamond\xi_2 \wedge \square\Diamond\xi_3 \wedge \square\Diamond(\xi_4 \wedge \Diamond(\xi_5 \wedge \Diamond\xi_6)) \\ & \wedge \Diamond\xi_7 \wedge \square\Diamond\xi_8 \wedge (!\xi_7 \mathcal{U} \xi_8). \end{aligned} \quad (19)$$

The LTL formula is satisfied if (i) ξ_1 is true infinitely often; (ii) ξ_2 is true infinitely often; (iii) ξ_3 is true infinitely often; (iv) ξ_4 , ξ_5 and ξ_6 are true in this order infinitely often; (v) ξ_7 is true eventually; (vi) ξ_8 is true infinitely often; (vii) ξ_7 is false until ξ_8 becomes true for the first time. The subformula ξ_e takes the Boolean form of $\bigwedge_{i=1}^m \pi_i^{\ell_j}$ that involves a subteam of robots. For instance, when $m = 3$, ξ_1 can be $\xi_1 = \pi_1^{\ell_3} \wedge \pi_3^{\ell_5} \wedge \pi_5^{\ell_1}$, which is true if (i) robot 1 is in region ℓ_3 ; (ii) robot 3 is in region ℓ_5 and (iii) robot 5 is in region ℓ_1 . All other Boolean formulas ξ_e are defined similarly. The corresponding NBA has 33 states and 348 edges. Given a team of robots, we randomly divide it into overlapping robot subteams in a way such that each robot belongs to at least one subteam. Then, we associate each subteam of robots with one subformula ξ_e .⁶

Along with pruning edges from the NBA as discussed in Section V, we further delete those feasible edges with labels in the form $\wedge_e \xi_e$, provided that $\wedge_e \xi_e$ does not contain the negation of any subformula ξ_e , that require more than one subformulas to be true simultaneously. The reason is that as m grows, each subformula will include more robots, thus, it will become harder to satisfy multiple subformulas simultaneously. After deletion, the resulting NBA has 111 edges, a dramatic drop compared to the original size. Note that the problem is still feasible since edges labeled with a single formula are intact and they contain a solution.⁷ Given a team of N robots, 5 different tasks ϕ_{spec} are generated randomly. It takes on average 2 seconds to prune the NBA. For each task, 5 sets of initial locations are randomly generated from $\mathcal{W}_{\text{free}} \setminus \cup_{i=1}^6 \ell_i$ with no collisions between robots. For each set of initial locations, we run the biased TL-RRT* 5 times, each terminating when a feasible plan is found, and compare with the SMC method in [38]. We also tested the synergistic method in [43], which failed to generate a plan within 1000 seconds, which agrees with the results in [38]. For each N , we record the runtime and cost of the SMC

⁶This formulation amounts to a generator of tasks rather than a specific task instance. It provides a systematic approach to testing the scalability by increasing the number of robots. E.g., an intermittent communication task [9] can be generated as $\square\Diamond(\pi_1^{\ell_1} \wedge \pi_3^{\ell_1}) \wedge \square\Diamond(\pi_2^{\ell_2} \wedge \pi_3^{\ell_2}) \dots$

⁷We can safely and automatically delete those edges since there is no conjunction of subformulas $\wedge_e \xi_e$ in the LTL formula (19). If there is a conjunction in (19), e.g., $\xi_1 \wedge \xi_2$, we can define an additional subformula $\xi_+ = \xi_1 \wedge \xi_2$ to replace the conjunction. We conducted simulations with and without such edge deletions, and observed that when these edges are deleted from the NBA, feasible plans can be found faster due to the smaller size of the pruned NBA and the edge labels that can be more easily satisfied.

TABLE IV
RUNTIMES AND COST FOR TASKS WITH INCREMENTAL COMPLEXITY

Task	TL-RRT*		SMC-based		
	$T_{\text{total}}(\text{s})$	$J(\tau)$	$T_{\text{total}}(\text{s})$	$J(\tau)$	$T_{\text{total}}^{(1)}(\text{s})$
ϕ_8	5.7±3.2 3.3±1.5	2.42±0.74 2.99±0.79	8.4±2.9	3.30±0.65	12.07
ϕ_{16}	56.4±50.2 17.5±18.1	6.81±1.53 7.73±1.04	89.8±16.1	8.27±0.67	131.66
ϕ_{24}	133.2±132.9 17.4±10.6	6.85±0.89 8.75±0.78	170.5±8.9	9.93±0.92	251.43
ϕ_{32}	288.6±184.4 75.8±85.0	9.73±0.90 13.80±1.35	321.3±70.5	11.81±1.54	470.07
ϕ_{40}	601.1±326.4 97.0±52.8	11.51±1.22 13.45±1.43	1025.6±529.0	14.16±1.32	1599.50
ϕ_{48}	1278.5±567.5 148.4±82.0	14.61±1.61 15.91±1.18	960.84±188.7	17.19±1.76	1380.63
ϕ_{56}	2245.3±419.7 374.1±491.4	16.50±1.81 16.69±1.68	1329.53±354.8	17.53±3.07	1632.21

Simulation Study III: ϕ_N denotes a task that involves N robots. For biased TL-RRT*, T_{total} is the total runtime needed to prune the NBA and find the first feasible prefix and suffix plan. For SMC, T_{total} represents the total runtime needed by the SAT solver and the CPLEX solver, with “perfect” initial horizons. $T_{\text{total}}^{(1)}$ is the average total runtime needed if the initial horizon is 1 step shorter than the smallest horizon that provides a feasible plan.

method, averaged over 25 trials, if starting with the “perfect” initial horizon as well as the average runtime if the initial horizon is 1 step shorter than the “perfect” initial horizon. The results are averaged over 125 experiments and are reported in Table IV. For TL-RRT*, the first row in each block shows results for the connection radius (17) and the second row for the connection radius 0. Observe that, with connection radius in (17), biased TL-RRT* achieves the lowest cost for the first feasible plans. We found that the connection radius in (17) decreases slowly, and, therefore, more nodes are considered by the `Extend` and `Rewire` functions when adding the new state to the tree and further improving the tree. While this increases the runtime of TL-RRT*, it is still less than the runtime of the SMC method for most tasks. On the other hand, biased TL-RRT* with connection radius 0 outperforms the SMC method with “perfect” or “imperfect” initial horizons in terms of both runtime and cost for each task. The reason is that not only is the node q_P^{closest} the only one that is considered when adding the new state to the tree, but also that, for robot i , $\mathbf{x}_i^{\text{new}}$ is the second point along the path from $\mathbf{x}_i^{\text{closest}}$ to $\mathcal{L}(i)$ if the step-size η is large enough. Therefore, connecting q_P^{new} to q_P^{closest} ensures that progress towards $\mathcal{L}(i)$ is made, and, thus, a successful transition to the next state with Büchi component $q_B^{\text{succ},2}$ is more likely to be made at each iteration, accelerating the detection of a feasible path. Due to the use of geodesic shortest paths, the cost achieved by biased TL-RRT* with connection radius 0 is close to that achieved using the connection radius in (17). In our simulation, we compute geodesic paths sequentially for all robots involved in one edge label. Therefore, the runtime can be further improved if these computations are done in parallel.

VIII. CONCLUSION

The majority of existing LTL planning approaches rely on a discrete abstraction of robot mobility to construct a product automaton which is then used to synthesize discrete motion plans. The limitation of these approaches is that both the abstraction process and the control synthesis are computation-

ally expensive and that the resulting discrete plans are only optimal given the discrete abstraction that was used to generate them. In this paper, we proposed a new sampling-based LTL planning algorithm, with unbiased and biased sampling, which does not require any discrete abstraction of robot mobility and avoids the construction of a product automaton. Instead, it builds incrementally a tree that can explore the workspace and the state-space of an NBA that captures a given LTL specification, simultaneously. We showed that our algorithm is probabilistically complete and asymptotically optimal, and provided numerical experiments that showed that our method outperforms relevant temporal planning methods.

