

Receding Horizon Temporal Logic Planning

Tichakorn Wongpiromsarn, Ufuk Topcu, and Richard M. Murray

Abstract—We present a methodology for automatic synthesis of embedded control software that incorporates a class of linear temporal logic (LTL) specifications sufficient to describe a wide range of properties including safety, stability, progress, obligation, response and guarantee. To alleviate the associated computational complexity of LTL synthesis, we propose a receding horizon framework that effectively reduces the synthesis problem into a set of smaller problems. The proposed control structure consists of a goal generator, a trajectory planner, and a continuous controller. The goal generator reduces the trajectory generation problem into a sequence of smaller problems of short horizon while preserving the desired system-level temporal properties. Subsequently, in each iteration, the trajectory planner solves the corresponding short-horizon problem with the currently observed state as the initial state and generates a feasible trajectory to be implemented by the continuous controller. Based on the simulation property, we show that the composition of the goal generator, trajectory planner and continuous controller and the corresponding receding horizon framework guarantee the correctness of the system with respect to its specification regardless of the environment in which the system operates. In addition, we present a response mechanism to handle failures that may occur due to a mismatch between the actual system and its model. The effectiveness of the proposed technique is demonstrated through an example of an autonomous vehicle navigating an urban environment. This example also illustrates that the system is not only robust with respect to exogenous disturbances but is also capable of properly handling violation of the environment assumption that is explicitly stated as part of the system specification.

Index Terms—Autonomous systems, control architecture, linear temporal logic, receding horizon control.

I. INTRODUCTION

Design and verification of modern engineered systems with a tight link between computational and physical elements have become increasingly complex due to the interleaving between the high-level logics and the low-level dynamics. Consider, for example, an autonomous driving problem, particularly the 2007 DARPA Urban Challenge. In this competition, the competing vehicles had to navigate, in a fully autonomous manner, through a partially known urban-like environment populated with static and dynamic obstacles and perform different tasks such as road and off-road driving, parking and visiting certain areas while obeying traffic rules. For the vehicles to successfully complete the race, they need to be capable of negotiating an intersection, handling changes in the environment or operating condition and replanning in response to such changes. Hence, the high-level logic that governs the behavior of the vehicles needs to be properly integrated with the low-level controller that regulates the physical hardware.

A schematic of the hierarchical protocol stack implemented in Team Caltech’s entry in the challenge is shown in Fig. 1. The stack comprises multiple software modules that are responsible for reasoning at different levels of abstraction [1], [2], [3]. Path Follower, for example, computes control signals based on the continuous model of the vehicle such that the vehicle closely follows the path generated by Path Planner. Mission Planner generates a route to complete the mission based only on the discrete model of the connectivity between different road segments. Finally, Traffic Planner, implemented as a set of finite state machines, determines how the vehicle should navigate this route incorporating the traffic rules. This paper particularly focuses on the traffic planner and path follower layers of such protocol stacks.

Protocol stacks on autonomous vehicles constitute an example of a broader class of systems, namely embedded control systems that incorporate continuous and discrete decision making and interact with the (potentially dynamic and a priori unknown) environment. These systems are typically very complex, yet a lot of them are still designed and implemented in an ad-hoc manner. Even though the individual components may be formally verified, the whole system is typically verified only through simulations and tests. Hence, there is no formal guarantee that the system would work as desired. In fact, a design bug in Team Caltech’s entry in the Urban Challenge related to a mismatch in the abstraction of the physical system used at different levels of the hierarchy had never been discovered in thousands of hours of our extensive simulations and over three hundred miles of field testing [4]. In addition, once the bug was uncovered, it was difficult to modify and verify the design due to the complexity of the system and the lack of sufficient time. Although it might be impractical to simplify such a system, part of the complexity could be avoided if the system had been designed in a systematic way.

Motivated by the difficulty of modifying and verifying complex embedded control systems, in this paper, we investigate the following control protocol synthesis problem.

Problem Description: Given a model for the system and its specification expressed in a formal language, synthesize a control protocol that, by construction, ensures that the system satisfies its specification for all valid environment behaviors.

In particular, we consider discrete-time linear time-invariant system and use linear temporal logic [5], [6], [7] as the specification language. Environment refers to the factors over which the system has no control such as obstacles, weather condition, software and hardware faults and failures, etc. We assume that the system respects its model and ensure that an execution described by the model, rather than the actual execution of the system, satisfies the specification. (Checking that the model accurately describes the actual system is a validation problem which is not in the scope of the paper.)

T. Wongpiromsarn is with the Singapore-MIT Alliance for Research and Technology, Singapore 117543, Singapore. nok@smart.mit.edu

U. Topcu and R. M. Murray are with the California Institute of Technology, Pasadena, CA 91125, USA. {utopcu, murray}@cds.caltech.edu

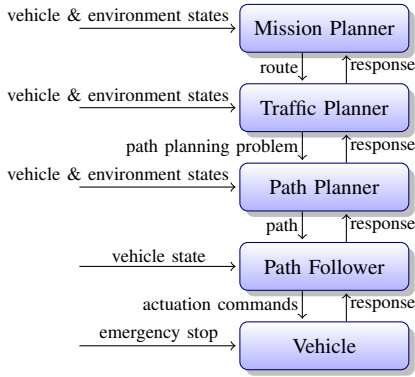


Fig. 1. A hierarchical protocol stack.

For simplicity, we refer to a model of the system as “system” for the rest of the paper. The key definitions and notations are provided in Section II and the previously described synthesis problem is formally formulated in Section III.

A Solution Approach and Associated Issues: A common approach to the above synthesis problem is to construct a finite transition system that serves as an abstract model of the physical system (which typically has infinitely many states) [8], [9], [10], [11], [12], [13], [14], [15], [16]. Then based on this abstract model, synthesize a strategy, represented by a finite state automaton, satisfying the specification. This leads to a hierarchical, two-layer design with a discrete planner computing a discrete plan based on the abstract model and a continuous controller computing control signal based on the physical model to continuously implement the discrete plan. Simulations/bisimulations [17] provide a proof that the continuous execution preserves the correctness of the discrete plan. We describe this hierarchical approach in detail in Section IV.

One of the main challenges of this hierarchical approach is in the abstraction of continuous, infinite-state systems into equivalent (in the simulation sense) finite state models. Several abstraction methods have been proposed based on a fixed abstraction. For example, a continuous-time, time-invariant model was considered in [9], [10] and [11] for special cases of fully actuated ($\dot{s}(t) = u(t)$), kinematic ($\dot{s}(t) = A(s(t))u(t)$) and piecewise affine (PWA) dynamics, respectively. A discrete-time, time-invariant model was considered in [14] and [12] for special cases of PWA and controllable linear systems, respectively. Reference [13] deals with more general dynamics by relaxing the bisimulation requirement and using the notion of approximate simulation [18]. More recently, a sampling-based method has been proposed for μ -calculus specifications [8].

Another issue is the computational complexity in the synthesis of finite state automata from a temporal logic specification, especially in the presence of the dynamic and a priori unknown environment. Piterman et al. [19] treated this problem as a two-player game between the system and the environment and proposed an algorithm for the synthesis of a finite state automaton that satisfies its specification for any environment behavior. Although for a certain class of properties, known as *Generalized Reactivity*[1], such an automaton can be computed in polynomial time, for practical problems, the rapid

increase in computational complexity is still a limiting factor. **Contributions of the paper:** This paper concerns both the abstraction and the computational complexity issues. First, in Section V, we propose an approach to automatically compute a finite state abstraction for a discrete-time linear time-invariant system, taking into account exogenous disturbances. Our approach differs from those presented in [9], [10], [11], [12], [13], [14] as we explicitly take into account the presence of exogenous disturbances. Although exogenous disturbances have been considered in [20], the state space is partitioned based on a fixed geometric shape (e.g., a hypercube of fixed size). Hence, the size of the resulting state space may be unnecessarily large. To alleviate this shortcoming, we start with a predicate-based partition and refine it based on the reachability relation to polytopic regions of different size and shape. This allows us to control the size of the state space and identify the set of initial states starting from which a control law that ensures the satisfaction of the desired properties cannot be found.

Second, in Section VI, we propose a receding horizon framework for executing finite state automata while ensuring system correctness with respect to a given linear temporal logic specification. This framework essentially reduces the synthesis problem into a set of smaller problems of shorter horizon. It relies on the partial order relation among the discrete states. The partial order plays a role similar to the contraction constraints [21], [22], which practically induce an order in the state space, in receding horizon control literature. The implementation of the receding horizon framework, presented in Section VII, leads to the decomposition of the discrete planner into a goal generator and a trajectory planner. The goal generator reduces the synthesis problem to a sequence of short-horizon problems while preserving the desired system-level temporal properties. Subsequently, in each iteration, the trajectory planner solves the corresponding short-horizon problem with the currently observed state as the initial state and generates a feasible trajectory to be implemented by the continuous controller.

Finally, we present a response mechanism that potentially increases the robustness of the system with respect to a mismatch between the actual system and its model and violation of the environment assumptions. The proposed technique is demonstrated through an example of an autonomous vehicle navigating an urban-like environment. This example also illustrates that the resulting hierarchical control structure is not only robust with respect to exogenous disturbances but also capable of handling violation of the environment assumptions.

Preliminary versions of this work have partially appeared in [14], [15], [16].

II. PRELIMINARIES

We use linear temporal logic (LTL) to specify properties of systems. In this section, we first give formal definitions of terminology and notations used throughout the paper. Then, based on these definitions, we briefly describe LTL and some important classes of LTL formulas.

Definition 1. A system consists of a set V of variables. The

domain of V , denoted by $dom(V)$, is the set of valuations of V . A *state* of the system is an element $v \in dom(V)$.

