

Motion and Action Planning under LTL Specifications using Navigation Functions and Action Description Language

Meng Guo, Karl H. Johansson and Dimos V. Dimarogonas

Abstract— We propose a novel framework to combine model-checking-based motion planning with action planning using action description languages, aiming to tackle task specifications given as Linear Temporal Logic (LTL) formulas. The specifications implicitly require both sequential regions to visit and the desired actions to perform at these regions. The robot’s motion is abstracted based on sphere regions of interest in the workspace and the structure of navigation function(NF)-based controllers, while the robot’s action map is constructed based on precondition and effect functions associated with the actions. An optimal planner is designed that generates the discrete motion-and-action plan fulfilling the specification, as well as the low-level hybrid controllers that implement this plan. The whole framework is demonstrated by a case study.

I. INTRODUCTION

Navigation functions proposed by Rimon and Koditschek in [16] provide an easy-to-implement and provably correct point-to-point navigation algorithm, which has been successfully applied in both single [21] and multi-agent [8] navigation under different geometric constraints. Formal high-level languages such as Linear Temporal Logic (LTL) and Computation Tree Logic (CTL) allow us to describe more complex planning objectives [3], [9], [17], [27]. Attempts to combine the strengths of both frameworks have recently appeared in [10]. A high-level discrete plan is synthesized by off-the-shelf model-checking algorithms given the finite abstraction of robot’s motion and the LTL task specification over the abstraction. This approach has been modified to take into account other navigating techniques like probabilistic roadmap method [7], [14] and rapidly-exploring random trees [1] under certain complex motion objectives.

However, to solve planning problems of practical interest it is often necessary to perform various actions at different regions to achieve a goal. In other words, the purpose of “going somewhere” is to “do something”. The plan would be a combination of transitions among different places and performing various sequential actions. It would be inadequate to carry out the motion planning and action planning independently since the motion plan and action plan are closely related, i.e., “where to go” is motivated by “what to do there” and “what to do now” depends on “where it has been”. Another observation is that some actions can only be performed when certain conditions are fulfilled

and as a result certain state variables might be changed. Action description languages like STRIPS [11], ADL [23] and PDDL [22] provide an intuitive and powerful way for describing the preconditions and effects of different actions. There has not been much work about how to incorporate the action description formalism with motion planning in general, and NF-based navigating techniques in particular.

Some relevant work takes into account this aspect. In [27], since the underlying actions can only be performed at fixed regions, the specification is reinterpreted in terms of regions to visit. In [18], since the actions can be turned on and off at any time, independent atomic propositions are created for each behavior. The above approaches will not be applicable if some behaviors can only be performed when the workspace and the robot itself satisfy certain conditions, or choices have to be made among all the behaviors.

We propose to separate the domain-specific knowledge [19] such as the workspace model and the robot’s mobility in the workspace, from the domain-independent knowledge such as the action map based on the actions the robot is capable of (given its on-board hardware and preprogrammed functionalities). One advantage is the increased modularity that our framework is adaptable whenever the workspace is modified or the task specification is changed. Another benefit is the considerably reduced size of the planning problem (compared with the exponential statespace [5] of classic planning using STRIPS), by avoiding unfolding certain state variables. The main contribution of this work is a generic framework to derive the complete description of the robot’s functionalities within a certain workspace, such that any LTL specification in terms of desired motions and actions can be treated. Furthermore, an optimal discrete plan that satisfies the specification and a hybrid controller that implements this plan are designed in a fully automated manner.

The rest of the paper is organized as follows: Section II provides the essential preliminaries. In Sections III and IV, we discuss about the discretized abstraction of robot’s mobility and actions. The way to synthesize the discrete plan and the low-level hybrid controller is in Section VII-C. Numerical simulations are presented in Section VII.

II. PRELIMINARIES

A. Navigation Function

The navigation function proposed by Rimon and Koditschek in [16] is given by $\Phi(q) = \frac{\gamma}{(\gamma^k + \beta)^{\frac{1}{k}}}$, where $k > 0$ is a design parameter, $q \in \mathbb{R}^n$ and $\Phi \in [0, 1]$. In more detail, $\gamma = \|q - q_d\|^2$ represents an attractive potential

The authors are with the ACCESS Linnaeus Center, School of Electrical Engineering, KTH Royal Institute of Technology, SE-100 44, Stockholm, Sweden. mengg, kallej, dimos@kth.se. This work was supported by the Swedish Research Council (VR) and EU STREP RECONFIG: FP7-ICT-2011-9-600825. The first and third authors are also affiliated with the KTH Centre for Autonomous Systems. The authors would like to thank Prof. John S. Baras for helpful suggestions.

from the goal where $q_d \in \mathbb{R}^n$ is the desired goal position, $\beta = \prod_{j=0}^M \beta_j$ is a repulsive potential from the sphere obstacles, where $\beta_0 \triangleq r_0^2 - \|q - q_0\|^2$ and $\beta_j \triangleq \|q - q_j\|^2 - r_j^2$. q_0, r_0 are the center and radius of the allowed workspace $\mathcal{W}_0 = \{q \in \mathbb{R}^n \mid \beta_0 > 0\}$; q_j, r_j are the center and radius of the sphere obstacles $\mathcal{O}_j = \{q \in \mathbb{R}^n \mid \beta_j < 0\}$, $j = 1, \dots, M$. It is assumed that \mathcal{W}_0 and \mathcal{W}_{obs} satisfy the definition of a valid workspace in [16]. By following the negated gradient $-\nabla_q \Phi$, a collision free path is guaranteed from almost any initial position in the free space (except a set of measure zero) to any goal position in the free space.

B. LTL and Büchi Automaton

We focus on the task specification φ given as an Linear Temporal Logic (LTL) formula. The basic ingredients of an LTL formula are a set of atomic propositions (APs) and several boolean and temporal operators. Atomic propositions are Boolean variables that can be either true or false. LTL formulas are formed according to the following grammar [2]: $\varphi ::= \text{True} \mid a \mid \varphi_1 \wedge \varphi_2 \mid \neg \varphi \mid \bigcirc \varphi \mid \varphi_1 \cup \varphi_2$, where $a \in AP$ and \bigcirc (*next*), \cup (*until*). For brevity, we omit the derivations of other useful operators like \square (*always*), \diamond (*eventually*), \Rightarrow (*implication*) and refer the readers to Chapter 5 of [2].