REFERENCES

- [1] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.
- [2] H. Choset, K. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. Kavraki, and T. S., "Principles of robot motion: theory, algorithms, and implementations," Boston, MA, 2005.
- [3] I. Arvanitakis, K. Giannousakis, and A. Tzes, "Mobile robot navigation in unknown environment based on exploration principles," in *2016 IEEE Conference on Control Applications (CCA)*. IEEE, 2016, pp. 493–498.
- [4] J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 2. IEEE, 2000, pp. 995–1001.
- [5] L. E. Kavraki, P. Vestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [6] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [7] G. E. Fainekos, H. Kress-Gazit, and G. J. Pappas, "Temporal logic motion planning for mobile robots," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*. IEEE, 2005, pp. 2020–2025.
- [8] M. Guo and M. M. Zavlanos, "Distributed data gathering with buffer constraints and intermittent communication," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 279–284.
- [9] Y. Kantaros and M. M. Zavlanos, "Distributed intermittent connectivity control of mobile robot networks," *IEEE Transactions on Automatic Control*, vol. 62, no. 7, pp. 3109–3121, 2017.
- [10] K. Leahy, D. Zhou, C.-I. Vasile, K. Oikonomopoulos, M. Schwager, and C. Belta, "Persistent surveillance for unmanned aerial vehicles subject to charging and temporal logic constraints," *Autonomous Robots*, vol. 40, no. 8, pp. 1363–1378, 2016.
- [11] C. Baier and J.-P. Katoen, *Principles of model checking*. MIT press Cambridge, 2008, vol. 26202649.
- [12] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.
- [13] —, "Where's waldo? sensor-based temporal logic motion planning," in *Proceedings 2007 IEEE International Conference on Robotics and Automation*. IEEE, 2007, pp. 3116–3121.
- [14] Y. Chen, X. C. Ding, and C. Belta, "Synthesis of distributed control and communication schemes from global ltl specifications," in *2011 50th IEEE Conference on Decision and Control and European Control Conference*. IEEE, 2011, pp. 2718–2723.
- [15] Y. Chen, X. C. Ding, A. Stefanescu, and C. Belta, "Formal approach to the deployment of distributed robotic teams," *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 158–171, 2012.
- [16] E. M. Clarke, O. Grumberg, and D. Peled, *Model checking*. MIT press, 1999.
- [17] S. L. Smith, J. Tůmová, C. Belta, and D. Rus, "Optimal path planning for surveillance with temporal-logic constraints," *The International Journal of Robotics Research*, vol. 30, no. 14, pp. 1695–1708, 2011.
- [18] M. Guo and D. V. Dimarogonas, "Multi-agent plan reconfiguration under local ltl specifications," *The International Journal of Robotics Research*, vol. 34, no. 2, pp. 218–235, 2015.
- [19] M. Kloetzer and C. Belta, "Automatic deployment of distributed teams of robots from temporal logic motion specifications," *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 48–61, 2010.
- [20] A. Ulusoy, S. L. Smith, X. C. Ding, C. Belta, and D. Rus, "Optimality and robustness in multi-robot path planning with temporal logic constraints," *The International Journal of Robotics Research*, vol. 32, no. 8, pp. 889–911, 2013.
- [21] A. Ulusoy, S. L. Smith, and C. Belta, "Optimal multi-robot path planning with ltl constraints: guaranteeing correctness through synchronization," in *Distributed Autonomous Robotic Systems*. Springer, 2014, pp. 337–351.
- [22] D. C. Conner, A. A. Rizzi, and H. Choset, "Composition of local potential functions for global robot control and navigation," in *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)(Cat. No. 03CH37453)*, vol. 4. IEEE, 2003, pp. 3546–3551.
- [23] C. Belta and L. Habetz, "Constructing decidable hybrid systems with velocity bounds," in *2004 43rd IEEE Conference on Decision and Control (CDC)(IEEE Cat. No. 04CH37601)*, vol. 1. IEEE, 2004, pp. 467–472.
- [24] C. Belta, V. Isler, and G. J. Pappas, "Discrete abstractions for robot motion planning and control in polygonal environments," *IEEE Transactions on Robotics*, vol. 21, no. 5, pp. 864–874, 2005.
- [25] M. Kloetzer and C. Belta, "Reachability analysis of multi-affine systems," in *International Workshop on Hybrid Systems: Computation and Control*. Springer, 2006, pp. 348–362.
- [26] D. Boskos and D. V. Dimarogonas, "Decentralized abstractions for multi-agent systems under coupled constraints," *European Journal of Control*, vol. 45, pp. 1–16, 2019.
- [27] Y. Kantaros and M. M. Zavlanos, "Intermittent connectivity control in mobile robot networks," in *49th Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, USA, November, 2015, pp. 1125–1129.
- [28] —, "Sampling-based control synthesis for multi-robot systems under global temporal specifications," in *2017 ACM/IEEE 8th International Conference on Cyber-Physical Systems (ICCPs)*. IEEE, 2017, pp. 3–14.
- [29] —, "Sampling-based optimal control synthesis for multirobot systems under global temporal tasks," *IEEE Transactions on Automatic Control*, vol. 64, no. 5, pp. 1916–1931, 2018.
- [30] —, "Temporal logic optimal control for large-scale multi-robot systems: 10 400 states and beyond," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 2519–2524.
- [31] —, "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.
- [32] X. Luo and M. M. Zavlanos, "Transfer planning for temporal logic tasks," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 5306–5311.
- [33] Y. Kantaros and M. M. Zavlanos, "Distributed optimal control synthesis for multi-robot systems under global temporal tasks," in *Proceedings of the 9th ACM/IEEE International Conference on Cyber-Physical Systems*. IEEE Press, 2018, pp. 162–173.
- [34] Y. Kantaros, B. V. Johnson, S. Chowdhury, D. J. Cappelleri, and M. M. Zavlanos, "Control of magnetic microrobot teams for temporal micromanipulation tasks," *IEEE Transactions on Robotics*, no. 99, pp. 1–18, 2018.
- [35] Y. E. Sahin, P. Nilsson, and N. Ozay, "Provably-correct coordination of large collections of agents with counting temporal logic constraints," in *2017 ACM/IEEE 8th International Conference on Cyber-Physical Systems (ICCPs)*. IEEE, 2017, pp. 249–258.
- [36] S. Karaman and E. Frazzoli, "Linear temporal logic vehicle routing with applications to multi-uav mission planning," *International Journal of Robust and Nonlinear Control*, vol. 21, no. 12, pp. 1372–1395, 2011.
- [37] E. M. Wolff, U. Topcu, and R. M. Murray, "Optimization-based trajectory generation with linear temporal logic specifications," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 5319–5325.
- [38] Y. Shoukry, P. Nuzzo, A. Balkan, I. Saha, A. L. Sangiovanni-Vincentelli, S. A. Seshia, G. J. Pappas, and P. Tabuada, "Linear temporal logic motion planning for teams of underactuated robots using satisfiability modulo convex programming," in *2017 IEEE 56th Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 1132–1137.
- [39] Y. Shoukry, P. Nuzzo, A. L. Sangiovanni-Vincentelli, S. A. Seshia, G. J. Pappas, and P. Tabuada, "Smc: Satisfiability modulo convex programming," *Proceedings of the IEEE*, vol. 106, no. 9, pp. 1655–1679, 2018.

- [40] S. Karaman and E. Frazzoli, “Sampling-based motion planning with deterministic μ -calculus specifications,” in *Proceedings of the 48th IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*. IEEE, 2009, pp. 2222–2229.
- [41] —, “Sampling-based algorithms for optimal motion planning with deterministic μ -calculus specifications,” in *American Control Conference (ACC)*, Montreal, Canada, June 2012, pp. 735–742.
- [42] C. I. Vasile and C. Belta, “Sampling-based temporal logic path planning,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Tokyo, Japan, November 2013, pp. 4817–4822.
- [43] A. Bhatia, L. E. Kavraki, and M. Y. Vardi, “Sampling-based motion planning with temporal goals,” in *International Conference on Robotics and Automation (ICRA)*, Anchorage, AL, May 2010, pp. 2689–2696.
- [44] K. He, M. Lahijanian, L. E. Kavraki, and M. Y. Vardi, “Towards manipulation planning with temporal logic specifications,” in *IEEE International Conference on Robotics and Automation, ICRA 2015, Seattle, WA, USA, 26-30 May, 2015*, 2015, pp. 346–352.
- [45] M. Y. Vardi and P. Wolper, “An automata-theoretic approach to automatic program verification,” in *1st Symposium in Logic in Computer Science (LICS)*. IEEE Computer Society, 1986.
- [46] M. Kloetzer and C. Belta, “A fully automated framework for control of linear systems from temporal logic specifications,” *IEEE Transactions on Automatic Control*, vol. 53, no. 1, pp. 287–297, 2008.
- [47] K. Solovey, L. Janson, E. Schmerling, E. Frazzoli, and M. Pavone, “Revisiting the asymptotic optimality of rrt,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2189–2195.
- [48] B. Lacerda, F. Faruq, D. Parker, and N. Hawes, “Probabilistic planning with formal performance guarantees for mobile service robots,” *The International Journal of Robotics Research*, vol. 38, no. 9, pp. 1098–1123, 2019.
- [49] Y. Akdoğan, C. Kuş, A. Asgharzadeh, İ. Kınacı, and F. Sharafi, “Uniform-geometric distribution,” *Journal of Statistical Computation and Simulation*, vol. 86, no. 9, pp. 1754–1770, 2016.
- [50] Y. Kantaros and M. M. Zavlanos, “Global planning for multi-robot communication networks in complex environments,” *IEEE Transactions on Robotics*, vol. 32, no. 5, pp. 1045–1061, 2016.
- [51] X. Li, Z. Serlin, G. Yang, and C. Belta, “A formal methods approach to interpretable reinforcement learning for robotic planning,” *Science Robotics*, vol. 4, no. 37, 2019.
- [52] R. Khodayi-mehr, Y. Kantaros, and M. M. Zavlanos, “Distributed state estimation using intermittently connected robot networks,” *IEEE Transactions on Robotics*, vol. 35, no. 3, pp. 709–724, 2019.
- [53] D. E. Soltero, S. L. Smith, and D. Rus, “Collision avoidance for persistent monitoring in multi-robot systems with intersecting trajectories,” in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 3645–3652.
- [54] Y. Zhou, H. Hu, Y. Liu, and Z. Ding, “Collision and deadlock avoidance in multirobot systems: A distributed approach,” *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 47, no. 7, pp. 1712–1726, 2017.
- [55] M. Van Kreveld, O. Schwarzkopf, M. de Berg, and M. Overmars, *Computational geometry algorithms and applications*. Springer, 2000.
- [56] P. Gastin and D. Oddoux, “Fast LTL to büchi automata translation,” in *International Conference on Computer Aided Verification*. Springer, 2001, pp. 53–65.
- [57] L. Janson, E. Schmerling, and M. Pavone, “Monte carlo motion planning for robot trajectory optimization under uncertainty,” in *Robotics Research*. Springer, 2018, pp. 343–361.
- [58] L. De Moura and N. Björner, “Z3: An efficient smt solver,” in *International conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 2008, pp. 337–340.
- [59] A. Biere, K. Heljanko, T. Junttila, T. Latvala, and V. Schuppan, “Linear encodings of bounded ltl model checking,” *arXiv preprint cs/0611029*, 2006.
- [60] W. Rudin, *Real and complex analysis*. Tata McGraw-Hill Education, 2006.
- [61] F. Aurenhammer, “Voronoi diagrams—a survey of a fundamental geometric data structure,” *ACM Computing Surveys (CSUR)*, vol. 23, no. 3, pp. 345–405, 1991.