We describe an *execution* of a system by an infinite sequence of its states. Specifically, for a discrete-time system, its execution σ can be written as $\sigma = v_0 v_1 v_2 \dots$ where $v_t \in dom(V)$ is the state of the system at time t .

Definition 2. A *finite transition system* is a tuple $\mathbb{T} := (\mathcal{V}, \mathcal{V}_0, \rightarrow)$ where \mathcal{V} is a finite set of states, $\mathcal{V}_0 \subseteq \mathcal{V}$ is a set of initial states, and $\rightarrow \subseteq \mathcal{V} \times \mathcal{V}$ is a transition relation. Given states $\nu_i, \nu_j \in \mathcal{V}$, we write $\nu_i \rightarrow \nu_j$ if there is a transition from ν_i to ν_j in \mathbb{T} .

In this paper, we use ν to represent a state of a finite transition system and v to represent a state of a general, possibly non-finite state system.

Definition 3. A *partially ordered set* (V, \leq) consists of a set V and a binary relation \leq over the set V satisfying the following properties: for any $v_1, v_2, v_3 \in V$, (a) $v_1 \leq v_1$; (b) if $v_1 \leq v_2$ and $v_2 \leq v_1$, then $v_1 = v_2$; and (c) if $v_1 \leq v_2$ and $v_2 \leq v_3$, then $v_1 \leq v_3$.

Definition 4. An *atomic proposition* is a statement on system variables v that has a unique truth value (*True* or *False*) for a given value of v . Let $v \in dom(V)$ be a state of the system and p be an atomic proposition. We write $v \models p$ if p is *True* at the state v . Otherwise, we write $v \not\models p$.

LTL is a powerful specification language for unambiguously and concisely expressing a wide range of properties of systems [5], [6], [7]. It is built up from (a) a set of atomic propositions, (b) the logic connectives: negation (\neg), disjunction (\vee), conjunction (\wedge) and material implication (\implies), and (c) the temporal modal operators: next (\circ), always (\square), eventually (\diamond) and until (\mathcal{U}). An LTL formula is defined inductively as follows: (1) any atomic proposition p is an LTL formula; and (2) given LTL formulas φ and ψ , $\neg\varphi$, $\varphi \vee \psi$, $\circ\varphi$ and $\varphi \mathcal{U} \psi$ are also LTL formulas. Other operators can be defined as follows: $\varphi \wedge \psi \triangleq \neg(\neg\varphi \vee \neg\psi)$, $\varphi \implies \psi \triangleq \neg\varphi \vee \psi$, $\diamond\varphi \triangleq \text{True } \mathcal{U} \varphi$, and $\square\varphi \triangleq \neg\diamond\neg\varphi$. A *propositional formula* is one that does not include temporal operators. Given a set of LTL formulas $\varphi_1, \dots, \varphi_n$, their *Boolean combination* is an LTL formula formed by joining $\varphi_1, \dots, \varphi_n$ with logic connectives.

Semantics of LTL: An LTL formula is interpreted over an infinite sequence of states. Given an execution $\sigma = v_0 v_1 v_2 \dots$ and an LTL formula φ , we write $v_i \models \varphi$ if φ holds at position $i \geq 0$ of σ . The semantics of LTL is defined inductively as follows: (a) For an atomic proposition p , $v_i \models p$ if and only if (iff) $v_i \models p$; (b) $v_i \models \neg\varphi$ iff $v_i \not\models \varphi$; (c) $v_i \models \varphi \vee \psi$ iff $v_i \models \varphi$ or $v_i \models \psi$; (d) $v_i \models \circ\varphi$ iff $v_{i+1} \models \varphi$; and (e) $v_i \models \varphi \mathcal{U} \psi$ iff there exists $j \geq i$ such that $v_j \models \psi$ and $\forall k \in [i, j), v_k \models \varphi$. Based on this definition, $\circ\varphi$ holds at position v_i iff φ holds at the next state v_{i+1} , $\square\varphi$ holds at position i iff φ holds at every position in σ starting at position i , and $\diamond\varphi$ holds at position i iff φ holds at some position $j \geq i$ in σ .

Definition 5. An execution $\sigma = v_0 v_1 v_2 \dots$ *satisfies* φ , denoted by $\sigma \models \varphi$, if $v_0 \models \varphi$.

Definition 6. Let Σ be the set of all executions of a system. The system is said to be *correct* with respect to the specification φ , written $\Sigma \models \varphi$, if all its executions satisfy φ , that is, $(\Sigma \models \varphi)$ iff $(\forall \sigma, (\sigma \in \Sigma) \implies (\sigma \models \varphi))$.

Remark 1. Properties typically studied in the control and hybrid systems domains are safety (usually in the form of constraints on the system state) and stability (i.e., convergence to an equilibrium or a desired state). However, these properties are not rich enough to describe certain desired properties of, for example, an autonomous vehicle such as staying in the travel lane unless there is an obstacle blocking the lane, and visiting a certain area infinitely often. In Section VIII, we show how such properties can be easily expressed in LTL. Examples of widely-used LTL formulas include safety, guarantee, obligation, progress, response and stability.

III. PROBLEM FORMULATION

We consider a system that comprises the physical component, which we refer to as the plant, and the (potentially dynamic and a priori unknown) environment in which the plant operates. Specifically, we define the system \mathbb{S} and the specification φ as follows.

System: Consider a system \mathbb{S} with a set $V = S \cup E$ of variables where S and E are disjoint sets that represent, respectively, the set of plant variables that are regulated by the control protocol and the set of environment variables whose values may change arbitrarily throughout an execution. The domain of V is given by $dom(V) = dom(S) \times dom(E)$ and a state of the system can be written as $v = (s, e)$ where $s \in dom(S) \subseteq \mathbb{R}^n$ and $e \in dom(E)$. We call s the *controlled state* and e the *environment state*.

The controlled state evolves according to the following discrete-time linear time-invariant state space model: for $t \in \{0, 1, 2, \dots\}$,

$$\begin{aligned} s[t+1] &= As[t] + B_u u[t] + B_d d[t], \\ u[t] &\in U, d[t] \in D, s[0] \in dom(S), \end{aligned} \quad (1)$$

where $U \subseteq \mathbb{R}^m$ is the set of admissible control inputs, $D \subseteq \mathbb{R}^p$ is the set of exogenous disturbances and $s[t]$, $u[t]$ and $d[t]$ are the controlled state, control signal and disturbance at time t .

Example 1. Take, for example, a scenario where a robot needs to navigate an environment populated with (potentially dynamic) obstacles and explore certain areas of interest. S typically includes the state (e.g. position and velocity) of the robot while E typically includes the positions of obstacles (which are generally not known a priori and may change over time). The evolution of the controlled state (i.e., the state of the robot) is governed by its equations of motion.

System Specification: We assume that the specification φ is of the form

$$\varphi = (\varphi_{init} \wedge \varphi_e) \implies \varphi_s. \quad (2)$$

where φ_{init} specifies system's initial conditions, φ_e describes the knowledge about the allowable environment behavior and the desired behavior of the systems is encoded in φ_s .

Let Π be a finite set of atomic propositions of variables from V . Each of the atomic propositions in Π essentially captures the states of interest. We assume that the desired behavior φ_s is an LTL specification built from Π and can be expressed as a conjunction of safety, guarantee, obligation, progress, response and stability properties as follows

$$\begin{aligned} \varphi_s = & \bigwedge_{j \in J_1} \square p_{1,j}^s \wedge \bigwedge_{j \in J_2} \diamond p_{2,j}^s \wedge \\ & \bigwedge_{j \in J_3} (\square p_{3,j}^s \vee \diamond q_{3,j}^s) \wedge \bigwedge_{j \in J_4} \square \diamond p_{4,j}^s \wedge \\ & \bigwedge_{j \in J_5} \square (p_{5,j}^s \implies \diamond q_{5,j}^s) \wedge \bigwedge_{j \in J_6} \diamond \square p_{6,j}^s, \end{aligned} \quad (3)$$

where J_1, \dots, J_6 are finite sets and for any i and j , $p_{i,j}^s$ and $q_{i,j}^s$ are propositional formulas of variables from V that are built from Π .

Furthermore, we assume that φ_{init} is a propositional formula built from Π and φ_e can be expressed as a conjunction of safety and justice requirements as follows

$$\varphi_e = \bigwedge_{i \in I_1} \square p_{f,i}^e \wedge \bigwedge_{i \in I_2} \square \diamond p_{s,i}^e, \quad (4)$$

where $p_{f,i}^e$ and $p_{s,i}^e$ are propositional formulas built from Π and only contain variables from E . Note that we restrict φ_s and φ_e to be of the form (3) and (4), respectively, for the clarity of presentation. Our framework only requires that the specification (2) can be reduced to the form of equation (6), presented later.

Example 2. Consider the robot motion planning problem described in Example 1. Suppose the workspace of the robot is partitioned into cells C_1, \dots, C_M and the robot needs to visit cells C_1 and C_2 infinitely often. We assume that one of the cells C_1, \dots, C_M may be occupied by an obstacle at any given time. In addition, this obstacle-occupied cell may change arbitrarily throughout an execution but infinitely often, C_1 and C_2 are not occupied. Let s and o represent the position of the robot and the obstacle, respectively. In this case,

$$\begin{aligned} \varphi_s = & \square \diamond (s \in C_1) \wedge \square \diamond (s \in C_2) \wedge \\ & \square ((o \in C_1) \implies (s \notin C_1)) \wedge \\ & \square ((o \in C_2) \implies (s \notin C_2)) \wedge \dots \wedge \\ & \square ((o \in C_M) \implies (s \notin C_M)). \end{aligned}$$

Assuming that initially, the robot does not occupy the same cell as the obstacle,

$$\begin{aligned} \varphi_{init} = & ((o \in C_1) \implies (s \notin C_1)) \wedge ((o \in C_2) \implies (s \notin C_2)) \\ & \wedge \dots \wedge ((o \in C_m) \implies (s \notin C_m)). \end{aligned}$$

Finally, the assumption on the environment can be expressed as $\varphi_e = \square \diamond (o \notin C_1) \wedge \square \diamond (o \notin C_2)$.