Given an LTL formula φ over a set of atomic propositions AP , there exists a Nondeterministic Büchi automaton (NBA) \mathcal{A}_φ over 2^{AP} corresponding to φ , which is defined as:

$$\mathcal{A}_\varphi = (Q, 2^{AP}, \delta, Q_0, \mathcal{F}), \quad (1)$$

where Q is a finite set of states; Q_0 is the initial state, 2^{AP} is an alphabets; $\delta \subseteq Q \times 2^{AP} \times Q$ is a transition relation and $\mathcal{F} \subseteq Q$ is a set of accepting states. An infinite run r of a NBA is an infinite sequence of states and is called accepting if $\text{Inf}(r) \cap \mathcal{F} \neq \emptyset$ where $\text{Inf}(r)$ is the set of states that appear in r infinitely often. There are fast translation algorithms [12] from an LTL formula to NBA. This translation process can be done in time and space $2^{\mathcal{O}(|\varphi|)}$ [2].

III. ABSTRACTION OF ROBOT MOBILITY

A. The Workspace Model

The workspace we consider is bounded by a large sphere $\pi_0 = \{q \in \mathbb{R}^n \mid \|q - q_0\| \leq r_0\}$, within which there exists N smaller spheres around the points of interest: $\pi_i = \mathcal{B}_{r_i}(q_i) = \{q \in \mathbb{R}^n \mid \|q - q_i\| \leq r_i\}$, where $q_i \in \mathbb{R}^n$ and $r_i > 0$ are the center and radius of the n -dimensional sphere areas $\mathcal{B}_{r_i}(q_i)$; r_i represent the margins with respect to the points of interest, $\forall i = 1, \dots, N$. Denote by $\Pi = \{\pi_1, \dots, \pi_N\}$ the set of smaller sphere areas. It is assumed that π_0 and Π satisfy the conditions of a valid workspace [16]. Compared with other cell decomposition schemes like triangles [4], polygons [3] and hexagons [24], this sphere-area-based approach typically reduces the size of the resulting abstraction since they represent regions of interest, rather than a complete partition of the workspace. Moreover, in order to indicate the robot's position, we define $\Psi_r = \{\Psi_{r,i}\}$, where $\Psi_{r,i} = \text{True}$ if $q \in \pi_i$ and False , otherwise. $\Psi_{r,i}$ can be evaluated by measurements from a real-time position system.

Beside the location of these regions, we also would like to know the properties satisfied by each region, denoted by $\Psi_p = \{\Psi_{p,1}, \dots, \Psi_{p,l}\}$. Ψ_p may reflect the designer's concern, or is relevant to robot's actions discussed in Section IV.

B. Robot Dynamics

The robot is assumed to be holonomic with the single-integrator dynamics: $\dot{q} = u$, where $q, u \in \mathbb{R}^n$ are the position and control signal. We denote by $\pi_g, \pi_s \in \Pi$ the goal and initial region. It is important to notice that the workspace remains *valid* no matter how π_g and π_s are chosen. There always exist a feasible path that leads the robot from one point in π_s to one point in π_g , without crossing other regions in Π and always staying in π_0 . The feasible path can be generated by following the negated gradient:

$$u = -\nabla_q \frac{\gamma_g}{(\gamma_g^k + \beta_{gs})^{\frac{1}{k}}}, \quad (2)$$

where $\beta_{gs} = \prod_{j=0, j \neq s, g}^M \beta_j$. The way to construct γ_g, β_j and compute the gradient $\nabla_q \Phi$ is introduced in Section II-A. Note that the asymptotic stability of the above controller guarantees the convergence to π_g in finite time [21].

Remark 1: [20] also provides the NF-based control strategy for non-holonomic vehicles. The rest of the framework still applies by replacing (2) with the one proposed in [20].

C. Robot's Mobility

Given the workspace model and the robot dynamics, the model of robot's mobility is built as a weighted finite transition system (wFTS):

$$\mathcal{M} = (\Pi_{\mathcal{M}}, \text{act}_{\mathcal{M}}, \longrightarrow_{\mathcal{M}}, \Pi_{\mathcal{M},0}, \Psi_{\mathcal{M}}, L_{\mathcal{M}}, W_{\mathcal{M}}), \quad (3)$$

where (i) $\Pi_{\mathcal{M}} = \{\pi_1, \dots, \pi_N\}$ is the finite set of states; (ii) $\text{act}_{\mathcal{M}}$ represents the controller (2); (iii) $\longrightarrow_{\mathcal{M}} \subseteq \Pi_{\mathcal{M}} \times \text{act}_{\mathcal{M}} \times \Pi_{\mathcal{M}}$ is the transition relation; (iv) $\Pi_{\mathcal{M},0} \subseteq \Pi_{\mathcal{M}}$ is the initial region the robot starts from; (v) $\Psi_{\mathcal{M}} = \Psi_r \cup \Psi_p$ is the set of APs; (vi) $L_{\mathcal{M}} : \Pi_{\mathcal{M}} \rightarrow 2^{\Psi_{\mathcal{M}}}$ is the labeling function for the properties that are satisfied at each region and $\Psi_{r,i} \in L_{\mathcal{M}}(\pi_i)$, $i = 1, \dots, N$; (vii) $W_{\mathcal{M}} : \longrightarrow_{\mathcal{M}} \rightarrow \mathbb{R}^+$ represents the implementation (energy/time) cost [13], approximated by the line distance $\|q_i - q_j\| - r_i - r_j$, $\forall \pi_i, \pi_j \in \Pi$, $i \neq j$.

Remark 2: The workspace π_0 is not included in $\Pi_{\mathcal{M}}$ as it is guaranteed by (2) that the robot always stays within the workspace. Moreover, $\Psi_{r,i} \in L_{\mathcal{M}}(\pi_i)$ by definition.