APPENDIX A

PROOF OF THEOREM 6.3

To prove completeness of TL-RRT*, we cannot directly adopt the proof for RRT [4] since in the product state space

considered here the transition relation differs from that in the continuous space. Particularly, while in continuous space a transition is valid between two points if they are connected by an obstacle-free straight line, for LTL task planning a transition is valid if the obstacle-free straight line additionally does not cross any labeled region more than once and the corresponding label enables a transition to at least one Büchi state (otherwise the LTL formula is violated). Nevertheless, we can still draw on the same high-level idea in [4] to show completeness by first proving that the tree can grow arbitrarily close to any state that is one-hop-reachable from the root, and then extending this concept to multi-hop-reachable states.

Before stating our main results, we provide some necessary notations. Let $\mathcal{R}_P(q_P)$ denote the one-hop reachable set of q_P in \mathcal{Q}_P , i.e., $\mathcal{R}_P(q_P) = \{q'_P \in \mathcal{Q}_P \mid q_P \rightarrow_P q'_P\}$. Then, let P_r denote the subspace of $\mathcal{W}_{\text{free}}^N \times \mathcal{Q}_B$, consisting of states that can be reached from the root through a multi-hop path. Furthermore, P_r can be divided into two sets, P_s and P_v , where $P_s = \{q_P \in P_r \mid \mathcal{R}_P(q_P) \neq \emptyset\}$ and $P_v = \{q_P \in P_r \mid \mathcal{R}_P(q_P) = \emptyset\}$. Given a state $q_P \in P_s$, let $q_P^{\text{cp}} \in \mathcal{V}_{\mathcal{T}}$ be its “virtual” candidate parent if there exists a feasible transition from q_P^{cp} to q_P . Recall that q_P^{near} is a candidate parent of q_P^{new} in the `Extend` function if both $\mathbf{x}^{\text{near}} \rightarrow \mathbf{x}^{\text{new}}$ and $(q_B^{\text{near}}, L(\mathbf{x}^{\text{near}}), q_B^{\text{new}}) \in \rightarrow_B$. Here, we use “virtual” since actually q_P is not necessarily sampled but q_P^{new} is. For simplicity, we use the term candidate parent below. Furthermore, let $D_k(q_P)$ denote a random variable whose value is the Euclidean distance between q_P and its nearest candidate parent node $q_P^{\text{ncp}} = (\mathbf{x}^{\text{ncp}}, q_B^{\text{ncp}})$ in a tree of k nodes. One specific realization of $D_k(q_P)$ is denoted by $d_k(q_P)$. Finally, we will use the following theorem.

Theorem A.1: ([60]) Suppose $\{a_n\}$ is a sequence of real numbers and $0 \leq a_n < 1$. Then $\prod_{n=1}^{\infty} (1 - a_n) > 0 \iff \sum_{n=1}^{\infty} a_n < \infty$.

In what follows we first present results that are necessary to show that TL-RRT* is probabilistically complete. For simplicity, D_k and d_k denote $D_k(q_P)$ and $d_k(q_P)$, respectively. Lemmas A.2 to A.5 describe how D_k evolves as the tree grows. Together they establish Propositions A.6 and A.7, which bear a similarity to Lemma 1 in [4] in that both focus on states that can be reached from the root in one hop. In Lemma 1 in [4], the states belong to a convex set whereas Propositions A.6 and A.7 consider one-hop-reachable states explicitly due to different definitions of the transition relations.

Lemma A.2: Let $q_P = (\mathbf{x}, q_B) \in P_s$ be any state that can be reached from the root (\mathbf{x}_0, q_B^0) in a one-hop transition, namely, $q_P \in \mathcal{R}_P(q_B^0)$. Then, $\mathbb{E}(D_{k+1} \mid D_k = d_k) < d_k$.

Proof: Consider any fixed product state $q_p \in \mathcal{R}_P(q_p^0)$. Observe that d_1 is equal to the distance between the root and state q_P when the tree only contains the root. Therefore, $d_k < \infty$ for $k \geq 1$ since $q_P \in \mathcal{R}_P(q_p^0)$ and the sequence $\{d_k\}$ is non-increasing. Given the realization d_k , if the new node q_P^{new} is not a candidate parent of q_P , we have $D_{k+1} = d_k$. Otherwise, there are two cases. First, the distance between q_P^{new} and q_P may be larger than or equal to d_k . In this case, the nearest candidate parent of q_P remains the same, so that $D_{k+1} = d_k$. Second, the distance between q_P^{new} and q_P is

smaller than d_k . In this case, $D_{k+1} < d_k$. Let p_k denote the probability of the event $\{D_{k+1} < d_k\}$, so that the event $\{D_{k+1} = d_k\}$ occurs with probability $1 - p_k$. Then, the expectation $\mathbb{E}(D_{k+1})$ can be written as

$$\mathbb{E}(D_{k+1}|D_k = d_k) = (1 - p_k) d_k + p_k d \quad (20)$$

where $d < d_k$ is a positive number.

In what follows, we show that p_k is strictly positive. Specifically, we define a ball $\mathcal{B}_{d_k}(q_P)$ with radius d_k that is equal to the distance between q_P and its nearest candidate parent node q_P^{ncp} . If $k = 1$, then $q_P^{\text{ncp}} = q_P^0$, because $q_P \in \mathcal{R}_P(q_P^0)$. As the tree grows, it is possible that another node becomes a new q_P^{ncp} for q_P . Since q_P is reachable from q_P^{ncp} , there exists a straight path between \mathbf{x} and \mathbf{x}^{ncp} , and the feasible transition $q_P^{\text{ncp}} \rightarrow_P q_P$ can be potentially an edge/path added to the tree. By Assumption 6.2, for every point (\mathbf{x}', q'_B) on this path, there exists a ball with radius $\delta_{\mathbf{x}'}$ in which all states have the same Büchi state component as (\mathbf{x}', q'_B) . Let δ_{\min} denote the minimum value of all $\delta_{\mathbf{x}'}$, i.e.,

$$\delta_{\min} = \inf\{\delta_{\mathbf{x}'} | \mathbf{x}' = \lambda \mathbf{x} + (1 - \lambda) \mathbf{x}^{\text{ncp}}, 0 \leq \lambda \leq 1\}. \quad (21)$$

Then, a tube neighborhood with radius δ_{\min} along the path connecting q_P^{ncp} and q_P can be constructed, denoted by $\mathcal{N}_{d_k}(q_P)$, which is convex and obstacle-free; see also Fig. 7. Moreover, define the Voronoi partition [61] of the space based on the nodes of the tree and let $\mathcal{C}(q_P^{\text{ncp}})$ denote the Voronoi cell around state q_P^{ncp} . Since the set $\mathcal{N}_{d_k}(q_P)$ has positive measure, i.e., $\mu(\mathcal{N}_{d_k}(q_P)) > 0$, we have that the set $\mathcal{I}_k = \mathcal{C}(q_P^{\text{ncp}}) \cap \mathcal{N}_{d_k}(q_P)$ also has positive measure. Note that we will use the notations $\mathcal{N}_{d_k}(q_P)$, $\mathcal{C}(q_P^{\text{ncp}})$, \mathcal{I}_k and their variations throughout the rest of this paper. Since the unbiased sampling function is supported in the free workspace, the probability that \mathbf{x}^{new} lies within \mathcal{I}_k is non-zero. If this happens, since \mathcal{I}_k is obstacle-free, the condition $\mathbf{x}^{\text{ncp}} \rightarrow \mathbf{x}^{\text{new}}$ is satisfied, which infers that the path $(\mathbf{x}^{\text{ncp}}, q_B^{\text{ncp}}) \rightarrow_P (\mathbf{x}^{\text{new}}, q_B^{\text{new}}) \rightarrow_P (\mathbf{x}, q_B)$ is valid, where $(\mathbf{x}, q_B) = q_P$ and q_B^{new} is either q_B^{ncp} or q_B . Note that node $(\mathbf{x}^{\text{ncp}}, q_B^{\text{ncp}})$ is in the set $\mathcal{Q}_P^{\text{nearest}}$ with respect to q_P^{new} , because \mathbf{x}^{new} is in $\mathcal{C}(q_P^{\text{ncp}})$. Therefore, the state q_P^{new} will be added to the tree, and it will become the new nearest candidate parent of q_P . Hence, $p_k > 0$ and (20) becomes:

$$\begin{aligned} \mathbb{E}(D_{k+1}|D_k = d_k) &= (1 - p_k) d_k + p_k d \\ &< (1 - p_k) d_k + p_k d_k = d_k \end{aligned} \quad (22)$$

which completes the proof. \blacksquare

Remark A.3: In the proof of Lemma A.2, we only consider when q_P^{new} is inside the intersection of \mathcal{I}_k and $\mathcal{B}_{d_k}(q_P)$. However, \mathcal{I}_k involves q_P and the nearest candidate parent q_P^{ncp} . Indeed, the true value of p_k is larger than the one in Lemma A.2. For example, consider when there is another candidate parent q_P^{cp} of q_P which is further from q_P than q_P^{ncp} is. If q_P^{new} is located inside the intersection of $\mathcal{I}'_k = \mathcal{C}(q_P^{\text{cp}}) \cap \mathcal{N}'_{d_k}(q_P)$ and $\mathcal{B}_{d_k}(q_P)$, where $\mathcal{N}'_{d_k}(q_P)$ is generated by q_P^{cp} and q_P , we can also have that $D_{k+1} < d_k$; see Fig. 7.