Control Protocol Synthesis Problem: Given a system \mathbb{S} and specification φ , synthesize a control protocol that generates a sequence of control signals $u[0], u[1], \dots \in U$ to the plant to ensure that starting from any initial condition, φ is satisfied for any sequence of exogenous disturbances $d[0], d[1], \dots \in D$ and any sequence of environment states.

Note that the control objective is to ensure that the specification φ is satisfied for any initial condition and environment, including those that violate the assumptions φ_{init} and φ_e . However, according to (2), φ holds in any execution where φ_{init} or φ_e is violated. Hence, the desired behavior φ_s only needs to be ensured when φ_{init} and φ_e are satisfied.

IV. HIERARCHICAL APPROACH

We take a hierarchical approach to solve the Control Protocol Synthesis Problem stated in Section III. First, we construct a finite transition system \mathbb{D} that serves as an abstract model of \mathbb{S} . The problem is then decomposed into (a) synthesizing a discrete planner that computes a discrete plan satisfying the specification φ based on the abstract, finite-state model \mathbb{D} , and (b) designing a continuous controller that implements the discrete plan. The success of this abstraction-based approach thus relies on the following two critical steps:

- (a) an abstraction of an infinite-state system into an equivalent (in the simulation sense) finite state model such that any discrete plan generated by the discrete planner can be implemented (i.e., *simulated*; see, for example, [23] for the exact definition of *simulation*) by the continuous controller, provided that the evolution of the controlled state satisfies (1), and
- (b) synthesis of a discrete planner (i.e., a strategy), represented by a finite state automaton, that ensures the correctness of the discrete plan.

In Section V, we present an approach for step (a), assuming that the physical system is modeled as described in Section III. For step (b) and ensure the system correctness for any initial condition and environment behavior, we apply the two-player game approach presented in [19] to synthesize a discrete planner as in [9], [14]. We refer the reader to [19] and references therein for a detailed discussion. In summary, consider a class of LTL formulas of the form

$$\left(\psi_{init} \wedge \square \psi_e \wedge \bigwedge_{i \in I_f} \square \diamond \psi_{f,i} \right) \implies \left(\bigwedge_{i \in I_s} \square \psi_{s,i} \wedge \bigwedge_{i \in I_g} \square \diamond \psi_{g,i} \right), \quad (5)$$

known as *Generalized Reactivity[1]* (GR[1]) formulas. Here, ψ_{init} , $\psi_{f,i}$ and $\psi_{g,i}$ are propositional formulas of variables from V ; ψ_e is a Boolean combination of propositional formulas of variables from V and expressions of the form $\square \psi_e^t$ where ψ_e^t is a propositional formula of variables from E that describes the assumptions on the transitions of environment states; and $\psi_{s,i}$ is a Boolean combination of propositional formulas of variables from V and expressions of the form $\square \psi_s^t$ where ψ_s^t is a propositional formula of variables from V that describes the constraints on the transitions of controlled states. The approach presented in [19] allows checking the realizability of this class of specifications and synthesizing the corresponding finite state automaton to be performed in $O(|\mathcal{V}|^3)$ time where $|\mathcal{V}|$ is the number of states of the finite state abstraction \mathbb{D} .

Proposition 1. A specification of the form (2)–(4) can be reduced to a subclass of GR[1] formula of the form

$$\left(\psi_{init} \wedge \square \psi_e^e \wedge \bigwedge_{i \in I_f} \square \diamond \psi_{f,i}^e \right) \implies \left(\bigwedge_{i \in I_s} \square \psi_{s,i} \wedge \bigwedge_{i \in I_g} \square \diamond \psi_{g,i} \right), \quad (6)$$

where ψ_{init} , $\psi_{s,i}$ and $\psi_{g,i}$ are as defined above and ψ_e^e and $\psi_{f,i}^e$ are propositional formulas of variables from E .

Throughout the paper, we refer to the left hand side and the right hand side of (6) as the “assumption” part and the “guarantee” part, respectively. The proof of Proposition 1 is

based on the fact that all safety, guarantee, obligation and response properties are special cases of progress formulas $\square\Diamond p$, provided that p is allowed to be a past formula [5]. Hence, these properties can be reduced to the guarantee part of (6) by introducing auxiliary Boolean variables. For example, the guarantee property $\Diamond p_{2,j}^s$ can be reduced to the guarantee part of (6) by introducing an auxiliary Boolean variable x , initialized to $p_{2,j}^s$. The formula $\Diamond p_{2,j}^s$ can then be equivalently expressed as a conjunction of $\square((x \vee p_{2,j}^s) \implies \bigcirc x)$, $\square(\neg(x \vee p_{2,j}^s) \implies \bigcirc(\neg x))$ and $\square\Diamond x$. Obligation and response properties can be reduced to the guarantee part of (6) using a similar idea. In addition, the stability property $\Diamond\square p_{6,j}^s$ can be reduced to the guarantee part of (6) by introducing an auxiliary Boolean variable y , initialized to *False*. The formula $\Diamond\square p_{6,j}^s$ can then be equivalently expressed as a conjunction of $\square(y \implies p_{6,j}^s)$, $\square(y \implies \bigcirc y)$, $\square(\neg y \implies (\bigcirc y \vee \bigcirc(\neg y)))$ and $\square\Diamond y$. Note that these reductions lead to equivalent specifications. However, for the case of stability $\Diamond\square p_{6,j}^s$, the reduction may lead to an unrealizable specification even though the original specification is realizable. Roughly speaking, this is because the auxiliary Boolean variable y needs to make clairvoyant (prophecy), non-deterministic decisions. For other properties, the realizability remains the same after the reduction since the synthesis algorithm [19] is capable of handling past formulas. The detail of this discussion is beyond the scope of this paper and we refer the reader to [19] for more detailed discussion on the synthesis of GR[1] specification.

In Section VI, we describe a receding horizon framework that incorporate LTL specification of the form (6) in order to reduce the computational complexity of the synthesis problem. Its implementation and a response mechanism that enables the system to handle certain failures and continue to exhibit a correct behavior are presented in Section VII.

V. COMPUTING FINITE STATE ABSTRACTION

To construct a finite transition system \mathbb{D} from the physical model \mathbb{S} , we first partition $dom(S)$ and $dom(E)$ into finite sets \mathcal{S} and \mathcal{E} , respectively, of equivalence classes or cells such that the partition is *proposition preserving* [17]. Roughly speaking, a partition is said to be proposition preserving if for any atomic proposition $\pi \in \Pi$ and any states v_1 and v_2 that belong to the same cell in the partition, v_1 satisfies π iff v_2 also satisfies π . We denote the resulting discrete domain of the system by $\mathcal{V} = \mathcal{S} \times \mathcal{E}$. We call $v \in dom(\mathcal{V})$ a *continuous state* and $\nu \in \mathcal{V}$ a *discrete state* of the system. For a discrete state $\nu \in \mathcal{V}$, we say that ν satisfies an atomic proposition $\pi \in \Pi$, denoted by $\nu \models_d \pi$, iff there exists a continuous state v contained in the cell labeled by ν such that v satisfies π . (Due to the proposition preserving property of the partition, the existence of such a state $v \models \pi$ implies that all the continuous states contained in the cell labeled by ν satisfies π .) Given an infinite sequence of discrete states $\sigma_d = \nu_0 \nu_1 \nu_2 \dots$ and LTL formula φ built from Π , we write $\nu_i \models_d \varphi$ if φ holds at position $i \geq 0$ of σ_d . With these definitions, the semantics of LTL for a sequence of discrete states can be derived from the general semantics of LTL.

Next, we need to determine the transition relations \rightarrow of \mathbb{D} . In Section V-A, we use a variant of the notion of *reachability* that is sufficient to guarantee that if a discrete controlled state ς_j is reachable from ς_i , the transition from ς_i to ς_j can be continuously *implemented* or *simulated* by a continuous controller. A computational scheme that provides a sufficient condition for reachability between two discrete controlled states and subsequently refines the state space partition is also presented in Sections V-B and V-C.

A. Finite-Time Reachability

Let $\mathcal{S} = \{\varsigma_1, \varsigma_2, \dots, \varsigma_l\}$ be a set of discrete controlled states. We define a map $T_s : dom(S) \rightarrow \mathcal{S}$ that sends a continuous controlled state to a discrete controlled state of its equivalence class. That is, $T_s^{-1}(\varsigma_i) \subseteq dom(S)$ is the set of all the continuous controlled states contained in the cell labeled by ς_i and $\{T_s^{-1}(\varsigma_i), \dots, T_s^{-1}(\varsigma_n)\}$ is the partition of $dom(S)$.

Definition 7. A discrete state ς_j is *reachable* from a discrete state ς_i , written $\varsigma_i \rightsquigarrow \varsigma_j$, only if starting from any point $s[0] \in T_s^{-1}(\varsigma_i)$, there exists a horizon length $N \in \{0, 1, \dots\}$ such that for any sequence of exogenous disturbances $d[0], d[1], \dots, d[N-1] \in D$, there exists a sequence of control signals $u[0], u[1], \dots, u[N-1] \in U$ that takes the system (1) to a point $s[N] \in T_s^{-1}(\varsigma_j)$ satisfying the constraint $s[t] \in T_s^{-1}(\varsigma_i) \cup T_s^{-1}(\varsigma_j)$ for all $t \in \{0, \dots, N\}$. We write $\varsigma_i \not\rightsquigarrow \varsigma_j$ if ς_j is not reachable from ς_i .