Up to this point, given the wFTS \mathcal{M} and an LTL task specification over $\Psi_{\mathcal{M}}$, various existing frameworks [3], [15], [27] could be utilized to synthesize a discrete motion plan in terms of sequences of regions to visit, as also our approach discussed in Section VI-A. However when the task specification is stated as requirements on desired actions within different regions, the mobility abstraction \mathcal{M} alone is not enough and a model of robot's actions is also needed.

IV. MODEL OF ROBOT ACTIONS

Classic planning formalisms, like STRIPS [11], ADL [23] and PDDL [22], provide an intuitive way to describe high-level actions the robot is capable of. Given a set of states

TABLE I
ACTIONS DESCRIPTION FOR SECTION VII

Action	Condition	Effect
$act_{\mathcal{B},0}$	True	$\Psi_{b,0} = \text{True}$, $\Psi_{b,\sim 0} = \text{False}$
$act_{\mathcal{B},1}$	$\Psi_{p,1} \& \neg \Psi_{s,1} \& \neg \Psi_{s,2}$	$\Psi_{s,1} = \text{True}$, $\Psi_{b,1} = \text{True}$, $\Psi_{b,\sim 1} = \text{False}$
$act_{\mathcal{B},2}$	$\Psi_{s,1}$	$\Psi_{s,1} = \text{False}$, $\Psi_{b,2} = \text{True}$, $\Psi_{b,\sim 2} = \text{False}$
$act_{\mathcal{B},3}$	$\Psi_{p,2} \& \neg \Psi_{s,2} \& \neg \Psi_{s,1}$	$\Psi_{s,2} = \text{True}$, $\Psi_{b,3} = \text{True}$, $\Psi_{b,\sim 3} = \text{False}$
$act_{\mathcal{B},4}$	$\Psi_{s,2}$	$\Psi_{s,2} = \text{False}$, $\Psi_{b,4} = \text{True}$, $\Psi_{b,\sim 4} = \text{False}$
$act_{\mathcal{B},5}$	True	$\Psi_{b,5} = \text{True}$, $\Psi_{b,\sim 5} = \text{False}$

and action names, each action is described by specifying its precondition and effect on the states. Here we utilize the same approach. Assume that the robot is capable of performing K different actions $\{act_{\mathcal{B},1}, \dots, act_{\mathcal{B},K}\}$, implemented by the corresponding low-level controllers $\{\mathcal{K}_k\}$, $k = 1, 2, \dots, K$. For brevity, denote by $Act_{\mathcal{B}} = \{act_{\mathcal{B},0}, act_{\mathcal{B},1}, \dots, act_{\mathcal{B},K}\}$, where $act_{\mathcal{B},0} \triangleq \text{None}$ indicates that none of these K actions is performed. Moreover, we introduce another two sets of atomic propositions:

- $\Psi_s = \{\Psi_{s,j}\}$, represents the internal states of the robot, $j = 1, 2, \dots, J$, e.g., “the robot has product A”.
- $\Psi_b = \{\Psi_{b,k}\}$ where $\Psi_{b,k} = \text{True}$ if and only if action k is performed. We assume that any two actions cannot be concurrent, i.e., only one element of Ψ_b can be true.

The subscripts of Ψ_s , Ψ_b stand for the “state” and “behavior” of the robot. With Ψ_p , Ψ_s and Ψ_b , we describe each action in $Act_{\mathcal{B}}$ by the precondition and effect functions.

A. Precondition and Effect

The precondition function $\text{Cond} : Act_{\mathcal{B}} \times 2^{\Psi_p} \times 2^{\Psi_s} \rightarrow \text{True/False}$, takes one action in $Act_{\mathcal{B}}$, subsets of 2^{Ψ_p} and 2^{Ψ_s} as inputs and returns a boolean value. In order to perform that action, the conditions on the properties of the workspace Ψ_p and the robot’s internal states Ψ_s have to be fulfilled. For instance, the action “pickup A” can only be performed when “the room has product A”. While some actions like “take pictures” might be performed without such constraints and then the condition is simply a tautology, e.g., $\text{Cond} = \text{True}$. Note the condition function for $act_{\mathcal{B},0}$ is defined as **True**.

The effect function $\text{Eff} : Act_{\mathcal{B}} \times (2^{\Psi_s} \times \Psi_b) \rightarrow (2^{\Psi_s} \times \Psi_b)$, represents the effect of the actions. As a result of performing action $act_{\mathcal{B},k}$, the robot’s internal states Ψ_s might be changed and Ψ_b is changed to indicate which action is performed. For example, once the action “pickup A” is performed, the propositions “the robot has A” and “‘pickup A’ is performed” become true. Note that the effect functions can not modify the properties of the workspace. More specifically,

- $\text{Eff}(act_{\mathcal{B},0}, w_s, \Psi_{b,k}) = (w_s, \Psi_{b,0})$, where $w_s \subseteq 2^{\Psi_s}$ and $\forall \Psi_{b,k} \in \Psi_b$. Performing $act_{\mathcal{B},0}$ does not change

the robot’s internal state and all elements in Ψ_b except $\Psi_{b,0}$ are set to false;

- $\text{Eff}(act_{\mathcal{B},k}, w_s, \Psi_{b,l}) = (w'_s, \Psi_{b,k})$, where $w_s, w'_s \subseteq 2^{\Psi_s}$ and $\Psi_{b,l}, \Psi_{b,k} \in \Psi_b$ by definition for $k \neq 0$.