Let $\{\mathbb{E}(D_k)\}$ denote the sequence $\mathbb{E}(D_1), \mathbb{E}(D_2), \dots$. We show that $\{\mathbb{E}(D_k)\}$ is strictly decreasing, which serves as an intermediate step to prove this sequence converges to 0.

Lemma A.4: $\mathbb{E}(D_{k+1}) < \mathbb{E}(D_k)$ holds.

Proof: Let $f(d_k)$ denote the probability density function of the random variable D_k . Multiplying both sides of (22) by

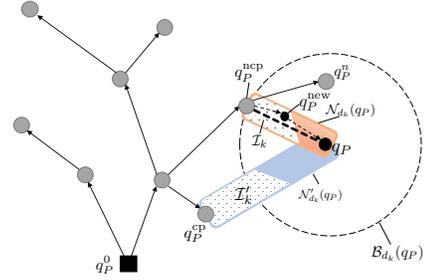


Fig. 7. An illustration of the scenario where the distance between q_P and its nearest candidate parent node drops. Although q_P^{new} is closer to q_P than q_P^{ncp} , it is not the candidate parent of q_P . q_P^{ncp} is the nearest candidate parent of both q_P^{new} and q_P . The orange region denotes the neighborhood around the feasible path depicted as the thick dashed line, from q_P^{ncp} to q_P , and the thin dashed line illustrates the alternative of the path in the tree. The dotted area denotes the intersection \mathcal{I}_k of the Voronoi cell $\mathcal{C}(q_P^{\text{ncp}})$ and the neighborhood around the path. q_P^{new} is inside \mathcal{I}_k . The blue region depicts the case for another candidate parent q_P^{cp} where $q_P^{\text{new}} \in \mathcal{I}'_k \cap \mathcal{B}_{d_k}(q_P)$.

$f(d_k)$ and integrating over the support of D_k , we have

$$\int_{D_k} \mathbb{E}(D_{k+1}|D_k = d_k) f(d_k) dd_k < \int_{D_k} d_k f(d_k) dd_k. \quad (23)$$

By the law of total expectation $\mathbb{E}(X) = \mathbb{E}(\mathbb{E}(X|Y))$, where X and Y are two random variables, the left-hand side in (23) is equivalent to $\mathbb{E}(D_{k+1})$. By definition, the right-hand side in (23) is $\mathbb{E}(D_k)$. Thus, $\mathbb{E}(D_{k+1}) < \mathbb{E}(D_k)$. \blacksquare

Note that the sequence $\{\mathbb{E}(D_k)\}$ is lower bounded by 0 since $q_P \in \mathcal{R}_P(q_P^0)$. The following lemma shows that the decreasing sequence $\{\mathbb{E}(D_k)\}$ in fact converges to zero.

Lemma A.5: We have that $\mathbb{E}(D_k) \rightarrow 0$ as $k \rightarrow \infty$.

Proof: We will show this proposition by contradiction. Specifically, assume that $\inf\{\mathbb{E}(D_k)\} = b > 0$, which means that for any $\epsilon > 0$, there exists $K \in \mathbb{N}^+$ so that for all $k > K$, we have $b < \mathbb{E}(D_k) < b + \epsilon$. Note that $\mathbb{E}(D_k) \neq b$, otherwise $\mathbb{E}(D_{k+1}) < b$ since $\{\mathbb{E}(D_k)\}$ is strictly decreasing by Lemma A.4. Moreover, let $k_\epsilon = \inf\{k : \mathbb{E}(D_k) < b + \epsilon\}$ denote the number of nodes in the tree when $\mathbb{E}(D_k) < b + \epsilon$ holds for the first time. Hence, $\mathbb{E}(D_{k_\epsilon}) < b + \epsilon$.

Define the ball $\mathcal{B}_{d_{k_\epsilon}/2}(q_P)$ with radius $d_{k_\epsilon}/2$, centered at q_P , where d_{k_ϵ} is the realization of D_{k_ϵ} ; see Fig. 8. Recall that \mathcal{I}_{k_ϵ} denotes the intersection of the Voronoi cell around q_P^{ncp} and the neighborhood of the feasible path from q_P^{ncp} to q_P ; see the proof of Lemma A.2. If q_P^{new} lies within \mathcal{I}_{k_ϵ} , then we have $D_{k_\epsilon+1} < d_{k_\epsilon}$. Let $\gamma_{k_\epsilon} = \sup\{\|q'_P\|_{\mathcal{W}_{\text{free}}^N} - \|\mathbf{x}^{\text{ncp}}\| : q'_P \in \mathcal{I}_{k_\epsilon} \cap \text{line}(q_P^{\text{ncp}}, q_P)\}$, where $\text{line}(q_P^{\text{ncp}}, q_P)$ is the line segment connecting q_P^{ncp} and q_P . The quantity γ_{k_ϵ} is the maximum possible decrease from d_{k_ϵ} if $q_P^{\text{new}} \in \mathcal{I}_{k_\epsilon} \cap \text{line}(q_P^{\text{ncp}}, q_P)$. Since $\mu(\mathcal{I}_{k_\epsilon}) > 0$, we have that $\gamma_{k_\epsilon} > 0$. Therefore, we can draw a ball of radius $d_{k_\epsilon} - \gamma_{k_\epsilon}/2$ centered at q_P that is contained in the ball $\mathcal{B}_{d_{k_\epsilon}}(q_P)$. If q_P^{new} is inside the intersection of \mathcal{I}_{k_ϵ} and $\mathcal{B}_{(d_{k_\epsilon} - \gamma_{k_\epsilon}/2)}(q_P)$, denoted by $\mathcal{I}\mathcal{B}_{k_\epsilon}$, we have that $D_{k_\epsilon+1} < d_{k_\epsilon} - \gamma_{k_\epsilon}/2$, whose probability satisfies $p_{k_\epsilon+1} > 0$.

Now let $\{D_{k_\epsilon+m}\}$ denote a sequence of random variables starting from D_{k_ϵ} for any $m \in \mathbb{N}^+$. Note that as the tree grows, more candidate parents of q_P will appear. For this, it suffices that q_P^{new} lies in the intersection of the Voronoi cell around one of the candidate parents q_P^{cp} , the neighborhood $\mathcal{N}_{d_{k_\epsilon}}(q_P)$ of the feasible path connecting q_P^{cp} and q_P , and the ball $\mathcal{B}_{(d_{k_\epsilon} - \gamma_{k_\epsilon}/2)}(q_P)$. Let $\mathcal{Q}_{k_\epsilon+i}^{\text{cp}}$ denote the set of all candidate parents of q_P for $i \in \{0, \dots, m-1\}$, and $\cup \mathcal{I}\mathcal{B}_{k_\epsilon+i}$

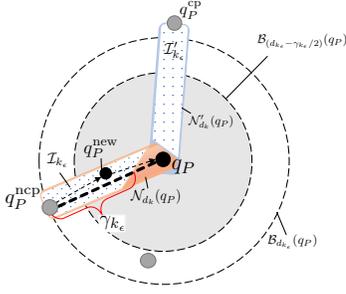


Fig. 8. An illustration of the scenario where $D_{k_{e+1}}$ drops below $d_{k_e} - \gamma_{k_e}/2$. The brown region illustrates the neighborhood around the feasible path depicted as the thick dashed line, from q_P^{ncp} to q_P , and the thin dashed line illustrates alternative of the path in the tree. The dotted area denotes the intersection \mathcal{I}_{k_e} of the Voronoi cell around q_P^{ncp} and the neighborhood around the path. q_P^{new} is inside \mathcal{IB}_{k_e} . The blue region depicts the case for another candidate parent q_P^{cp} where $q_P^{\text{new}} \in \mathcal{I}_{k_e} \cap \mathcal{B}_{(d_{k_e} - \gamma_{k_e}/2)}(q_P)$.

denote the union of the intersections $\mathcal{IB}_{k_{e+i}}^1, \dots, \mathcal{IB}_{k_{e+i}}^{|Q_{k_{e+i}}^{\text{cp}}|}$, associated with every candidate parent node of q_P in $Q_{k_{e+i}}^{\text{cp}}$. Following a similar argument as above, we can show that, if $q_P^{\text{new}} \in \cup \mathcal{IB}_{k_{e+i}}$, then $D_{k_{e+i+1}} < d_{k_e} - \gamma_{k_e}/2$, which occurs with probability $p_{k_{e+i+1}}$. Note that $\{D_{k_{e+m}} < d_{k_e} - \gamma_{k_e}/2\}$ will occur if there exists $i \in \{0, \dots, m-1\}$ such that $D_{k_{e+i+1}} < d_{k_e} - \gamma_{k_e}/2$. Then, conditioned on $D_{k_e} = d_{k_e}$, the event $\{d_{k_e} - \gamma_{k_e}/2 < D_{k_{e+m}} \leq d_{k_e}\}$ occurs with probability $\prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}})$. Moreover, with probability $1 - \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}})$, the event $\{D_{k_{e+m}} \leq d_{k_e} - \gamma_{k_e}/2\}$ occurs. Therefore, $\mathbb{E}(D_{k_{e+m}} | D_{k_e} = d_{k_e})$ can be bounded as:

$$\mathbb{E}(D_{k_{e+m}} | D_{k_e} = d_{k_e}) \leq d_{k_e} \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}}) + (d_{k_e} - \gamma_{k_e}/2) \left[1 - \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}}) \right]. \quad (24)$$

Multiplying both sides of (24) by $f(d_{k_e})$, integrating over the support of D_{k_e} , and denoting by $\Delta = \frac{1}{2} \int_{D_{k_e}} \gamma_{k_e} f(d_{k_e}) dd_{k_e} > 0$, we obtain

$$\begin{aligned} \mathbb{E}(D_{k_{e+m}}) &\leq \mathbb{E}(D_{k_e}) \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}}) \\ &\quad + (\mathbb{E}(D_{k_e}) - \Delta) \left[1 - \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}}) \right] \\ &< (b + \epsilon) \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}}) \\ &\quad + (b + \epsilon - \Delta) \left[1 - \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}}) \right] \\ &= (b + \epsilon) - \left[1 - \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}}) \right] \Delta. \end{aligned} \quad (25)$$

where the strict inequality comes from the fact $\mathbb{E}(D_{k_e}) < b + \epsilon$.