In general, for two discrete states ς_i and ς_j , verifying the reachability relation $\varsigma_i \rightsquigarrow \varsigma_j$ is hard because it requires searching for a proper horizon length N . Therefore, we consider the restricted case where the horizon length is fixed and given and U , D and $T_s^{-1}(\varsigma_i), i \in \{1, \dots, l\}$ are polyhedral sets. Our approach relies on solving the following problem.

Reachability Problem: Given an initial continuous controlled state $s[0] \in \mathbb{R}^n$, discrete controlled states $\varsigma_i, \varsigma_j \in \mathcal{S}$, the set of admissible control inputs $U \subseteq \mathbb{R}^m$, the set of exogenous disturbances $D \subseteq \mathbb{R}^p$, the matrices A , B_u and B_d as in (1), a horizon length $N \geq 0$, determine a sequence of control signals $u[0], u[1], \dots, u[N-1] \in \mathbb{R}^m$ such that for all $t \in \{0, \dots, N-1\}$ and $d[t] \in D$,

$$\begin{aligned} s[t+1] &= As[t] + B_u u[t] + B_d d[t], \\ s[t] \in T_s^{-1}(\varsigma_i), u[t] \in U, s[N] \in T_s^{-1}(\varsigma_j). \end{aligned} \quad (7)$$

If the above Reachability Problem is feasible, a solution $u[0], \dots, u[N-1]$ can be computed by formulating a constrained optimal control problem, which can be solved using off-the-shelf software such as MPT [24], YALMIP [25] or NTG [26].

B. Verifying the Reachability Relation

Given two discrete controlled states $\varsigma_i, \varsigma_j \in \mathcal{S}$, to determine whether $\varsigma_i \rightsquigarrow \varsigma_j$, we essentially have to verify that $T_s^{-1}(\varsigma_i) \subseteq S_0$ where S_0 is the set of $s[0]$ starting from which the Reachability Problem defined in Section V-A is feasible. In this section, we describe how S_0 can be computed using an idea from constrained robust optimal control [27].

We assume that U , D and $T_s^{-1}(\varsigma_i), i \in \{1, \dots, l\}$ are polyhedral sets, i.e., there exist matrices L_1, L_2 and L_3 and vectors M_1, M_2 and M_3 such that $T_s^{-1}(\varsigma_i) = \{s \in \mathbb{R}^n \mid L_1 s \leq M_1\}$, $U = \{u \in \mathbb{R}^m \mid L_2 u \leq M_2\}$ and $T_s^{-1}(\varsigma_j) = \{s \in \mathbb{R}^n \mid L_3 s \leq M_3\}$. Then, by substituting

$$s[t] = A^t s[0] + \sum_{k=0}^{t-1} (A^k B_u u[t-1-k] + A^k B_d d[t-1-k])$$

and replacing $s[t] \in T_s^{-1}(\varsigma_i)$, $u[t] \in U$ and $s[N] \in T_s^{-1}(\varsigma_j)$ with $L_1 s[t] \leq M_1$, $L_2 u[t] \leq M_2$ and $L_3 s[N] \leq M_3$, respectively, in (7), it can be easily checked that (7) can be rewritten in the form $L[s[0], \hat{u}]' \leq M - G\hat{d}$, where $\hat{u} \triangleq [u[0]', \dots, u[N-1]']' \in \mathbb{R}^{mN}$, $\hat{d} \triangleq [d[0]', \dots, d[N-1]']' \in D^N$ and the matrices $L \in \mathbb{R}^{r \times n+mN}$ and $G \in \mathbb{R}^{r \times pN}$ and the vector $M \in \mathbb{R}^r$ can be obtained from $L_1, L_2, L_3, M_1, M_2, M_3, A, B_u$ and B_d .

Using properties of polyhedral convexity, we can prove the following result.

Proposition 2. *Suppose D is a closed and bounded polyhedral subset of \mathbb{R}^p and \overline{D} is the set of all its extreme points. Let $P \triangleq \{y \in \mathbb{R}^{n+mN} \mid Ly \leq M - G\hat{d}, \forall \hat{d} \in \overline{D}^N\}$ and S_0 be the projection of P onto its first n coordinates, i.e.,*

$$S_0 = \left\{ s \in \mathbb{R}^n \mid \exists \hat{u} \in \mathbb{R}^{mN} \text{ s.t. } L \begin{bmatrix} s \\ \hat{u} \end{bmatrix} \leq M - G\hat{d}, \forall \hat{d} \in \overline{D}^N \right\}.$$

Then, the Reachability Problem defined in Section V-A is feasible for any $s[0] \in S_0$.

Using Proposition 2, the problem of computing S_0 such that the Reachability Problem is feasible for any $s[0] \in S_0$ is reduced to computing a projection of the intersection of finite sets.

C. State Space Discretization and Correctness of the System

In general, given the predicate-based partition of $dom(S)$ and $i, j \in \{1, \dots, n\}$, the reachability relation between ς_i and ς_j may not be established through the set S_0 of $s[0]$ starting from which the Reachability Problem defined in Section V-A is feasible. (Due to the constraints on u and a specific choice of the finite horizon N , $T_s^{-1}(\varsigma_i)$ is not necessarily covered by S_0 .) To partially alleviate this limitation, we refine the partition based on the reachability relation to increase the number of valid discrete state transitions of \mathbb{D} . The underlying idea is that starting with an arbitrary pair of ς_i and ς_j , we determine the set S_0 of feasible $s[0]$ for the Reachability Problem. Then, we partition $T_s^{-1}(\varsigma_i)$ into $T_s^{-1}(\varsigma_i) \cap S_0$, labeled by $\varsigma_{i,1}$, and $T_s^{-1}(\varsigma_i) \setminus S_0$, labeled by $\varsigma_{i,2}$, and obtain the following reachability relations: $\varsigma_{i,1} \rightsquigarrow \varsigma_j$ and $\varsigma_{i,2} \rightsquigarrow \varsigma_j$. This process is continued until some pre-specified termination criteria are met.

Table I shows the pseudo-code of the algorithm where a prescribed lower bound Vol_{min} on the volume of each cell in the new partition is used as a termination criterion. The algorithm terminates when no cell can be partitioned such that the volumes of the two resulting new cells are both greater than Vol_{min} . Larger Vol_{min} causes the algorithm to terminate sooner. Other termination criteria such as the maximum number of iterations can be used as well. Note that the point at which the algorithm terminates affects the reachability

TABLE I
DISCRETIZATION ALGORITHM

Discretization Algorithm
input: The lower bound on cell volume (Vol_{min}), the parameters A, B_u, B_d, U, D, N of the Reachability Problem, and the original partition ($\{T_s^{-1}(\varsigma_i) \mid i \in \{1, \dots, n\}\}$)
output: The new partition sol
$sol = \{T_s^{-1}(\varsigma_i) \mid i \in \{1, \dots, n\}\}; JJ = \{(i, j) \mid i, j \in \{1, \dots, n\}\};$ while ($size(JJ) > 0$)
Pick arbitrary ς_i and ς_j where $(i, j) \in JJ;$
Compute the set S_0 of $s[0]$ starting from which the Reachability Problem is feasible for the previously chosen ς_i and $\varsigma_j;$
if ($volume(sol[i] \cap S_0) > Vol_{min}$ and $volume(sol[i] \setminus S_0) > Vol_{min}$) then
Replace $sol[i]$ with $sol[i] \cap S_0$ and append $sol[i] \setminus S_0$ to $sol;$
For each $k \in \{1, \dots, size(sol)\}$, add $(i, k), (k, i), (size(sol), k)$ and $(k, size(sol))$ to $JJ;$
else
Remove (i, j) from $JJ;$
endif
endwhile

between discrete controlled states of the new partition and as a result, affects the realizability of the specification. Generally, a coarse partition may render the specification unrealizable whereas a fine partition increases the computational cost. A way to decide when to terminate the algorithm is to start with a coarse partition and keep refining it until the specification is realizable or the computational resources are exhausted.

It can be shown that the resulting partition of $dom(S)$, after applying the discretization algorithm, contains at most $N_{max} = volume(dom(S)) / Vol_{min}$ cells. Hence, the number of iterations does not exceed $\frac{N_{max}}{2} (N_{max} + 1)$. We denote the set of all the discrete controlled states corresponding to the resulting partition of $dom(S)$ by \mathcal{S}' . Since the partition obtained from the proposed algorithm is a refined partition of $\{T_s^{-1}(\mathcal{S}_1), \dots, T_s^{-1}(\mathcal{S}_n)\}$ and $\mathcal{V} = \mathcal{S} \times \mathcal{E}$ is proposition preserving, it is trivial to show that $\mathcal{V}' = \mathcal{S}' \times \mathcal{E}$ is also proposition preserving. For simplicity of notation, we call \mathcal{S}' as \mathcal{S} and \mathcal{V}' as \mathcal{V} for the rest of the paper.

We define the finite transition system \mathbb{D} that serves as the abstract model of \mathbb{S} as: (a) $\mathcal{V} = \mathcal{S} \times \mathcal{E}$ is the set of states of \mathbb{D} , and (b) $\nu_i \rightarrow \nu_j$ where $\nu_i, \nu_j \in \mathcal{V}$, $\nu_i = (\varsigma_i, \epsilon_i)$ and $\nu_j = (\varsigma_j, \epsilon_j)$ only if $\varsigma_i \rightsquigarrow \varsigma_j$. Using the abstract model \mathbb{D} , a discrete planner that guarantees the satisfaction of φ while ensuring that the discrete plans are restricted to those satisfying the reachability relations can be automatically constructed using the digital design synthesis tool [19].

From the stutter invariant property of φ [28], the formulation of the Reachability Problem and the proposition preserving property of \mathcal{V} , it is straightforward to prove the correctness of the hierarchical approach.