B. Action Map

Given Ψ_p , Ψ_s , Ψ_b and $Act_{\mathcal{B}}$, Cond , Eff , the *action map* is defined as a tuple

$$\mathcal{B} = (\Pi_{\mathcal{B}}, Act_{\mathcal{B}}, \Psi_p, \hookrightarrow_{\mathcal{B}}, \Pi_{\mathcal{B},0}, \Psi_{\mathcal{B}}, L_{\mathcal{B}}, W_{\mathcal{B}}), \quad (4)$$

where (i) $\Pi_{\mathcal{B}} \subseteq 2^{\Psi_s} \times \Psi_b$ is set of all assignments of Ψ_s and Ψ_b ; (ii) Ψ_p serves as the input propositions, and 2^{Ψ_p} is the finite set of possible input assignments; (iii) the conditional transition relation $\hookrightarrow_{\mathcal{B}}$ is defined by $\pi_{\mathcal{B}} \times \alpha_{\mathcal{B}} \times 2^{\Psi_p} \times \pi'_{\mathcal{B}} \subseteq \hookrightarrow_{\mathcal{B}}$ if the following conditions hold: (1) $\alpha_{\mathcal{B}} \in Act_{\mathcal{B}}$, $\pi_{\mathcal{B}}, \pi'_{\mathcal{B}} \in \Pi_{\mathcal{B}}$; (2) $\text{Cond}(\alpha_{\mathcal{B}}, 2^{\Psi_p}, \pi_{\mathcal{B}}) = \text{True}$; (3) $\pi'_{\mathcal{B}} \in \text{Eff}(\alpha_{\mathcal{B}}, \pi_{\mathcal{B}})$. (iv) $\Pi_{\mathcal{B},0} \subseteq 2^{\Psi_s} \times \Psi_{b,0}$ is the initial state; (v) $\Psi_{\mathcal{B}} = \Psi_s \cup \Psi_b$ is the set of atomic propositions; (vi) $L_{\mathcal{B}}(\pi_{\mathcal{B}}) = \{\pi_{\mathcal{B}}\}$, i.e., the labeling function is the state itself; (vii) $W_{\mathcal{B}} : \hookrightarrow_{\mathcal{B}} \rightarrow \mathbb{R}^+$ is the weight associated with each transition and estimated by the cost of action $\alpha_{\mathcal{B}}$.

Remark 3: The set of states $\Pi_{\mathcal{B}}$ is defined as $2^{\Psi_s} \times \Psi_b$ instead of $2^{\Psi_s} \times 2^{\Psi_b}$ because only one element in Ψ_b can be true. This reduces the size of the action map significantly.

Note that the action map is constructed independently of the structure of the workspace where the robot is deployed. Ψ_p can be viewed as external inputs [2] to the action map. Thus given an instance of Ψ_p , the action map \mathcal{B} is equivalent to a wFTS as all conditional transition relations can be verified or falsified based on the definition of $\hookrightarrow_{\mathcal{B}}$.

V. MODEL OF COMPLETE FUNCTIONALITIES

As mentioned earlier, the abstraction of robot’s mobility \mathcal{M} from (3) and the robot’s action map \mathcal{B} from (4) are adequate for the controller synthesis within certain problem domain. However, in order to consider richer and more complex tasks involving both regions to visit and actions to perform within these regions, we need a complete model of robot’s functionalities that combines these two parts. We propose the following way to compose \mathcal{M} and \mathcal{B} :

$$\mathcal{R} = (\Pi_{\mathcal{R}}, Act_{\mathcal{R}}, \rightarrow_{\mathcal{R}}, \Pi_{\mathcal{R},0}, \Psi_{\mathcal{R}}, L_{\mathcal{R}}, W_{\mathcal{R}}), \quad (5)$$

where (i) $\Pi_{\mathcal{R}} = \Pi_{\mathcal{M}} \times \Pi_{\mathcal{B}}$ is the set of states; (ii) $Act_{\mathcal{R}} = act_{\mathcal{M}} \cup Act_{\mathcal{B}}$ is the set of actions; (iii) $\rightarrow_{\mathcal{R}} \subseteq \Pi_{\mathcal{R}} \times Act_{\mathcal{R}} \times \Pi_{\mathcal{R}}$ is the transition relation, defined by the following rules:

- (1) $\langle \pi_{\mathcal{M}}, \pi_{\mathcal{B}} \rangle \xrightarrow{act_{\mathcal{M}}} \langle \pi'_{\mathcal{M}}, \pi'_{\mathcal{B}} \rangle$ if $\pi_{\mathcal{M}} \xrightarrow{act_{\mathcal{M}}} \pi'_{\mathcal{M}}$ and $\pi_{\mathcal{B}} \xrightarrow{act_{\mathcal{B},0}} \pi'_{\mathcal{B}}$;
- (2) $\langle \pi_{\mathcal{M}}, \pi_{\mathcal{B}} \rangle \xrightarrow{\alpha_{\mathcal{B}}} \langle \pi_{\mathcal{M}}, \pi'_{\mathcal{B}} \rangle$ if $\pi_{\mathcal{B}} \times \alpha_{\mathcal{B}} \times L_{\mathcal{M}}(\pi_{\mathcal{M}}) \times \pi'_{\mathcal{B}} \subseteq \hookrightarrow_{\mathcal{B}}$, where $\alpha_{\mathcal{B}} \in Act_{\mathcal{B}}$;

(iv) $\Pi_{\mathcal{R},0} = \Pi_{\mathcal{M},0} \times \Pi_{\mathcal{B},0}$ contains the robot’s initial region and initial internal state; (v) $\Psi_{\mathcal{R}} = \Psi_{\mathcal{M}} \cup \Psi_{\mathcal{B}}$ is the complete set of atomic propositions including Ψ_r , Ψ_p , Ψ_s and Ψ_a ; (vi) $L_{\mathcal{R}} : \Pi_{\mathcal{R}} \rightarrow 2^{\Psi_{\mathcal{R}}}$ is the labeling function, $L_{\mathcal{R}}(\langle \pi_{\mathcal{M}}, \pi_{\mathcal{B}} \rangle) = L_{\mathcal{M}}(\pi_{\mathcal{M}}) \cup L_{\mathcal{B}}(\pi_{\mathcal{B}})$; (vii) $W_{\mathcal{R}} : \rightarrow_{\mathcal{R}} \rightarrow \mathbb{R}^+$, is the weight function on each transition, defined as:

$$(1) \quad W_{\mathcal{R}}(\langle \pi_{\mathcal{M}}, \pi_{\mathcal{B}} \rangle, act_{\mathcal{M}}, \langle \pi'_{\mathcal{M}}, \pi'_{\mathcal{B}} \rangle) = W_{\mathcal{M}}(\pi_{\mathcal{M}}, act_{\mathcal{M}}, \pi'_{\mathcal{M}});$$

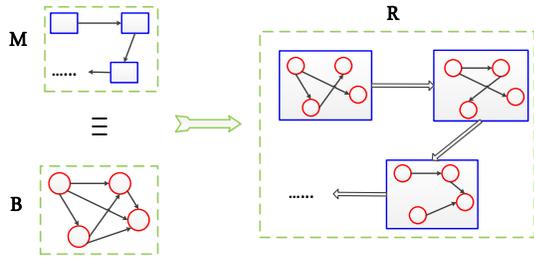


Fig. 1. The action map \mathcal{B} is composed with each region of \mathcal{M} , giving a complete description of robot's functionalities.