By Assumption 6.2, the sequence $\{p_{k_{e+i+1}}\}_{i=0}^{m-1}$ is bounded away from zero for all $i \in \{0, \dots, m-1\}$ for any $m \in \mathbb{N}^+$. It does not converge to 0 since the ball $\mathcal{B}_{d_{k_e}}(q_P)$ is fixed given k_e and increasingly more candidate parents of q_P will appear as the tree grows, their position components forming a dense subset of \mathcal{W}^N . By Assumption 6.2, $\mathcal{IB}_{k_{e+i}}$ covers a subset of \mathcal{W}^N with positive measure. Therefore $\sum_{i=0}^{\infty} p_{k_{e+i+1}} = \infty$. Moreover, since $0 < p_{k_{e+i+1}} < 1$, by Theorem A.1, we can conclude that $\lim_{m \rightarrow \infty} \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}}) = 0$. Hence, $\lim_{m \rightarrow \infty} 1 - \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}}) = 1$. Then, as $m \rightarrow \infty$, (25) becomes $\mathbb{E}(D_{k_{e+m}}) < (b + \epsilon) - \Delta$. Since ϵ can be any arbitrary small non-negative number, we can select $\epsilon = \Delta/2$. Then $\mathbb{E}(D_{k_{e+m}}) < b - \Delta/2 < b$ as $m \rightarrow \infty$, which contradicts the assumption that b is the infimum. Hence, 0 is the infimum of the sequence $\{\mathbb{E}(D_k)\}$. By the monotone convergence theorem, we get that $\mathbb{E}(D_k) \rightarrow 0$. ■

Using the previous lemmas for unbiased sampling, the

following proposition asserts that the tree built by TL-RRT* rooted at $q_P^0 = (\mathbf{x}^0, q_B^0)$ will grow arbitrarily close to any state $q_P = (\mathbf{x}, q_B) \in P_s$ that is one-hop-reachable from the root.

Proposition A.6: For any state $q_P = (\mathbf{x}, q_B) \in P_s \cap \mathcal{R}_P(q_P^0)$ in the tree and for any positive constant $\delta > 0$, we have that $\lim_{k \rightarrow \infty} \mathbb{P}(D_k(q_P) \leq \delta) = 1$.

Proof: According to Lemma A.5, $\mathbb{E}(D_k) \rightarrow 0$ as $k \rightarrow \infty$. Following the same logic, we can get that $\mathbb{E}(D_k^2) \rightarrow 0$. Hence, the variance $\text{Var}(D_k) = \mathbb{E}(D_k^2) - [\mathbb{E}(D_k)]^2 \rightarrow 0$. Since $\mathbb{E}(D_k) \rightarrow 0$ and $D_k \geq 0$, for any small number $\delta > 0$ there exists a sufficiently large $k \in \mathbb{N}^+$, so that $\mathbb{E}(D_k) < \delta/2$. Therefore, we obtain $\mathbb{P}(D_k > \delta) \leq \mathbb{P}(|D_k - \mathbb{E}(D_k)| > \delta/2) \leq 4 \text{Var}(D_k)/\delta^2$, where the first inequality holds because (i) $D_k \leq |D_k - \mathbb{E}(D_k)| + \mathbb{E}(D_k)$ by the triangle inequality, and (ii) $\mathbb{E}(D_k) < \delta/2$, and the second is obtained by Chebyshev inequality. Since $\text{Var}(D_k) \rightarrow 0$, we get that $\mathbb{P}(|D_k| > \delta) \rightarrow 0$ and further that $D_k \rightarrow 0$ in probability, namely, for any $\delta > 0$, it holds that $\lim_{k \rightarrow \infty} \mathbb{P}(D_k \leq \delta) = 1$. ■

Next, we show that the nodes in the tree that are arbitrarily close to q_P share the same Büchi state as q_P .

Proposition A.7: Given any state $q_P = (\mathbf{x}, q_B) \in P_s \cap \mathcal{R}_P(q_P^0)$, the probability that there exists a node in the tree constructed by TL-RRT*, which shares the same Büchi state as q_P and lies within distance $\delta > 0$ from q_P satisfies

$$\lim_{k \rightarrow \infty} \mathbb{P} \left(\left\{ \exists q_P' = (\mathbf{x}', q_B) \in \mathcal{V}_k^{\text{TL-RRT}^*} : \|\mathbf{x} - \mathbf{x}'\| \leq \delta \right\} \right) = 1,$$

where $\mathcal{V}_k^{\text{TL-RRT}^*}$ is the set of nodes in the tree constructed by TL-RRT* when the number of nodes is k .

Proof: To see this, recall from the proof of the Lemma A.5 that with probability $1 - \prod_{i=0}^{m-1} (1 - p_{k_{e+i+1}})$, the event $\{D_{k_{e+m}} \leq d_{k_e} - \gamma_{k_e}/2\}$ occurs. Next, we show this is also the probability that there is a node inside $\mathcal{B}_{(d_{k_e} - \gamma_{k_e}/2)}(q_P)$ which has the same Büchi state as q_P . The reason is that if the new point \mathbf{x}^{new} lies within the set $\cup \mathcal{IB}_{k_{e+i}}$, defined in the proof of Lemma A.5, then q_B can be paired with \mathbf{x}^{new} to create a new node, because (i) $\mathbf{x}^{\text{cp}} \rightarrow \mathbf{x}^{\text{new}}$ holds since \mathbf{x}^{new} is in $\mathcal{I}_{k_{e+i}}$, and (ii) $q_B^{\text{cp}} \xrightarrow{L(\mathbf{x}^{\text{cp}})} q_B$ since $q_P^{\text{cp}} \rightarrow_P q_P$. When $m \rightarrow \infty$, this probability goes to 1,

$$\lim_{k \rightarrow \infty} \mathbb{P} \left(\left\{ \exists q_P' = (\mathbf{x}', q_B) \in \mathcal{V}_k^{\text{TL-RRT}^*} : \|\mathbf{x} - \mathbf{x}'\| < \delta \right\} \right) = 1,$$

which completes the proof. ■

Propositions A.6 and A.7 consider the states in P_s that can be reached from the root in one hop. Next we extend Propositions A.6 and A.7 to any reachable state in P_s , which is similar to Lemma 2 in [4].

Proposition A.8: Let Assumptions 6.1 and 6.2 hold and further assume that sampling of the free workspace is unbiased. For any state $q_P = (\mathbf{x}, q_B) \in P_s$ and any positive number $\delta > 0$, we have that $\lim_{k \rightarrow \infty} \mathbb{P}(D_k(q_P) < \delta) = 1$. Also, the probability that there exists a node in the tree that shares the same Büchi state as q_P and lies within a ball of radius δ centered at q_P , goes to 1 as $k \rightarrow \infty$.

Proof: We prove this inductively, using Propositions A.6 and A.7. Specifically, let q_P^0 denote the root of the tree and assume there exists a feasible path between q_P and the root. Then, there exists a sequence of states $\{q_P^0, q_P^1, \dots, q_P^l\}$, where $q_P^i = (\mathbf{x}^i, q_B^i)$, $q_P^l = q_P$, and q_P^i can make one-hop transition to $q_P^{i+1} \forall i \in \{0, \dots, l-1\}$.