Proposition 3. *Let $\sigma_d = \nu_0 \nu_1 \dots$ be an infinite sequence of discrete states of \mathbb{D} where for each natural number k , $\nu_k \rightarrow \nu_{k+1}$, $\nu_k = (\varsigma_k, \epsilon_k)$, $\varsigma_k \in \mathcal{S}$ is the discrete controlled state and $\epsilon_k \in \mathcal{E}$ is the discrete environment state. If $\sigma_d \models_d \varphi$, then by applying a sequence of control signals, each corresponding to a solution of the Reachability Problem with $\varsigma_i = \varsigma_k$ and $\varsigma_j = \varsigma_{k+1}$, the infinite sequence of continuous states $\sigma = \nu_0 \nu_1 \nu_2 \dots$ satisfies φ .*

Remark 2. In verifying the reachability relation, we consider the restricted case where the horizon length N is fixed and given. In addition, the discretization algorithm terminates when a user-defined termination criterion is met. This may lead to a conservative result, i.e., cell ζ_j (or part of it) may be reachable from cell ζ_i according to Definition 7 even though the discretization algorithm declares that $\zeta_i \not\rightsquigarrow \zeta_j$. Hence, the resulting partition may render the specification unrealizable even though there exists a provably correct control protocol.

VI. RECEDING HORIZON FRAMEWORK

The main limitation of the synthesis of finite state automata from their LTL specifications [19] is the state explosion problem. In the worst case, the resulting automaton may contain all the possible states of the system. For example, if the system has $|V|$ variables, each can take any value in $\{1, \dots, M\}$, then there may be as many as $M^{|V|}$ nodes in the automaton. This type of computational complexity limits the application of synthesis to relatively small problems.

Similar computational complexity is also encountered in the area of constrained optimal control. In the controls domain, an effective and well-established technique to address this issue is to design and implement control strategies in a receding horizon manner, i.e., optimize over a *shorter* horizon, starting from the currently observed state, implement the initial control action, move the horizon one step ahead, and re-optimize. This approach reduces the computational complexity by essentially solving a sequence of *smaller* optimization problems, each with a specific initial condition (as opposed to optimizing with *any* initial condition in traditional optimal control). Under certain conditions, receding horizon control strategies are known to lead to closed-loop stability [26], [22], [29]. See, e.g., [30] for a detailed discussion on constrained optimal control, including finite horizon optimal control and receding horizon control.

To reduce computational complexity in the synthesis of finite state automata, we apply an idea similar to the traditional receding horizon control. First, we observe that in many applications, it is not necessary to plan for the whole execution, taking into account all the possible behaviors of the environment since a state that is very far from the current state of the system typically does not affect the near future plan. Consider, for example, the robot motion planning problem described in Example 2. Suppose C_1 or C_2 is very far away from the initial position of the robot. Under certain conditions, it may be sufficient to only plan out an execution for only a short segment ahead and implement it in a receding horizon fashion, i.e., re-compute the plan as the robot moves, starting from the currently observed state (rather than from all initial conditions satisfying φ_{init} as the original specification (2) suggests). In this section, we present a sufficient condition and a receding horizon strategy that allows the synthesis to be performed on a smaller domain; thus, substantially reduces the number of states (or nodes) of the automaton while still ensuring the system correctness with respect to the LTL specification (2).

We assume that a finite state abstraction \mathbb{D} of the system \mathbb{S} has been constructed using, for example, the discretization




Fig. 2. Illustration of the receding horizon framework showing the relationships between the states of \mathcal{V} and between the subsets $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$

algorithm presented in Section V-C. Let \mathcal{V} be the finite set of states of \mathbb{D} . We consider a specification of the form (6) since, from Proposition 1, the specification (2) can be reduced to this form. Let Φ be a propositional formula of variables from V such that $\psi_{init} \implies \Phi$ is a tautology, i.e., any state $\nu \in \mathcal{V}$ that satisfies ψ_{init} also satisfies Φ . For each progress property $\square \diamond \psi_{g,i}$, $i \in I_g$, suppose there exists a collection of subsets $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$ of \mathcal{V} such that

- $\mathcal{W}_0^i \cup \mathcal{W}_1^i \cup \dots \cup \mathcal{W}_p^i = \mathcal{V}$,
- $\psi_{g,i}$ is satisfied for any $\nu \in \mathcal{W}_0^i$, i.e., \mathcal{W}_0^i is the set of the states that constitute the progress of the system, and
- $\mathcal{P}^i := (\{\mathcal{W}_0^i, \dots, \mathcal{W}_p^i\}, \leq_{\psi_{g,i}})$ is a partially ordered set defined such that $\mathcal{W}_0^i <_{\psi_{g,i}} \mathcal{W}_j^i, \forall j \neq 0$.

Define a map $\mathcal{F}^i : \{\mathcal{W}_0^i, \dots, \mathcal{W}_p^i\} \rightarrow \{\mathcal{W}_0^i, \dots, \mathcal{W}_p^i\}$ such that $\mathcal{F}^i(\mathcal{W}_j^i) <_{\psi_{g,i}} \mathcal{W}_j^i, \forall j \neq 0$.

Consider a simple case where $\{\nu_1, \dots, \nu_{10}\}$ is the set \mathcal{V} of states, ν_{10} satisfies $\psi_{g,i}$, and the states in \mathcal{V} are organized into 5 subsets $\mathcal{W}_0^i, \dots, \mathcal{W}_4^i$. The relationships between the states in \mathcal{V} and the subsets $\mathcal{W}_0^i, \dots, \mathcal{W}_4^i$ are illustrated in Fig. 2. The partial order may be defined as $\mathcal{W}_0^i < \mathcal{W}_1^i < \dots < \mathcal{W}_4^i$ and the map \mathcal{F}^i may be defined as $\mathcal{F}^i(\mathcal{W}_j^i) = \mathcal{W}_{j-2}^i, \forall j \geq 2$, $\mathcal{F}^i(\mathcal{W}_1^i) = \mathcal{W}_0^i$ and $\mathcal{F}^i(\mathcal{W}_0^i) = \mathcal{W}_0^i$. Suppose ν_1 is the initial state of the system. The key idea of the receding horizon framework, as described later, is to plan from ν_1 to a state in $\mathcal{F}^i(\mathcal{W}_4^i) = \mathcal{W}_2^i$, rather than planning from the initial state ν_1 to the goal state ν_{10} in one shot, taking into account all the possible behaviors of the environment. Once a state in \mathcal{W}_3^i , i.e., ν_5 or ν_6 is reached, we then replan from that state to a state in $\mathcal{F}^i(\mathcal{W}_3^i) = \mathcal{W}_1^i$. We repeat this process until ν_{10} is reached. Under certain sufficient conditions presented later, this strategy ensures the correctness of the overall execution of the system.

Formally, with the above definitions of Φ , $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$ and \mathcal{F}^i , we define a short-horizon specification Ψ_j^i associated with \mathcal{W}_j^i for each $i \in I_g$ and $j \in \{0, \dots, p\}$ as

$$\begin{aligned} \Psi_j^i &\triangleq ((\nu \in \mathcal{W}_j^i) \wedge \Phi \wedge \square \psi_e^e \wedge \bigwedge_{k \in I_f} \square \diamond \psi_{f,k}^e) \\ &\implies (\bigwedge_{k \in I_s} \square \psi_{s,k} \wedge \square \diamond (\nu \in \mathcal{F}^i(\mathcal{W}_j^i))) \wedge \square \Phi, \end{aligned} \quad (8)$$

where ν is the state of the system and ψ_e^e , $\psi_{f,k}^e$ and $\psi_{s,k}$ are defined as in (6). An automaton \mathcal{A}_j^i satisfying Ψ_j^i defines a strategy for going from a state $\nu_1 \in \mathcal{W}_j^i$ to a state $\nu_2 \in \mathcal{F}^i(\mathcal{W}_j^i)$ while satisfying the safety requirements $\bigwedge_{i \in I_s} \square \psi_{s,i}$ and maintaining the invariant Φ . The roles of \mathcal{P}^i , \mathcal{F}^i and Φ are discussed later in Section VI-A.

Receding Horizon Strategy: For each $i \in I_g$ and $j \in \{0, \dots, p\}$, construct an automaton \mathcal{A}_j^i satisfying Ψ_j^i . Let $I_g = \{i_1, \dots, i_n\}$ and define a corresponding ordered set (i_1, \dots, i_n) . This order only affects the sequence of progress




Fig. 3. A graphical description of the receding horizon strategy for a special case where for each $i \in I_g$, $\mathcal{W}_j^i \prec_{\psi_{g,i}} \mathcal{W}_k^i, \forall j < k$, $F^i(\mathcal{W}_j^i) = \mathcal{W}_{j-1}^i, \forall j > 0$ and $F^i(\mathcal{W}_0^i) = \mathcal{W}_0^i$.

properties $\psi_{g,i_1}, \dots, \psi_{g,i_n}$ that the system tries to satisfy and can be picked arbitrarily without affecting the correctness of the strategy.