(2) $W_{\mathcal{R}}(\langle \pi_{\mathcal{M}}, \pi_{\mathcal{B}} \rangle, \alpha_{\mathcal{R}}, \langle \pi'_{\mathcal{M}}, \pi'_{\mathcal{B}} \rangle) = W_{\mathcal{B}}(\pi_{\mathcal{B}}, \alpha_{\mathcal{R}}, \pi'_{\mathcal{B}})$, if $\alpha_{\mathcal{R}} \in Act_{\mathcal{B}}$.

Remark 4: In the definition of $\hookrightarrow_{\mathcal{B}}$, $act_{\mathcal{B},0}$ is released automatically whenever the controller (2) is activated, because whenever the robot moves to a new region, this automatically indicates that no actions within $act_{\mathcal{B},k}$ are performed as we assume non-concurrent actions.

Figure 1 illustrates the process of parallel composition above. Blue squares represent the states of \mathcal{M} and red cycles encode the states of \mathcal{B} . When composing them into \mathcal{R} , N copies of \mathcal{B} are first created, corresponding to the N regions within the workspace. At the same time, the conditional transition relations in these copies are verified or falsified by verifying the conditions on the properties of each region.

It is possible to construct the complete model \mathcal{R} directly. But there are several advantages in constructing \mathcal{M} and \mathcal{B} first and then composing them into \mathcal{R} . Normally an abstraction for the robot's mobility is only valid for a specific workspace and it has to be modified whenever the structure of the workspace is changed. In contrast, the abstraction for robot's actions is relatively fixed, considering commercial robots with pre-programmed functionalities. By building \mathcal{M} and \mathcal{B} separately, each of them independently can serve as the input model to the controller synthesis machinery in Section VI-A. Moreover, \mathcal{R} is more difficult to build manually than \mathcal{M} and \mathcal{B} separately and then taking their composition automatically. Another advantage is that the number of states in \mathcal{R} is greatly reduced (compared with the normal exponential complexity [5] of classic planning). We avoid unfolding those propositional variables associated with the workspace properties as they are invariant and use the fact that only one element in Ψ_r or Ψ_b can be true.

The composed system \mathcal{R} is a wFTS over the set of atomic propositions $\Psi_{\mathcal{R}}$. Recall that $\Psi_{\mathcal{R}} = \Psi_r \cup \Psi_p \cup \Psi_s \cup \Psi_b$. Among them, Ψ_r, Ψ_p are commonly seen in related work [3], [4], [15] and [27], but Ψ_s, Ψ_b allow us to express richer requirements on the robot's internal states and actions directly, for example where these actions are desired and the preferred sequence. We now state the problem we consider:

Problem 1: Given the weighted finite transition system \mathcal{R} and an LTL formula φ over $\Psi_{\mathcal{R}}$, construct a discrete motion and action plan such that φ is satisfied and also the hybrid controller that implements this discrete plan.

VI. MOTION AND ACTION PLANNER

We provide the solution to Problem 1 in this section. First an optimal motion-and-action plan is derived by searching for the optimal accepting run in the product automaton [2], [6]. Then the hybrid controller that implements this plan is synthesized in an automated manner.

A. Model Checking with Optimality

An infinite path fragment $\tau_{\mathcal{R}}$ of \mathcal{R} is an infinite sequence of states $\pi_{\mathcal{R},0}\pi_{\mathcal{R},1}\pi_{\mathcal{R},2}\pi_{\mathcal{R},3}\dots$, where $\pi_{\mathcal{R},0} \in \Pi_{\mathcal{R},0}$ and $(\pi_{\mathcal{R},i}, \pi_{\mathcal{R},i+1}) \in \rightarrow_{\mathcal{R}}, \forall i > 0$. Its trace $\text{trace}(\tau_{\mathcal{R}})$ defines a *word* of \mathcal{R} , given by the sequence of atomic propositions that are true in the states along this path. Namely, $\text{trace}(\tau_{\mathcal{R}}) = L_{\mathcal{R}}(\pi_{\mathcal{R},0})L_{\mathcal{R}}(\pi_{\mathcal{R},1})\dots$. The NBA $\mathcal{A}_{\varphi} = (Q, 2^{AP}, \delta, Q_0, \mathcal{F})$ from Section II-B allows us to check whether $\tau_{\mathcal{R}}$ satisfies φ by checking if $\text{trace}(\tau_{\mathcal{R}})$ is accepted by \mathcal{A}_{φ} . We use the automaton-based model-checking approach, see [6] and Algorithm 11 in [2]. The product Büchi automaton is defined as a tuple

$$\mathcal{A}_{\mathcal{P}} = \mathcal{R} \otimes \mathcal{A}_{\varphi} = (Q_{\mathcal{P}}, \delta_{\mathcal{P}}, Q_{\mathcal{P},0}, \mathcal{F}_{\mathcal{P}}, W_{\mathcal{P}}), \quad (6)$$

which consists of (i) $Q_{\mathcal{P}} = \Pi_{\mathcal{R}} \times Q$; (ii) the transition relation $(\langle \pi_{\mathcal{R}}, q \rangle, \langle \pi'_{\mathcal{R}}, q' \rangle) \in \delta_{\mathcal{P}}$ iff $(\pi_{\mathcal{R}}, \pi'_{\mathcal{R}}) \in \rightarrow_{\mathcal{R}}$ and $(q, L_{\mathcal{R}}(\pi_{\mathcal{R}}), q') \in \delta$; (iii) the set of initial states $Q_{\mathcal{P},0} = \Pi_{\mathcal{R},0} \times Q_0$; (iv) the set of accepting states $\mathcal{F}_{\mathcal{P}} = \Pi_{\mathcal{R}} \times \mathcal{F}$; (v) $W_{\mathcal{P}} : \delta_{\mathcal{P}} \rightarrow \mathbb{R}^+$, $W_{\mathcal{P}}(\langle \pi_{\mathcal{R}}, q \rangle, \langle \pi'_{\mathcal{R}}, q' \rangle) = W_{\mathcal{R}}(\pi_{\mathcal{R}}, \pi'_{\mathcal{R}})$.