By Proposition A.7, since q_P^1 is reachable from

q_P^0 through a one-hop transition, we have that $\lim_{k \rightarrow \infty} \mathbb{P}(\{\exists (\mathbf{x}', q_B^1) \in \mathcal{V}_k^{\text{TL-RRT}^*} : \|\mathbf{x}_1 - \mathbf{x}'_1\| < \delta\}) = 1$ for any $\delta > 0$. In other words, as $k \rightarrow \infty$, the tree will contain a node arbitrarily close to q_P^1 with the same Büchi state as q_P^1 . Now, consider any pair of nodes in the sequence (\mathbf{x}_i, q_B^i) and $(\mathbf{x}_{i+1}, q_B^{i+1})$ such that $(\mathbf{x}_{i+1}, q_B^{i+1}) \in \mathcal{R}_P(\mathbf{x}_i, q_B^i)$. Viewing (\mathbf{x}'_i, q_B^i) as a new initial state, which is very close to (\mathbf{x}_i, q_B^i) , we can apply Propositions A.6 and A.7 to get that (\mathbf{x}'_i, q_B^i) can reach a node $(\mathbf{x}'_{i+1}, q_B^{i+1})$, possibly through multiple edges in the tree, that is arbitrarily close to $(\mathbf{x}_{i+1}, q_B^{i+1})$. Since $q_P = q_P^1$, we conclude that $\lim_{k \rightarrow \infty} \mathbb{P}(D_k(q_P) < \delta) = 1$, and $\lim_{k \rightarrow \infty} \mathbb{P}(\{\exists q'_P = (\mathbf{x}', q_B) \in \mathcal{V}_k^{\text{TL-RRT}^*} : \|\mathbf{x} - \mathbf{x}'\| < \delta\}) = 1$. ■

Using Proposition A.8, we can show the probabilistic completeness of TL-RRT* with unbiased sampling in Theorem 6.3. Specifically, if feasible plans exist, then accepting states exist for both the prefix and suffix plans. For the construction of the prefix plan, the accepting states should be in the set \mathcal{P}_s , satisfying the LTL specification and reachable from the root. By Proposition A.8, the tree will grow arbitrarily close to an accepting state and it will contain a node with the same Büchi state as the accepting state. This completes the construction of the prefix plan. Moreover, after the predetermined number of iterations, the set \mathcal{P} is returned. Then, for the construction of the suffix plans, the same analysis can be applied assuming that the root becomes one certain element of \mathcal{P} . Therefore a feasible plan will be found with probability 1.

APPENDIX B PROOF OF COROLLARY 6.4

The proof is based on a similar analysis as in Lemmas A.2-A.5 that are used to derive Propositions A.6 and A.7. As in Appendix A, we exclude `Extend` and `Rewire` functions and connect q_P^{new} to q_P^{closest} directly. The key idea in Lemmas A.2 to A.5, is that the probability that \mathbf{x}^{new} lies within the ball $\mathcal{B}_{d_k}(q_P)$ and is close enough to one candidate parent q_P^{cp} of q_P is bounded away from 0 at each iteration. As a result, $q_P^{\text{new}} = (\mathbf{x}^{\text{new}}, q_B^{\text{cp}})$ or $q_P^{\text{new}} = (\mathbf{x}^{\text{new}}, q_B)$ is the new nearest candidate parent of q_P . Using biased sampling, this argument still holds. Recall that the biased TL-RRT* differs from the unbiased one in the use of the biased sampling and the way to detect a cycle around the root. Since RRT* is used to detect the cycle and it is probabilistically complete, it suffices to focus on the use of the biased sampling. First, since $p_{\text{closest}} \in (0.5, 1)$, following the distribution in (11), any node has a non-zero probability to be selected as q_P^{closest} . Thus, increasingly more candidate parents appear in the tree, as in the case of unbiased sampling. These candidate parents form a dense subset of \mathcal{W}^N , so that the probability that any one of these candidate parents q_P^{cp} is selected as q_P^{closest} is bounded away from 0. In the biased sampling, the selection of q_P^{closest} comes first and then the sampling of \mathbf{x}^{new} follows. Thus, we need to ensure first that a candidate parent is selected as q_P^{closest} . Second, recall that $\mathcal{N}_{d_k}(q_P)$ denotes the neighborhood around the path connecting q_P^{cp} and q_P . Since $p_{\text{rand}} \in (0.5, 1)$ and the collection of $\mathcal{N}_{d_k}(q_P)$ covers a subset of \mathcal{W}_N , the probability that \mathbf{x}^{new} lies within the neighborhood region $\mathcal{N}_{d_k}(q_P)$ is bounded away from 0, as well. Here, unlike the unbiased case, we consider $\mathcal{N}_{d_k}(q_P)$ rather than \mathcal{I}_k , the intersection of the Voronoi cell

$\mathcal{C}(q_P^{\text{cp}})$ and $\mathcal{N}_{d_k}(q_P)$. This is because once q_P^{cp} is selected as q_P^{closest} , those samples lying within $\mathcal{N}_{d_k}(q_P)$ will produce a new nearest candidate parent of q_P , while in the unbiased case, a stricter constraint is imposed which requires the candidate parent to be the nearest node to \mathbf{x}^{new} , which means \mathbf{x}^{new} should lie within $\mathcal{C}(q_P^{\text{cp}})$ first. For these two reasons, the same arguments as in the unbiased case can be made here too, and counterparts of Propositions A.6 and A.7 for biased sampling can be obtained. The proof of Corollary 6.4 can be completed by inheriting the proofs of Proposition A.8 and Theorem 6.3.

APPENDIX C PROOF OF THEOREM 6.5

The proof for the asymptotic optimality of TL-RRT* is based on that for RRT* [6], [47], where [47] fixes a logical gap in the proof in [6]. Our proof differs from [47] mainly in that we consider here a product state space and, thus, the definitions of transition relations differ.

1) *Construction of the product path p_ϵ* : This step differs from [6], [47] in that here we build a product path that lives in the combined continuous and discrete state space. Specifically, let τ^* be the optimal continuous path that optimizes (2). Since Problem 1 is robustly feasible, following [47] there exists a robust path τ_ϵ such that $J(\tau_\epsilon) \leq (1 + \epsilon/4)J(\tau^*)$, where $\epsilon \in (0, 1)$. Next, as discussed in Section VI-B, we can construct a product path p_ϵ from τ_ϵ .

2) *Construction of the sequence of sets of balls $\{\mathcal{B}_n\}_{n \in \mathbb{N}}$* : This step differs from [6], [47] in that here we construct a set of balls along the product path p_ϵ instead of the continuous path τ_ϵ . Specifically, given the product path p_ϵ , we define a set of M_n balls $\mathcal{B}_n = \{\mathcal{B}_{n,1}, \dots, \mathcal{B}_{n,M_n}\}$ that cover the whole product path p_ϵ , each with radius $q_n = r_n(\mathcal{V}_T)/(2 + \theta)$ and center on p_ϵ , where $\theta \in (0, 1/4)$. The centers of two consecutive balls are $l_n = \theta q_n$ apart. The center of ball $\mathcal{B}_{n,1}$ is the initial state of the product path, and the center of ball \mathcal{B}_{n,M_n} is the accepting state; see also Fig. 5. The construction of the balls here is the same, in a geometric sense, as that in [6], [47] where a set of balls is constructed along τ_ϵ rather than p_ϵ . Thus, we can adapt Lemmas 53 and 54 in [6] to obtain that for any two points, $(\mathbf{x}_m, \cdot) \in \mathcal{B}_{n,m}$ and $(\mathbf{x}_{m+1}, \cdot) \in \mathcal{B}_{n,m+1}$ in any two consecutive balls, the following properties hold: (i) the Euclidean distance between \mathbf{x}_m and \mathbf{x}_{m+1} is no larger than the radius $r_n(\mathcal{V}_T)$ used in Alg. 3 and 4; (ii) the line $\overline{\mathbf{x}_m \mathbf{x}_{m+1}}$ lies entirely within $\mathcal{W}_{\text{free}}^N$.

Note that the radii of the balls in the set \mathcal{B}_n decrease to 0 as $r_n \rightarrow 0$, and the distance between the centers of any two consecutive balls in \mathcal{B}_n also goes to 0. By Assumption 6.2, for every point on the subpath connecting any two consecutive centers, $(\mathbf{x}_i, q_B^i) \in \mathcal{B}_{n,i}$ and $(\mathbf{x}_{i+1}, q_B^{i+1}) \in \mathcal{B}_{n,i+1}$, $\forall i \in [M_n - 1]$, there exists a ball in which all states have the same Büchi state. Similar to the definition of the neighborhood $\mathcal{N}_{d_k}(q_P)$ in (21), which is shown in Fig. 7, we denote by $\mathcal{N}_{n,i}$ the neighborhood along the line connecting centers (\mathbf{x}_i, q_B^i) and $(\mathbf{x}_{i+1}, q_B^{i+1})$. We can conclude that there exists a large $N_0 \in \mathbb{N}$, such that $\mathcal{B}_{n,i} \subset \mathcal{N}_{n,i}$ and $\mathcal{B}_{n,i+1} \subset \mathcal{N}_{n,i}$ for any iteration $n > N_0$ and any two consecutive balls, because both the radii of $\mathcal{B}_{n,i}$ and $\mathcal{B}_{n,i+1}$ go to 0 and the distance between the two centers also goes to 0. Furthermore, since the balls in

\mathcal{B}_n are constructed along p_ϵ , when $n > N_0$, the center (\mathbf{x}_i, q_B^i) can transition to the center $(\mathbf{x}_{i+1}, q_B^{i+1})$, i.e. (i) $\mathbf{x}_i \rightarrow \mathbf{x}_{i+1}$, and (ii) $q_B^i \xrightarrow{L(\mathbf{x}_i)} q_B^{i+1}$. In what follows, we assume $n > N_0$.

Moreover, we define a tree \mathcal{G}_n that contains an edge starting from (\mathbf{x}, q_B) and ending at (\mathbf{x}', q_B') if (R1): $\|\mathbf{x} - \mathbf{x}'\| \leq r_n(\mathcal{V}_T)$; (R2): \mathbf{x} is sampled before \mathbf{x}' ; (R3): $\mathbf{x} \rightarrow \mathbf{x}'$; and (R4): $q_B \xrightarrow{L(\mathbf{x})} q_B'$. Given the same sequence of samples, the cost of the best path in \mathcal{G}_n from the root to any node is not less than that generated by TL-RRT*. This is because the `Extend` function in TL-RRT* will connect the new node (\mathbf{x}', q_B') to the node (\mathbf{x}, q_B) in the tree that incurs the least cost as long as rules (R1)-(R4) are satisfied, and the `Rewire` function will further decrease the cost of the nodes in the tree, again as long as these rules are satisfied. Thus, if the algorithm generating \mathcal{G}_n is asymptotically optimal, so is TL-RRT*.