- (1) Determine the index j_1 such that the current state $\nu_0 \in \mathcal{W}_{j_1}^{i_1}$. If $j_1 \neq 0$, then execute automaton $\mathcal{A}_{j_1}^{i_1}$ until the system reaches a state $\nu_1 \in \mathcal{W}_k^{i_1}$ where $\mathcal{W}_k^{i_1} \prec_{\psi_{g,i_1}} \mathcal{W}_{j_1}^{i_1}$. Note that since the union of $\mathcal{W}_1^{i_1}, \dots, \mathcal{W}_p^{i_1}$ is the set \mathcal{V} of all the states, given any $\nu_0, \nu_1 \in \mathcal{V}$, there exist $j_1, k \in \{0, \dots, p\}$ such that $\nu_0 \in \mathcal{W}_{j_1}^{i_1}$ and $\nu_1 \in \mathcal{W}_k^{i_1}$.
- (2) If the current state $\nu_1 \notin \mathcal{W}_0^{i_1}$, switch to automaton $\mathcal{A}_k^{i_1}$ where the index k is chosen such that the current state $\nu_1 \in \mathcal{W}_k^{i_1}$. Execute $\mathcal{A}_k^{i_1}$ until the system reaches a state that is smaller in the partial order \mathcal{P}^{i_1} . Repeat this process until a state $\nu_2 \in \mathcal{W}_0^{i_1}$ is reached.
- (3) Switch to automaton $\mathcal{A}_{j_2}^{i_2}$ where the index j_2 is chosen such that the current state $\nu_2 \in \mathcal{W}_{j_2}^{i_2}$. Repeat step (2) with i_1 replaced by i_2 for the partial order \mathcal{P}^{i_2} until a state $\nu_3 \in \mathcal{W}_0^{i_2}$ is reached. Repeat this process with i_2 replaced by i_3, i_4, \dots, i_n until a state $\nu_n \in \mathcal{W}_0^{i_n}$ is reached.
- (4) Repeat steps (1)–(3).

A graphical description of this strategy is depicted in Figure 3. Starting from a state ν_0 , the system executes the automaton $\mathcal{A}_{j_1}^{i_1}$ where the index j_1 is chosen such that $\nu_0 \in \mathcal{A}_{j_1}^{i_1}$. Step (2) ensures that a state $\nu_2 \in \mathcal{W}_0^{i_1}$ (i.e., a state satisfying ψ_{g,i_1}) is eventually reached. This state ν_2 belongs to some set, say, $\mathcal{W}_{j_2}^{i_2}$ in the partial order \mathcal{P}^{i_2} . The system then works through this partial order until a state $\nu_3 \in \mathcal{W}_0^{i_2}$ (i.e., a state satisfying ψ_{g,i_2}) is reached. This process is repeated until a state ν_n satisfying ψ_{g,i_n} is reached. At this point, for each $i \in I_g$, a state satisfying $\psi_{g,i}$ has been visited at least once in the execution. In addition, the state ν_n belongs to some set in the partial order \mathcal{P}^{i_1} and the whole process is repeated, ensuring that for each $i \in I_g$, a state satisfying $\psi_{g,i}$ is visited infinitely often in the execution.

Theorem 1. Suppose Ψ_j^i is realizable for each $i \in I_g, j \in \{0, \dots, p\}$. Then the receding horizon strategy ensures that the system is correct with respect to the specification (6), i.e., any execution of the system satisfies (6).

Proof: Consider an arbitrary execution σ of the system that satisfies the assumption part of (6). We want to show that the safety properties $\psi_{s,i}, i \in I_s$, hold throughout the execution

and for each $i \in I_g$, a state satisfying $\psi_{g,i}$ is reached infinitely often.

Let $\nu_0 \in \mathcal{V}$ be the initial state of the system and let the index j_1 be such that $\nu_0 \in \mathcal{W}_{j_1}^{i_1}$. From the tautology of $\psi_{init} \implies \Phi$, it is easy to show that σ satisfies the assumption part of $\Psi_{j_1}^{i_1}$ as defined in (8). Thus, if $j_1 = 0$, then $\mathcal{A}_0^{i_1}$ ensures that a state ν_2 satisfying ψ_{g,i_1} is eventually reached and the safety properties $\psi_{s,i}, i \in I_s$, hold at every position of σ up to and including the point where ν_2 is reached. Otherwise, $j_1 \neq 0$ and $\mathcal{A}_{j_1}^{i_1}$ ensures that eventually, a state $\nu_1 \in \mathcal{W}_k^{i_1}$ where $\mathcal{W}_k^{i_1} \prec_{\psi_{g,i_1}} \mathcal{W}_{j_1}^{i_1}$ is reached, i.e., ν_1 is the state of the system at some position l_1 of σ . In addition, the invariant Φ and all the safety properties $\psi_{s,i}, i \in I_s$, are guaranteed to hold at all the positions of σ up to and including the position l_1 . According to the receding horizon strategy, the planner switches to the automaton $\mathcal{A}_k^{i_1}$ at position l_1 of σ . Since $\nu_1 \in \mathcal{W}_k^{i_1}$ and ν_1 satisfies Φ , the assumption part of $\Psi_k^{i_1}$ as defined in (8) is satisfied. Using the previous argument, we get that $\Psi_k^{i_1}$ ensures that all the safety properties $\psi_{s,i}, i \in I_s$, hold at every position of σ starting from position l_1 up to and including position l_2 at which the planner switches the automaton (from $\mathcal{A}_k^{i_1}$) and Φ holds at position l_2 . By repeating this procedure and using the finiteness of the set $\{\mathcal{W}_0^{i_1}, \dots, \mathcal{W}_p^{i_1}\}$ and its partial order condition, eventually the automaton $\mathcal{A}_0^{i_1}$ is executed which ensures that σ contains a state ν_2 satisfying ψ_{g,i_1} and step (2) terminates.

Applying the previous argument to step (3), we get that step (3) terminates and before it terminates, the safety properties $\psi_{s,i}, i \in I_s$, and the invariant Φ hold throughout the execution and for each $i \in I_g$, a state satisfying $\psi_{g,i}$ has been reached at least once. By continually repeating steps (1)–(3), the receding horizon strategy ensures that $\psi_{s,i}, i \in I_s$, hold throughout the execution and for each $i \in I_g$, a state satisfying $\psi_{g,i}$ is reached infinitely often. ■

Remark 3. For each $i \in I_g$ and $j \in \{0, \dots, p\}$, checking the realizability of Ψ_j^i requires considering all the initial conditions in \mathcal{W}_j^i satisfying Φ . However, as will be further discussed in Section VII, when a strategy (i.e., a finite state automaton satisfying Ψ_j^i) is to be extracted, only the currently observed state needs to be considered as the initial condition. Typically, the realizability can be checked symbolically and enumeration of states is only required when a strategy needs to be extracted [19]. Symbolic methods are known to handle large number of states, in practice, significantly better than enumeration-based methods. Hence, state explosion usually occurs in the synthesis (i.e., strategy extraction) stage rather than the realizability checking stage. By considering only the currently observed state as the initial state in the synthesis stage, the receding horizon strategy delays state explosion both by considering a short-horizon problem and a specific initial state.

A. Remarks on the Partial Order \mathcal{P}^i and the Invariant Φ

Subsets $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$ of \mathcal{V} , the corresponding partial order \mathcal{P}^i , the map \mathcal{F}^i and the propositional formula Φ are at the core of the receding horizon framework. Roughly speaking, \mathcal{P}^i provides a measure of “closeness” to the states satisfying $\psi_{g,i}$. Since each short-horizon specification Ψ_j^i asserts that

the system eventually reaches a state that is smaller in the partial order, it ensures that each automaton \mathcal{A}_j^i brings the system “closer” to the states satisfying $\psi_{g,i}$. The map \mathcal{F}^i thus defines the horizon length for these short-horizon problems. In general, the size of \mathcal{A}_j^i increases with the horizon length. However, with too short horizon, the specification Ψ_j^i may not be realizable. Since the computational complexity increases with the size of \mathcal{A}_j^i , a good practice is to choose the shortest horizon, subject to the realizability of Ψ_j^i , so that the automaton \mathcal{A}_j^i contains as small number of states as possible.

The propositional formula Φ can be viewed as an invariant of the system. It adds an assumption on the initial state of each automaton \mathcal{A}_j^i and is introduced in order to make Ψ_j^i realizable. Without Φ , the set of initial states of \mathcal{A}_j^i includes all states $\nu \in \mathcal{W}_j^i$. However, starting from some “bad” state (e.g. unsafe state) in \mathcal{W}_j^i , there may not exist a strategy for the system to satisfy Ψ_j^i . Φ is essentially used to eliminate the possibility of starting from these “bad” states.

Computation of these critical elements requires insights for each problem domain. (See the example presented in Section VIII for such insights for an autonomous driving problem.) However, automatic construction of certain elements is possible, given other elements. For example, given an invariant Φ and subsets $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$ of \mathcal{V} , construction of the partial order \mathcal{P}^i and the map \mathcal{F}^i can be automated. Such an automatic construction is discussed in detail in Section VII.

On the other hand, given a partially order set \mathcal{P}^i and a map \mathcal{F}^i , one way to determine Φ is to start with $\Phi \equiv \text{True}$ and check the realizability of the resulting Ψ_j^i . If there exist $i \in I_g$ and $j \in \{0, \dots, p\}$ such that Ψ_j^i is not realizable, the synthesis process provides the initial state ν^* of the system starting from which there exists a set of moves of the environment such that the system cannot satisfy Ψ_j^i . This information provides guidelines for constructing Φ , i.e., we can add a propositional formula to Φ that prevents the system from reaching the state ν^* . This procedure can be repeated until Ψ_j^i is realizable for any $i \in I_g$ and $j \in \{0, \dots, p\}$ or until Φ excludes all the possible states, in which case either the original specification is unrealizable or the proposed receding horizon strategy cannot be applied with the given partially order set \mathcal{P}^i and map \mathcal{F}^i . Such an automatic construction of Φ has been implemented in TuLiP, a Python-based software toolbox for receding horizon temporal logic planning [31].

In Section VIII, we illustrate a simple computation of $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$ for a particular system where the notion of “distance” to the goals can be easily defined. A systematic approach to construct $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$ for a general system, however, is subject to current research. Automatic construction of Φ given \mathcal{P}^i and \mathcal{F}^i and automatic construction of \mathcal{P}^i and \mathcal{F}^i given $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$ and Φ motivates the following iterative approach for computing these critical elements, provided that $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$ are given. First, start with an initial guess for Φ (e.g. only exclude the unsafe states) and compute the corresponding \mathcal{P}^i and \mathcal{F}^i . If the resulting \mathcal{P}^i and \mathcal{F}^i render Ψ_j^i unrealizable for some $i \in I_g, j \in \{0, \dots, p\}$, we recompute Φ for this \mathcal{P}^i and \mathcal{F}^i . Such an iterative approach is subject to current study.