It is proven in [2] that there exists an infinite path of \mathcal{R} satisfying φ if and only if $\mathcal{A}_{\mathcal{P}}$ has at least one accepting run. Then this accepting run could be projected to an infinite path in \mathcal{R} , the trace of which should satisfy φ automatically. In this paper, we consider the accepting runs with the following *prefix-suffix* structure: $r_{\mathcal{P}} = p_0 p_1 \dots (p_k \dots p_n p_k)^{\omega}$, where $p_0 \in Q_{\mathcal{P},0}$ and $p_k \in \mathcal{F}_{\mathcal{P}}$. Namely, $r_{\mathcal{P}}$ consists of two parts: the prefix part that is executed only once from an initial state p_0 to one accepting state p_k and the suffix part which is repeated infinitely from this accepting state back to itself [2], [27]. An accepting run with the prefix-suffix structure has a finite representation, and more importantly it allows us to define the *prefix-suffix cost* of an accepting run:

$$\begin{aligned} \text{Cost}(r_{\mathcal{P}}) &= \left(\sum_{i=0}^{k-1} W_{\mathcal{P}}(p_i, p_{i+1}) \right) \\ &+ \gamma \left(\sum_{i=k}^{n-1} W_{\mathcal{P}}(p_i, p_{i+1}) + W_{\mathcal{P}}(p_n, p_k) \right), \end{aligned} \quad (7)$$

of which the first term is the sum over the weights of transitions along the prefix and the second summation and steady response to the task specification [27].

Remark 5: The prefix-suffix structure is more of a way to formulate the total cost (7), rather than a conservative assumption. If an accepting run exists, by definition at least one accepting state should appear in it infinitely often. Among all the finite number of cycles starting for this accepting state and back to itself there is one with the minimal cost. Thus an accepting run of the prefix-suffix form can be built using this minimal cycle as the periodic suffix.

Algorithm 1: Function $\text{optRun}(G, I, F)$

Input: a weighted graph G, I, F .

Output: the optimal accepting run $r_{\mathcal{P}, \text{opt}}$.

1. Compute the path with minimal cost from every initial vertex in I to every accepting vertex in F .

$(D_{IF}, P_{IF}) = \text{MinPath}(G, I, F)$.

2. Compute the path with minimal cost from every accepting vertex in F and back to itself:

$(D_{FF}, P_{FF}) = \text{MinCycl}(G, F)$.

3. For each column of D_{IF} , find the element with the minimal value and the corresponding cell in P_{IF} (with the same index). Save them sequentially in $1 \times M$ matrix D_{iF} and $1 \times M$ cell P_{iF} .

4. Find the element with the minimal value in $D_{iF} + \gamma D_{FF}$ and its index f_{\min} .

5. Optimal accepting run $r_{\mathcal{P}, \text{opt}}$, prefix: the f_{\min} -th element of P_{iF} ; suffix: the f_{\min} -th element P_{FF} .

The optimal accepting run $r_{\mathcal{P}, \text{opt}}$ can be found by Algorithm 1, which takes as inputs the weighted state graph [2] $G(\mathcal{A}_{\mathcal{P}})$, the set of initial vertices $I = Q_{\mathcal{P}, 0}$ and the set of accepting vertices $F = \mathcal{F}_{\mathcal{P}}$. It utilizes Dijkstra's algorithm [19] for computing the shortest path between pairs of vertices. In particular, denote the number of elements in I and F by $|I| = L$ and $|F| = M$. Function MinPath takes the same inputs and outputs a $L \times M$ matrix D_{IF} , with the $(i_{\text{th}}, j_{\text{th}})$ element containing the value of the minimal cost from I_i to F_j ; and a $L \times M$ cell P_{IF} , with the $(i_{\text{th}}, j_{\text{th}})$ cell containing the sequence of vertices appearing in that minimal path. Function MinCycl is a variant of function MinPath , which outputs a $1 \times M$ matrix D_{FF} , with the j_{th} element containing the value of the minimal cost from F_j back to F_j ; and a $1 \times M$ cell P_{FF} with the j_{th} cell containing the sequence of vertices appearing in that minimal path.

B. Hybrid Controller Synthesis

The optimal accepting run $r_{\mathcal{P}, \text{opt}}$ from Algorithm 1 can be projected to an infinite path $\tau_{\mathcal{R}}$ of \mathcal{R} by projecting $r_{\mathcal{P}}$ onto \mathcal{R} . The trace of $\tau_{\mathcal{R}}$ then automatically satisfies φ . $\tau_{\mathcal{R}}$ also fulfills the prefix-suffix structure [26], which gives a finite representation. Thus the underlying low-level control strategy can be synthesized by sequentially implementing the continuous controller associated with the actions along $\tau_{\mathcal{R}}$. In particular, if $\alpha_{\mathcal{R}} = \text{act}_{\mathcal{B}, k}$, the controller $\{\mathcal{K}_k\}$ that implements $\text{act}_{\mathcal{B}, k}$ is activated. If $\alpha_{\mathcal{R}} = \text{act}_{\mathcal{M}}$, the NF-based controller (2) is applied to drive the robot from one point in the starting region to one point in the goal region.

C. Complexity and Overall Framework

The correctness of the proposed solutions in Section VI follows from the problem formulation and the correctness of the Dijkstra's shortest path algorithm. Let $|\mathcal{M}|$ and $|\mathcal{B}|$ denote the size of the robot's mobility model and action map. The size of $\mathcal{A}_{\mathcal{P}}$ by (6) is $|\mathcal{A}_{\mathcal{P}}| = |\mathcal{M}| \cdot |\mathcal{B}| \cdot 2^{|\varphi|}$. Algorithm 1 runs in $\mathcal{O}(|\mathcal{A}_{\mathcal{P}}| \cdot \log |\mathcal{A}_{\mathcal{P}}| \cdot |Q_{\mathcal{P}, 0}| \cdot |\mathcal{F}_{\mathcal{P}}|)$.