3) *Connecting nodes in consecutive balls in $\{\mathcal{B}_n\}_{n \in \mathbb{N}}$* : To prove the optimality of \mathcal{G}_n , we show that a path exists in \mathcal{G}_n that is arbitrarily close to p_ϵ . To do so, we need to show that eventually every ball in \mathcal{B}_n contains at least one node of \mathcal{G}_n .

Given a sequence of uniformly sampled position states, we can reason about the existence of samples inside the projections of the balls $\{\mathcal{B}_n\}_{n \in \mathbb{N}}$ onto the continuous space and the order in which they were sampled. We denote by \mathbf{X}_j the sample point in continuous space at the j -th iteration and by $V = \{\mathbf{X}_1, \dots, \mathbf{X}_n\}$ the sample set up to the n -th iteration. Next, we partition the set V into $M_n + 1$ subsets V_0, \dots, V_{M_n} respecting the order in which the points in V were sampled. Specifically, V_0 contains the first n' samples, where $n' = \kappa n$, $\kappa \in (0, 1)$, and the subsequent subsets V_i contain $z = \lfloor (n - n')/M_n \rfloor$ samples each, except for the last subset which contains the remaining points in V , that is, $V_i = \{\mathbf{X}_j \mid j \in \{n' + (i-1) \cdot z, \dots, n' + i \cdot z\}\}$, for $i \in \{1, \dots, M_n - 1\}$. Then, Lemma 2 in [47] proves that with probability one there is an ascending subsequence V' of samples in V , one sample per V_i , that correspond one-to-one to the balls in $\{\mathcal{B}_n\}$ in terms of order. Next we define conditions under which rules (R1)-(R4) hold so that \mathcal{G}_n can connect the samples in V' , along with their Büchi states, to generate a path that approximates p_ϵ . These conditions are different than those required in [47].

Specifically, consider the product path in its discretized form $p_\epsilon = (\mathbf{x}_1, q_B^1), (\mathbf{x}_2, q_B^2), \dots, (\mathbf{x}_{M_n}, q_B^{M_n})$, where each point is the center of a ball constructed in 2). For simplicity, consider the first two balls. Assume the center (\mathbf{x}_1, q_B^1) of the first ball can take a one-hop transition to the second center (\mathbf{x}_2, q_B^2) , therefore, $\mathbf{x}_1 \rightarrow \mathbf{x}_2$ holds. Consider also a point \mathbf{x}'_2 located inside $\mathcal{B}_{n,2}$ that is sampled after \mathbf{x}_1 . Then, $\mathbf{x}_1 \rightarrow \mathbf{x}'_2$ does not necessarily hold, even though $\mathcal{B}_{n,1} \subset \mathcal{N}_{n,1}, \mathcal{B}_{n,2} \subset \mathcal{N}_{n,1}$. This is because more than one boundaries $\partial \ell_j$ can cross the balls $\mathcal{B}_{n,1}$ and $\mathcal{B}_{n,2}$, in which case the straight line $\mathbf{x}_1 \mathbf{x}'_2$ will also cross these boundaries; see Fig. 9(a). As a result, rule (R3) is not satisfied and the graph \mathcal{G}_n can not connect \mathbf{x}_1 to \mathbf{x}'_2 . However, as n grows to infinity, the radii of the balls in the set \mathcal{B}_n go to 0, and, therefore, the probability that more than one boundaries cross the balls $\mathcal{B}_{n,1}$ and $\mathcal{B}_{n,2}$ goes to 0. Thus, as n grows, at most one boundary $\partial \ell_j$ of a labeled region ℓ_j can cross one or both balls $\mathcal{B}_{n,1}$ and $\mathcal{B}_{n,2}$; see Fig. 9(b). In

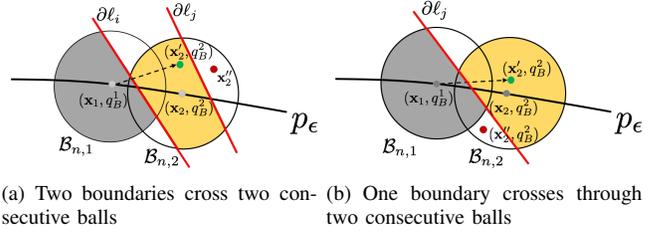


Fig. 9. Graphical depiction of cases where boundaries of regions cross balls. In Fig. 9(a), two different boundaries $\partial \ell_i$ and $\partial \ell_j$ cross the first two consecutive balls. In Fig. 9(b), a boundary crosses through two consecutive balls, the shaded parts illustrate where the centers belong.

this case, rule (R3) is satisfied and \mathbf{x}_1 can be connected to \mathbf{x}'_2 in $\mathcal{B}_{n,2}$. In what follows, we show that the exact location of \mathbf{x}'_2 in $\mathcal{B}_{n,2}$ is important in order to further satisfy rule (R4). Specifically, when n is large, this boundary can be locally approximated by a hyperplane, which divides $\mathcal{B}_{n,2}$ into two parts. Suppose \mathbf{x}'_2 and \mathbf{x}_2 are in different parts of the ball $\mathcal{B}_{n,2}$. Then, the Büchi state q_B^2 can be paired with \mathbf{x}'_2 , because in Alg. 2, the Büchi state space is searched exhaustively. But since $L(\mathbf{x}_2) \neq L(\mathbf{x}'_2)$, it is possible that q_B^2 can not transition at all or can transition to a state other than q_B^3 , i.e., the Büchi state of the center of $\mathcal{B}_{n,3}$. In this case, the plan under construction will not proceed along p_ϵ . To avoid this situation, \mathbf{x}'_2 should lie within that part of $\mathcal{B}_{n,2}$ that contains its center, as opposed to [47] where to follow τ_ϵ , it suffices that \mathbf{x}'_2 lies anywhere within $\mathcal{B}_{n,2}$. Following this logic, to satisfy rule (R4), the samples in V' should lie in those parts of the balls that contain their centers. Let E_n^1 be the event that the projection of that part of ball $\mathcal{B}_{n,i}$ that contains its center on the continuous space contains at least one sample in V_i , i.e., $E_n^1 = \{\forall 1 \leq i \leq M_n, \mathcal{B}'_{n,i} \cap V_i \neq \emptyset\}$, where $\mathcal{B}'_{n,i}$ is the part of the ball that contains its center so that $\mathcal{B}'_{n,i} = \mathcal{B}_{n,i}$ if $\mathcal{B}_{n,i}$ is not crossed by a boundary, and $\mathcal{B}'_{n,i} \subset \mathcal{B}_{n,i}$ otherwise, and the intersection $\mathcal{B}'_{n,i} \cap V_i$ is defined on the projection of $\mathcal{B}'_{n,i}$ on the continuous space. Assuming that the event E_n^1 occurs, next we show that if rules (R1)-(R4) are satisfied, then indeed \mathcal{G}_n can connect samples in the subsequence V' , augmented by Büchi states, to get a path that approximates p_ϵ .

Let n be large enough so that at most one boundary crosses the balls in the set $\{\mathcal{B}_n\}$, and consider the node $(\mathbf{x}'_3, q_B^3) \in \mathcal{B}'_{n,3}$ with \mathbf{x}'_3 drawn after \mathbf{x}'_2 . Then, since \mathbf{x}'_2 and \mathbf{x}'_3 are located in two consecutive balls and we have proved in step 2) that the distance between any states in any two consecutive balls is no more than $r_n(\mathcal{V}_T)$, (R1) is satisfied. Furthermore, since \mathbf{x}'_2 is sampled before \mathbf{x}'_3 , (R2) holds. Recall from step 2) that the straight line connecting any two states in any two consecutive balls is entirely inside $\mathcal{W}_{\text{free}}^N$, and that those balls are shrinking, so eventually there is at most one boundary crossing through two consecutive balls. Therefore, rule (R3) is also satisfied.

As for rule (R4), note that $q_B^2 \xrightarrow{L(\mathbf{x}'_2)} q_B^3$ since \mathbf{x}'_2 belongs to the same part of the ball that contains its center and the center (\mathbf{x}_2, q_B^2) can transition to the center (\mathbf{x}_3, q_B^3) , that is, $q_B^2 \xrightarrow{L(\mathbf{x}_2)} q_B^3$. Therefore, (R4) is also satisfied and $(\mathbf{x}'_3, q_B^3) \in \mathcal{B}'_{n,3}$ will be added to the tree \mathcal{G}_n . Iterating over the sequence of balls, we conclude that the tree \mathcal{G}_n will have nodes in each one of the parts $\mathcal{B}'_{n,i}$ that contain their centers. Since every $\mathcal{B}'_{n,i}$ contains one point in V' , we finally conclude that \mathcal{G}_n

can connect the samples in the subsequence V' .