B. The Computational Complexity and Completeness of the Receding Horizon Strategy

The receding horizon framework essentially reduces the original synthesis problem to a set of smaller problems. The partial order relation on $\{\mathcal{W}_0^i, \dots, \mathcal{W}_p^i\}$ induces a directed graph \mathbb{G}^i whose nodes are $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$. In this induced graph, there is an edge from \mathcal{W}_0^i to itself and to each of the other nodes. For each $j, k \in \{0, \dots, p\}$ such that $j \neq 0$, there is an edge from \mathcal{W}_j^i to \mathcal{W}_k^i if and only if $\mathcal{W}_j^i <_{\psi_{g,i}} \mathcal{W}_k^i$ and there does not exist $l \in \{0, \dots, p\}$ such that $\mathcal{W}_j^i <_{\psi_{g,i}} \mathcal{W}_l^i <_{\psi_{g,i}} \mathcal{W}_k^i$. Suppose for each $j \in \{0, \dots, p\}$, there is only one path from \mathcal{W}_j^i to $\mathcal{F}^i(\mathcal{W}_j^i)$ in \mathbb{G}^i . Then, we define the horizon length T_j^i for a short-horizon specification Ψ_j^i as the length of the path from \mathcal{W}_j^i to $\mathcal{F}^i(\mathcal{W}_j^i)$ in \mathbb{G}^i .

Let N_S and N_E be the number of possible controlled and environment states, respectively, for a short-horizon problem with horizon length 1. Since the environment states for different short-horizon problems may be completely independent (e.g., whether there is an obstacle in a certain cell may not depend on whether there is an obstacle in other cells), it can be shown that the size of a short-horizon problem with horizon length T can be as large as $TM_S M_E^T$. Recall that the computational complexity of the synthesis problem for GR[1] formula is $O(|\mathcal{V}|^3)$ where $|\mathcal{V}|$ is the size of the state space. Hence, the computational complexity of solving each short-horizon problem is $O((TM_S M_E^T)^3)$ where $T = \max_{i,j} T_j^i \leq p$. Since there are the total of $(p+1)|I_g|$ short-horizon problems where $|I_g|$ is the cardinality of I_g , the overall complexity of our receding horizon approach is $O((p+1)|I_g|(TM_S M_E^T)^3)$. (Note that each of the short-horizon problems can be solved independently; thus, our algorithm is easily parallelizable.) In comparison, solving the original synthesis problem without applying the receding horizon planning procedure may lead to the complexity of $O((pM_S M_E^p)^3)$. Thus, if the horizon length is chosen such that $T \ll p$, the receding horizon approach can significantly reduce the complexity of the synthesis problem.

In Section VIII, we provide an example of a synthesis problem that may not be solved without the receding horizon approach due to the size of the state space and illustrate the application of the receding horizon approach that allows such problem to be solved without excessive computational power. Other examples can be found in our previous work [14], [16].

On the other hand, the receding horizon approach is not complete. Even if there exists a control strategy that satisfies the original specification in (6), there may not exist an invariant Φ or a collection of subsets $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$ that allow the receding horizon strategy to be applied since the corresponding Ψ_j^i may not be realizable for all $i \in I_g$ and $j \in \{0, \dots, p\}$.

Remark 4. Traditional receding horizon control is known to not only reduce computational complexity but also increase the robustness of the system with respect to exogenous disturbances and modeling uncertainties [26]. With disturbances and modeling uncertainties, an actual execution of the system usually deviates from a reference trajectory s_d . Receding horizon control allows the current state of the system to be continually re-evaluated so s_d can be adjusted accordingly based on the externally received reference if the actual execution of the

system does not match it. Such an effect may be expected in our extension of the traditional receding horizon control. Verifying this property is subject to current study.

VII. IMPLEMENTATION OF THE RECEDING HORIZON FRAMEWORK

In order to implement the receding horizon strategy, a partial order \mathcal{P}^i and the corresponding map \mathcal{F}^i need to be defined for each $i \in I_g$. We now present an implementation of this strategy, allowing \mathcal{P}^i and \mathcal{F}^i to be automatically determined for each $i \in I_g$ while ensuring that all the short-horizon specifications $\Psi_j^i, i \in I_g, j \in \{0, \dots, p\}$, as defined in (8) are realizable.

Given an invariant Φ and subsets $\mathcal{W}_0^i, \dots, \mathcal{W}_p^i$ of \mathcal{V} for each $i \in I_g$, we first construct a finite transition system \mathbb{T}^i with the set of states $\{\mathcal{W}_0^i, \dots, \mathcal{W}_p^i\}$. For each $j, k \in \{0, \dots, p\}$, there is a transition $\mathcal{W}_j^i \rightarrow \mathcal{W}_k^i$ in \mathbb{T}^i only if $j \neq k$ and the specification in (8) is realizable with $\mathcal{F}^i(\mathcal{W}_j^i) = \mathcal{W}_k^i$. The finite transition system \mathbb{T}^i can be regarded as an abstraction of the finite state model \mathbb{D} of the physical system \mathbb{S} , i.e., a higher-level abstraction of \mathbb{S} .

Suppose Φ is defined such that there exists a path in \mathbb{T}^i from \mathcal{W}_j^i to \mathcal{W}_0^i for all $i \in I_g, j \in \{1, \dots, p\}$. (Verifying this property is basically a graph search problem. If a path does not exist, Φ can be re-computed using a procedure described in Section VI-A.) We propose a hierarchical control structure with three components (cf. Fig. 4): goal generator, trajectory planner, and continuous controller.

Goal generator: Pick a sequence¹ (i_1, \dots, i_n) for the elements of the unordered set $I_g = \{i_1, \dots, i_n\}$ and maintain an index $k \in \{1, \dots, n\}$ throughout the execution. Starting with $k = 1$, in each iteration, the goal generator performs the following tasks.

- (a1) Receive the currently observed state of the plant (i.e. the controlled state) and environment.
- (a2) If the abstract state corresponding to the currently observed state belongs to $\mathcal{W}_0^{i_k}$, update k to $(k \bmod n) + 1$.
- (a3) If k was updated in step (a2) or this is the first iteration, then based on the higher level abstraction \mathbb{T}^{i_k} of the physical system \mathbb{S} , compute a path from $\mathcal{W}_j^{i_k}$ to $\mathcal{W}_0^{i_k}$ where the index $j \in \{0, \dots, p\}$ is chosen such that the abstract state corresponding to the currently observed state belongs to $\mathcal{W}_j^{i_k}$.
- (a4) If a new path is computed in step (a3), then issue this path (i.e., a sequence $\mathcal{G} = \mathcal{W}_{l_0}^{i_k}, \dots, \mathcal{W}_{l_m}^{i_k}$ for some $m \in \{0, \dots, p\}$ where $l_0, \dots, l_m \in \{0, \dots, p\}, l_0 = j, l_m = 0, l_\alpha \neq l_{\alpha'}$ for any $\alpha \neq \alpha'$, and there exists a transition $\mathcal{W}_{l_\alpha}^{i_k} \rightarrow \mathcal{W}_{l_{\alpha+1}}^{i_k}$ in \mathbb{T}^{i_k} for any $\alpha < m$) to the trajectory planner.

The problem of finding a path in \mathbb{T}^{i_k} from $\mathcal{W}_j^{i_k}$ to $\mathcal{W}_0^{i_k}$ can be efficiently solved using any graph search algorithm [32], such as Dijkstra's and A*. To reduce the original synthesis problem into a set of problems with shorter horizon, the cost

¹As discussed in the description of the receding horizon strategy in Section VI, this sequence can be picked arbitrarily. In general, its definition affects a strategy the system chooses to satisfy the specification (6) as it corresponds to the sequence of progress properties $\psi_{g,i_1}, \dots, \psi_{g,i_n}$ the system tries to satisfy.

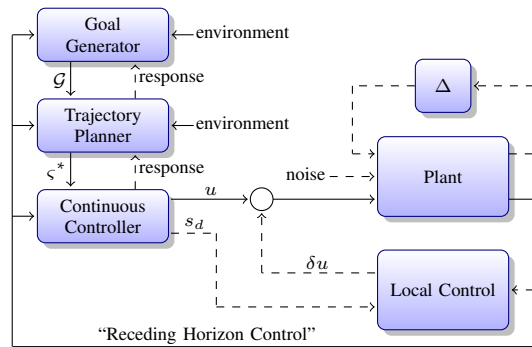


Fig. 4. A system with the control protocol implemented in a receding horizon manner. Besides the components discussed in this paper, Δ , which captures uncertainties in the plant model, may be added to make the model more realistic. Additionally, a local control may be implemented to account for the effect of noise, disturbances, and unmodeled dynamics. The inputs and outputs of these two components, not considered in this paper, are drawn in dashed.

on each edge $(\mathcal{W}_{l_\alpha}^{i_k}, \mathcal{W}_{l_{\alpha'}}^{i_k})$ of the graph built from \mathbb{T}^{i_k} may be defined, for example, as an exponential function of the “distance” between the sets $\mathcal{W}_{l_\alpha}^{i_k}$ and $\mathcal{W}_{l_{\alpha'}}^{i_k}$ so that a path with smaller cost contains segments of shorter “distance”.

Trajectory planner: The trajectory planner maintains the latest sequence $\mathcal{G} = \mathcal{W}_{l_0}^{i_k}, \dots, \mathcal{W}_{l_m}^{i_k}$ of goal states received from the goal generator, an index $q \in \{1, \dots, m\}$ of the current goal state in \mathcal{G} , a strategy \mathbb{F} represented by a finite state automaton, and the next abstract state ν^* throughout the execution. Starting with $q = 1$ and \mathbb{F} and ν^* initialized as an empty finite state automaton and a null state, respectively, in each iteration, the trajectory planner performs the following tasks.