Algorithm 2: Robot motion and action planning under LTL specifications

1. Construct the abstraction of robot mobility \mathcal{M} .

2. Construct the action map \mathcal{B} .

3. Compose \mathcal{M} and \mathcal{B} : $\mathcal{R} = \mathcal{M} ||| \mathcal{B}$.

4. Given a LTL task specification φ over $\Psi_{\mathcal{R}}$, construct its associated NBA \mathcal{A}_{φ} .

5. Compute the product automaton $\mathcal{A}_{\mathcal{P}} = \mathcal{R} \otimes \mathcal{A}_{\varphi}$ and the associated directed graph $G(\mathcal{A}_{\mathcal{P}}) = (Q_{\mathcal{P}}, \delta_{\mathcal{P}}, W_{\mathcal{P}})$.

6. Call Algorithm 1 to derive the optimal accepting run and project it to \mathcal{R} yielding $\tau_{\mathcal{R}}$.

7. Synthesize the hybrid controller as in Section VI-B.

To summarize, the overall framework is shown in Algorithm 2. It is worth mentioning that \mathcal{M} , \mathcal{B} are constructed only once for the robot within a certain workspace and φ can express any task specification in terms of required motions and actions. Steps 3, 5, 6 and 7 are performed automatically [2], [4]. Whenever a new task specification is given, \mathcal{R} remains unchanged and steps 5, 6, 7 are repeated to synthesize the corresponding plan. Whenever the workspace is modified, only \mathcal{M} needs to be revised but \mathcal{B} remains the same and is reused in the following procedures.

VII. CASE STUDY

In the case study, we consider an autonomous robot that repetitively delivers various products from a source region to destination regions while at the same time avoids the prohibited areas and surveils over certain regions. All simulations are carried out in MATLAB on a desktop computer.

A. System Model

We take into account a 2-D workspace for better visualization of the results. The workspace is bounded by region 0: $\pi_0 = \mathcal{B}_1([0.5 \ 0.5]^T)$, where $[0.5 \ 0.5]^T$ is the center point and 1 is the radius. Within π_0 there exist five sphere regions of interest: region 1: $\pi_1 = \mathcal{B}_{0.1}([0 \ 0]^T)$, region 2: $\pi_2 = \mathcal{B}_{0.1}([1 \ 0]^T)$, region 3: $\pi_3 = \mathcal{B}_{0.1}([1 \ 1]^T)$, region 4: $\pi_4 = \mathcal{B}_{0.1}([0 \ 1]^T)$, region 5: $\pi_5 = \mathcal{B}_{0.15}([0.5 \ 0.5]^T)$. $\Psi_r = \{\Psi_{r,i}\}$ reflects the robot's position, $i = 1, \dots, 5$. There are three properties of concern: $\Psi_{p,1} = \{\text{this region has product A}\}$, $\Psi_{p,2} = \{\text{this region has product B}\}$, $\Psi_{p,3} = \{\text{this region is a office area}\}$. It is assumed that region 1 has product A and B and region 5 is a office area. The robot starts from region 1. Then \mathcal{M} is constructed by (3).

The robot is capable of five actions: $\text{act}_{\mathcal{B},1} = \{\text{pickup A}\}$, $\text{act}_{\mathcal{B},2} = \{\text{drop A}\}$, $\text{act}_{\mathcal{B},3} = \{\text{pickup B}\}$, $\text{act}_{\mathcal{B},4} = \{\text{drop B}\}$, $\text{act}_{\mathcal{B},5} = \{\text{take pictures}\}$, and $\text{act}_{\mathcal{B},0} = \{\text{None}\}$. The associated costs are 20, 20, 20, 20, 15, 5 respectively. Ψ_b indicates which action is performed. Two propositions reflecting the robot's internal states are given by $\Psi_{s,1} = \{\text{the robot has A}\}$, $\Psi_{s,2} = \{\text{the robot has B}\}$. The effect and condition functions are listed in Table I after Section IV-B. Assume that the robot initially has no product A or B. The resulting action map \mathcal{B} is then constructed by (4), which has $6 \times 2^2 = 24$ states. We omit the digrams of \mathcal{M} and \mathcal{B} here due to limited space.

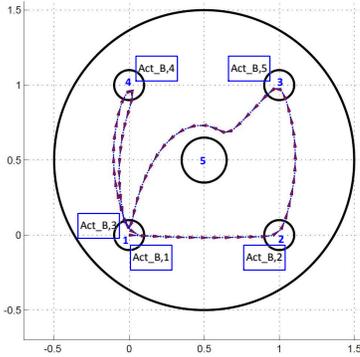


Fig. 2. The final action and motion trajectories fulfills the task specification. The robot stays within the workspace π_0 during the whole mission. Inside the blue boxes are the actions to perform at different regions.

B. Task Specification

The given task is to repeatedly transport product A from region 1 to 2 and product B from region 1 to 4, while region 3 need to be under surveillance. Moreover, all office areas should be avoided during the mission. Related the propositions we have defined, the task is reinterpreted as: Infinitely often, drop A in region 2, drop B in region 4, take pictures within region 3. Always, avoid office areas. This task can be expressed in LTL formula as $\varphi = \square\Diamond(\Psi_{r,2} \wedge \Psi_{b,2}) \wedge \square\Diamond(\Psi_{r,4} \wedge \Psi_{b,4}) \wedge \square\Diamond(\Psi_{r,3} \wedge \Psi_{b,5}) \wedge \square(\neg\Psi_{p,3})$. The NBA \mathcal{A}_φ is obtained from [12], with 4 states and 13 transitions. As can be seen here, the size of \mathcal{A}_φ is relatively small even though the specification is quite complex. Note there is no need to specify where to pickup A and B.