In what follows, we prove that the probability of event E_n^1 approaches 1. First, let E_n^2 denote the event that $\mathbf{x}^{\text{new}} = \mathbf{x}^{\text{rand}}$ at the n -th iteration. Letting $n \rightarrow \infty$ and using the fact that $\lim_{n \rightarrow \infty} \mathbb{P}(E_n^2) = 1$ shown in [47], we have that

$$\begin{aligned} \mathbb{P}(E_n^1) &= \mathbb{P}(E_n^1|E_n^2)\mathbb{P}(E_n^2) + \mathbb{P}(E_n^1|\overline{E_n^2})\mathbb{P}(\overline{E_n^2}) \\ &\geq \mathbb{P}(E_n^1|E_n^2)\mathbb{P}(E_n^2) \\ &= (1 - \mathbb{P}(\overline{E_n^1}|E_n^2))(1 - \mathbb{P}(\overline{E_n^2})) \\ &= 1 - \mathbb{P}(\overline{E_n^2}) - \mathbb{P}(\overline{E_n^1}|E_n^2) + \mathbb{P}(\overline{E_n^1}|E_n^2)\mathbb{P}(\overline{E_n^2}) \\ &= 1 - \mathbb{P}(\overline{E_n^1}|E_n^2), \end{aligned} \quad (26)$$

where $\overline{E_n^2}$ is the complement of the event E_n^2 . Next, let c_n denote the number of balls that are crossed by a boundary. Then, we have that

$$\mathbb{P}(\overline{E_n^1}|E_n^2) = \mathbb{P}(\exists 1 \leq i \leq M_n, \mathcal{B}'_{n,i} \cap V_i = \emptyset) \quad (27)$$

$$\leq \sum_{i=1}^{M_n} \mathbb{P}(\mathcal{B}'_{n,i} \cap V_i = \emptyset) = \sum_{i=1}^{M_n} \left(1 - \frac{\mu(\mathcal{B}'_{n,i})}{\mathcal{W}_{\text{free}}^N}\right)^{|V_i|} \quad (28)$$

$$\leq (M_n - c_n) \left(1 - \frac{\mu(\mathcal{B}_{n,i})}{\mathcal{W}_{\text{free}}^N}\right)^{|V_i|} + c_n \left(1 - \frac{\mu(\mathcal{B}_{n,i})}{2\mathcal{W}_{\text{free}}^N}\right)^{|V_i|} \quad (29)$$

$$\leq M_n \left(1 - \frac{\mu(\mathcal{B}_{n,i})}{\mathcal{W}_{\text{free}}^N}\right)^{|V_i|} + c_n \left(1 - \frac{\mu(\mathcal{B}_{n,i})}{2\mathcal{W}_{\text{free}}^N}\right)^{|V_i|}, \quad (30)$$

where $|V_i| = (n - n')/M_n$, (28) is obtained assuming that \mathbf{x}^{new} is sampled uniformly which holds for large n since $\mathbf{x}^{\text{new}} = \mathbf{x}^{\text{rand}}$, and to derive (29), we divide all balls into two classes based on whether they are crossed by a boundary or not. For those balls crossed by a boundary, we have $\mu(\mathcal{B}'_{n,i}) \geq \mu(\mathcal{B}_{n,i})/2$; otherwise, $\mu(\mathcal{B}'_{n,i}) = \mu(\mathcal{B}_{n,i})$. According to [47], the first term in (30) approaches 0 when $\gamma_{\text{TL-RRT}^*} \geq (2 + \theta) \left(\frac{(1+\epsilon/4)J(\tau^*)}{(\text{dim}+1)\theta(1-\kappa)} \cdot \frac{\mu(\mathcal{W}_{\text{free}}^N)}{\zeta_{\text{dim}}}\right)^{\frac{1}{\text{dim}+1}}$, and also $\left(1 - \frac{\mu(\mathcal{B}_{n,i})}{\mathcal{W}_{\text{free}}^N}\right)^{|V_i|} \leq \exp\{-\xi\gamma_{\text{TL-RRT}^*}^{\text{dim}+1} \log n\}$, where $\xi := \frac{\theta\zeta_{\text{dim}}(1-\kappa)}{J(\tau^*)(2+\theta)^{\text{dim}+1}\mu(\mathcal{W}_{\text{free}}^N)}$. Following the same logic, we can show that the fraction $\frac{1}{2}$ in the second term in (30) can be absorbed into ξ and $\left(1 - \frac{\mu(\mathcal{B}_{n,i})}{2\mathcal{W}_{\text{free}}^N}\right)^{|V_i|} \leq \exp\{-\frac{1}{2}\xi\gamma_{\text{TL-RRT}^*}^{\text{dim}+1} \log n\}$. Since c_n is a small number (because there are only a finite number of boundaries), we have that the second term in (30) approaches 0. Thus, $\mathbb{P}(\overline{E_n^1}|E_n^2)$ approaches 0. Substituting in (26) we have that $\mathbb{P}(E_n^1)$ converges to 1 as $n \rightarrow \infty$. This guarantees that \mathcal{G}_n will contain a path that approximates p_ϵ .

The rest of the proof is similar to that in [47]. Essentially, we show that the probability of the event E_n^3 that at most αM_n of the balls $\mathcal{B}_{n,i}^\beta$ do not contain any samples from V_i approaches 1, where $\alpha \in (0, \theta\epsilon/16)$, $\beta \in (0, \theta\epsilon/16)$, and $\mathcal{B}_{n,i}^\beta$ shares the same center with $\mathcal{B}_{n,i}$ but its radius is a fraction β of that of $\mathcal{B}_{n,i}$. In this way, we exclude the case where the tree \mathcal{G}_n has a path approximating p_ϵ but zig-zaging around it. We omit the details due to space limitations. Recall from step 1) that the cost of the path τ_ϵ is smaller than $(1 + \epsilon/4)J(\tau^*)$. Then, when events E_n^1 and E_n^3 occur, as shown in [47], the cost of the path returned by \mathcal{G}_n is less than $(1 + \epsilon)J(\tau^*)$, which is slightly larger than $(1 + \epsilon/4)J(\tau^*)$, since this path is an approximation of p_ϵ .

APPENDIX D

PROOF OF THEOREM 6.6

The key step to obtain $\gamma_{\text{TL-RRT}^*}^{\text{unbiased}}$ for uniform sampling lies in upper-bounding the probability $\mathbb{P}(\overline{E_n^1}|E_n^2)$ in (27) by a quantity that converges to 0. Although the biased samples

are guided towards given labeled regions, according to (14), with probability $(1 - y_{\text{rand}})^N$ the sample \mathbf{x}^{rand} still follows the uniform distribution. Thus, within the whole sequence of n biased samples, there is a subsequence that contains $(1 - y_{\text{rand}})^N \cdot n$ samples on average that were sampled uniformly. We focus only on this subsequence of uniform samples and proceed in the same way as in the proof for the unbiased sampling. Specifically, for this subsequence, $|V_i|$ in (30) will become $(1 - y_{\text{rand}})^N(n - n')/M_n$. Moreover, as before, the radius of the balls $\{\mathcal{B}_n\}$ in Fig. 5 will be proportional to the connection radius $r_n(\mathcal{V}_\tau)$ that depends on the length n of the whole sequence. Therefore, in (30), $|V_i|$ will change, and so will $\mu(\mathcal{B}_{n,i})$ and M_n since they are related to $\gamma_{\text{TL-RRT}^*}^{\text{biased}}$. Following a similar derivation as in [47] that shows that the first term in (30) approaches 0, we can show that the extra term $(1 - y_{\text{rand}})^N$ in $|V_i|$ will enter the denominator of $\gamma_{\text{TL-RRT}^*}^{\text{unbiased}}$, and the upper bound on $\mathbb{P}(\overline{E_n^1}|E_n^2)$ will still converge to 0. Note that the upper bound on $\mathbb{P}(\overline{E_n^1}|E_n^2)$ obtained using the whole sequence of biased samples will be smaller than that obtained using the subsequence of uniform samples. This is because the samples that are biased towards accepting states can also lie in the balls $\{\mathcal{B}_n\}$, further reducing the chance that a ball $\mathcal{B}_{n,i}$ does not contain a sample in V_i . Since this bound will be smaller, the probability of E_n^1 will converge faster to 1, which explains why biased TL-RRT* is faster than the unbiased version.



Xusheng Luo (S'19) received the B.Eng. and M.S.E. degrees in aerospace engineering from the Harbin Institute of Technology, Harbin, China, in 2015 and 2017, respectively, and the Ph.D. degree in mechanical engineering from Duke University, Durham, NC, in 2020. His research interest focuses on the formal control synthesis under high-level specifications with applications in robotics.



Yiannis Kantaros (S'14M'19) received the Diploma in Electrical and Computer Engineering in 2012 from the University of Patras, Patras, Greece, and the M.S.E and Ph.D. degrees in mechanical engineering from Duke University, Durham, NC, in 2017 and 2018, respectively. His research interests include distributed control, machine learning, formal methods and control synthesis with applications in robotics. He received the Best Student Paper Award at the 2nd IEEE Global Conference on Signal and Information Processing in 2014 and the 2017-18 Outstanding Dissertation Research Award from the Department of Mechanical Engineering and Materials Science, Duke University, NC.



Michael M. Zavlanos (S'05M'09SM'19) received the Diploma in mechanical engineering from the National Technical University of Athens, Greece, in 2002, and the M.S.E. and Ph.D. degrees in electrical and systems engineering from the University of Pennsylvania, Philadelphia, PA, in 2005 and 2008, respectively. He is currently an Associate Professor in the Department of Mechanical Engineering and Materials Science at Duke University, Durham, NC. His research focuses on control theory, optimization, and learning and, in particular, autonomous systems and robotics, networked and distributed control systems, and cyber-physical systems. Dr. Zavlanos is a recipient of various awards including the 2014 ONR YIP Award and the 2011 NSF CAREER Award.