C. Controller Synthesis

By Algorithm 2, the composition $\mathcal{R} = \mathcal{M} ||| \mathcal{B}$ is constructed, which has 90 states and 606 transitions (compared with 2^{15} states when unfolding $\Psi_{\mathcal{R}}$ blindly). The product automaton $\mathcal{A}_{\mathcal{P}} = \mathcal{R} \otimes \mathcal{A}_\varphi$ is given by (6), which has 480 states, out of which 120 are accepting states. Then Algorithm 1 is applied to find the optimal accepting path. The final discrete plan is obtained by projecting this accepting run onto \mathcal{R} , which is interpreted in terms of the following action sequence: pickup A in region 1 \rightarrow move to region 2 \rightarrow drop A \rightarrow move to region 3 \rightarrow take pictures \rightarrow move to region 1 \rightarrow pickup B \rightarrow move to region 4 \rightarrow drop B \rightarrow move to region 1. Note that this sequence is cyclic and can be repeated as many times as needed. The total cost of this discrete plan is 230. The corresponding hybrid controller is synthesized. In Figure 2, the final trajectories are shown by the red arrowed lines and the actions performed during the motion are indicated by action names in the blue boxes.

VIII. CONCLUSION

We presented a systematic way to synthesize a hybrid control strategy for motion and action planning of autonomous robots under LTL task specifications. The specifications take into account not only a sequence of regions to visit, but also the desired actions at the regions. Further research could involve reactive environments and multi-robot systems.

REFERENCES

- [1] A. Bhatia, M. R. Maly, L. E. Kavraki, M. Y. Vardi. Motion planning with complex goals. *IEEE Robotics & Automation Magazine*, 18(3): 55-64, 2011.
- [2] C. Baier, J.-P. Katoen. Principles of model checking. *The MIT Press*, 2008.
- [3] C. Belta, V. Isler, G. J. Pappas. Discrete abstractions for robot motion planning and control in polygonal environments. *IEEE Transactions on Robotics*, 21(5): 864-874, 2005.
- [4] C. Belta, A. Bicchi, M. Egerstedt, E. Frazzoli, E. Klavins, G. J. Pappas. Symbolic planning and control of robot motion. *IEEE Robotics and Automation Magazine*, 14: 61-71, 2007.
- [5] T. Bylander. The Computational Complexity of Propositional STRIPS Planning. *Artificial Intelligence*, 69(2): 165204, 1994.
- [6] E. M. Clarke, O. Grumberg, D. A. Peled. Model checking. *The MIT Press*, 1999.
- [7] X. Ding, M. Kloetzer, Y. Chen, C. Belta. Automatic deployment of robotic teams. *IEEE Robotics Automation Magazine*, 18: 75-86, 2011.
- [8] D. V. Dimarogonas and K. J. Kyriakopoulos. Decentralized Navigation Functions for Multiple Robotic Agents with Limited Sensing Capabilities. *Journal of Intelligent and Robotic Systems*. 48(3): 411-433, 2007.
- [9] G. E. Fainekos, A. Girard, H. Kress-Gazit, G. J. Pappas. Temporal Logic Motion Planning for Dynamic Mobile Robots. *Automatica*, 45(2): 343-352, 2009.
- [10] I. F. Filippidis, D. V. Dimarogonas, K. J. Kyriakopoulos. Decentralized Multi-Agent Control from Local LTL Specifications. *IEEE Conference on Decision and Control*, 2012.
- [11] R. E. Fikes, Nils J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial intelligence*, 2(3): 189-208, 1972.
- [12] P. Gastin, D. Oddoux. Fast LTL to Büchi automaton translation. *In Proceedings of the 13th International Conference on Computer Aided Verification (CAV'01)*, 2001.
- [13] M. Guo, K. H. Johansson, D. V. Dimarogonas. Revising Motion Planning under Linear Temporal Logic Specifications in Partially Known Workspaces. *IEEE International Conference on Robotics and Automation*, 2013.
- [14] S. Karaman, E. Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 30(7): 846-894, 2011.
- [15] M. Kloetzer, C. Belta. Automatic deployment of distributed teams of robots from temporal logic specifications. *IEEE Transactions on Robotics*, 26(1): 48-61, 2010.
- [16] D. E. Koditschek, E. Rimon. Robot navigation functions on manifolds with boundary. *Advances Appl. Math.*, 11:412-442, 1990.
- [17] H. Kress-Gazit, T. Wongpiromsarn, U. Topcu. Correct, reactive robot control from abstraction and temporal logic specifications. *IEEE Robotics and Automation Magazine*, 2011.
- [18] H. Kress-Gazit, G. E. Fainekos, G. J. Pappas. Where's Waldo? Sensor-based temporal logic motion planning. *IEEE International Conference on Robotics and Automation*, 2007.
- [19] S. M. LaValle. Planning Algorithms. *Cambridge University Press*, 2006.
- [20] S. G. Loizou, K. J. Kyriakopoulos. Closed loop navigation for multiple non-holonomic vehicles. *IEEE International Conference on Robotics and Automation*, 3: 4240-4245, 2003.
- [21] S. G. Loizou, A. Jadbabaie. Density Functions for Navigation Function Based Systems. *IEEE Conference on Decision and Control*, 2006.
- [22] M. Ghallab, C. Aeronautiques, C. K. Isi, D. Wilkins. PDDL: The Planning Domain Definition Language. *Tech. report CVC TR98003/DCS TR1165*, Yale Center for Computational Vision and Control, 1998.
- [23] M. Gelfond, V. Lifschitz. Action languages. *Transactions on AI*, 1998.
- [24] A. S. Oikonomopoulos, S. G. Loizou, K. J. Kyriakopoulos. Coordination of multiple non-holonomic agents with input constraints. *IEEE International Conference on Robotics and Automation*, 869-874, 2009.
- [25] T. Wongpiromsarn, U. Topcu, R. Murray. Receding horizon temporal logic planning. *IEEE Transactions on Automatic Control*, 2012.
- [26] A. Ulusoy, S. L. Smith, C. Belta. Optimal Multi-robot path planning with LTL constraints: guaranteeing correctness through synchronization. *International Symposium on Distributed Autonomous Robotic Systems*, 2012.
- [27] S. L. Smith, J. Tumova, C. Belta, D. Rus. Optimal path planning for surveillance with temporal logic constraints. *International Journal of Robotics Research*, 30(14): 1695-1708, 2011.