- (b1) Receive the currently observed state of the plant and environment.
- (b2) If a new sequence of goal states is received from the goal generator, update \mathcal{G} , q and ν^* to this latest sequence of goal states, 1, and null. Otherwise, if the abstract state corresponding to the currently observed state belongs to $\mathcal{W}_{l_q}^{i_k}$, update q and ν^* to $q + 1$ and null.
- (b3) If ν^* is null, then based on the abstraction \mathbb{D} of the physical system \mathbb{S} , synthesize a strategy that satisfies the specification in (8) with $\mathcal{F}^i(\mathcal{W}_j^i) = \mathcal{W}_{l_q}^{i_k}$, starting from the abstract state ν_0 corresponding to the currently observed state, i.e., replace the assumption $\nu \in \mathcal{W}_j^i$ with $\nu = \nu_0$. Assign this strategy to \mathbb{F} and update ν^* to the state following the initial state in \mathbb{F} based on the current environment state.
- (b4) If the controlled state ζ^* component of ν^* corresponds to the currently observed state of the plant, update ν^* to the state following ν^* in \mathbb{F} based on the current environment state.
- (b5) If ν^* was updated in step (b3) or (b4), then issue ζ^* to the continuous controller.

Continuous controller: The continuous controller maintains the most recent (abstract) final controlled state ζ^* from the trajectory planner. In each iteration, it receives the currently observed state s of the plant. Then, it computes a control signal u such that the continuous execution of the system eventually

reaches the cell of \mathbb{D} corresponding to ζ^* while always staying in the cell corresponding to the abstract controlled state ζ^* and the cell containing s . Essentially, the continuous execution has to *simulate* the abstract plan computed by the trajectory planner. As discussed at the end of Section V-A, such a control signal can be computed by formulating a constrained optimal control problem and solved using off-the-shelf software.

From the construction of $\mathbb{T}^i, i \in I_g$, it can be verified that the composition of the goal generator and the trajectory planner correctly implements the receding horizon strategy described in Section VI. Roughly speaking, the path \mathcal{G} from \mathcal{W}_j^i to \mathcal{W}_0^i computed by the goal generator essentially defines the partial order \mathcal{P}^i and the corresponding map \mathcal{F}^i . For a set $\mathcal{W}_{l_\alpha}^i \neq \mathcal{W}_0^i$ contained in \mathcal{G} , we simply let $\mathcal{W}_{l_{\alpha+1}}^i < \mathcal{W}_{l_\alpha}^i$ and $\mathcal{F}^i(\mathcal{W}_{l_\alpha}^i) = \mathcal{W}_{l_{\alpha+1}}^i$ where $\mathcal{W}_{l_{\alpha+1}}^i$ immediately follows $\mathcal{W}_{l_\alpha}^i$ in \mathcal{G} . In addition, since, by assumption, for any $i \in I_g$ and $l \in \{0, \dots, p\}$, there exists a path in \mathbb{T}^i from \mathcal{W}_l^i to \mathcal{W}_0^i , it can be easily verified that the specification Ψ_l^i is realizable with $\mathcal{F}(\mathcal{W}_l^i) = \mathcal{W}_0^i$. Thus, to be consistent with the previously described receding horizon framework, we assign $\mathcal{W}_l^i > \mathcal{W}_0^i$ and $\mathcal{F}(\mathcal{W}_l^i) = \mathcal{W}_0^i$ for a set \mathcal{W}_l^i not contained in \mathcal{G} . Note that such \mathcal{W}_l^i that is not in the path \mathcal{G} does not affect the computational complexity of the synthesis algorithm. With this definition of the partial order \mathcal{P}^i and the corresponding map \mathcal{F}^i , we can apply Theorem 1 to conclude that the abstract plan generated by the trajectory planner ensures the correctness of the system with respect to the specification in (6). In addition, since the continuous controller *simulates* this abstract plan, the continuous execution is guaranteed to preserve the correctness of the system.

The resulting system is depicted in Fig. 4. Observe how this design corresponds to the planner-controller subsystem in Fig. 1 with the continuous controller having similar functionality as Path Follower, the trajectory planner having similar functionality as the composition of Traffic Planner and Path Planner, the goal generator having similar functionality as Mission Planner, and each of the sets $\mathcal{W}_1^i, \dots, \mathcal{W}_p^i$ being an entire road. Note that since the system is guaranteed to satisfy the specification in (6), the desired behavior (i.e. the guarantee part of (6)) is ensured only when the environment and the initial condition respect their assumptions. To moderate the sensitivity to violation of these assumptions, the trajectory planner may send a response to the goal generator, indicating the failure of executing the last received sequence of goals as a consequence of assumption violation. The goal generator can then remove the problematic transition from the corresponding finite transition system \mathbb{T}^i and re-compute a new sequence \mathcal{G} of goals. This procedure will be illustrated in the example presented in Section VIII. Similarly, a response may be sent from the continuous controller to the trajectory planner to account for the mismatch between the actual system and its model. In addition, a local control may be added in order to account for the effect of the noise and unmodeled dynamics captured by Δ .

VIII. EXAMPLE

As an initial step toward correct-by-construction design of complex embedded control systems, we consider a simple

autonomous driving problem in an urban-like environment. The state of the vehicle is the position (x, y) whose evolution is governed by

$$\dot{x}(t) = u_x(t) + d_x(t) \quad \text{and} \quad \dot{y}(t) = u_y(t) + d_y(t) \quad (9)$$

where $u_x(t)$ and $u_y(t)$ are control signals and $d_x(t)$ and $d_y(t)$ are external disturbances at time t . The control effort is subject the constraints $u_x(t), u_y(t) \in [-1, 1], \forall t \geq 0$. We assume that the disturbances are bounded by $d_x(t), d_y(t) \in [-0.1, 0.1], \forall t \geq 0$. More complicated dynamics of an omnidirectional vehicle is considered in [14] for a simple road network.

We consider the road network shown in Fig. 5 with 3 intersections, I_1, I_2 and I_3 , and 6 roads, R_1, R_2 (joining I_1 and I_3), R_3, R_4 (joining I_2 and I_3), R_5 (joining I_1 and I_3) and R_6 (joining I_1 and I_2). Each of these roads has two lanes going in opposite directions. The positive and negative directions for each road are shown in Fig. 5. We partition the roads and intersections into $N = 282$ cells (cf. Fig. 5), each of which may be occupied by an obstacle.

The planner-controller subsystem in Fig. 1 is implemented in a hierarchical fashion and naturally follows our general framework for designing a control protocol (Fig. 4). However, such a subsystem is typically designed by hand and validated through extensive simulations and field tests. Although a correct-by-construction approach has been applied in [33], it is based on building a finite state abstraction of the physical system and synthesizing a planner that computes a strategy for the whole execution, taking into account all the possible behaviors of the environment. As discussed in Section IV, this approach fails to handle even modest size problems due to its computational complexity. In this section, we apply the receding horizon scheme to substantially reduce computational complexity of correct-by-construction approach.




Fig. 5. The road network and its partition for the autonomous vehicle example. The solid (black) lines define the states in the set \mathcal{V} of the finite state model \mathbb{D} used by the trajectory planner. Examples of subsets \mathcal{W}_j^i are drawn in dotted (red) rectangles. The stars indicate the positions that need to be visited infinitely often.

A. System Specification

Given the system in (9), we want to design a control protocol for the vehicle based on the following desired behavior and assumptions.

Desired Behavior: Following the terminology and notations used in Section III, the desired behavior φ_s in (2) includes the following properties.

- (P1) Each of the two cells marked by star needs to be visited infinitely often.
- (P2) No collision is allowed, i.e., the vehicle cannot occupy the same cell as an obstacle.
- (P3) The vehicle stays in the right lane unless there is an obstacle blocking the lane.
- (P4) The vehicle can only proceed through an intersection when the intersection is clear.

Assumptions: We assume that the vehicle starts from an obstacle-free cell on R_1 with at least one obstacle-free cell adjacent to it. This constitutes the assumption φ_{init} on the initial condition of the system. The environment assumption φ_e encapsulates the following statements which are assumed to hold in any execution: (A1) obstacles may not block a road; (A2) an obstacle is detected before the vehicle gets too close to it, i.e., an obstacle may not instantly pop up right in front of the vehicle; (A3) sensing range is limited, i.e., the vehicle cannot detect an obstacle that is away from it farther than certain distance. (A4) to make sure that the stay-in-lane property is achievable, we assume that an obstacle does not disappear while the vehicle is in its vicinity; (A5) obstacles may not span more than a certain number of consecutive cells in the middle of the road; (A6) each of the intersections is clear infinitely often; and (A7) each of the cells marked by star and its adjacent cells are not occupied by an obstacle infinitely often.

In this example, we let this sensing range be 2 cells ahead in the driving direction. It can be shown [34] that the properties (P2) and (P3) and the assumptions (A1)–(A4) can be expressed in the form of the guarantee and the assumption parts of (6). Property (P4) can be expressed as a safety formula and property (P1) is a progress property. Finally, assumption (A5) can be expressed as a safety assumption on the environment while assumptions (A6) and (A7) can be expressed as justice requirements on the environment.

B. Correct-by-Construction Control Protocol

We follow the approach described in Section IV. First, we compute a finite state abstraction \mathbb{D} of the system. Following the scheme in Section V, a state ν of \mathbb{D} can be written as $\nu = (\varsigma, \rho, o_1, o_2, \dots, o_M)$ where $\varsigma \in \{1, \dots, M\}$ and $\rho \in \{+, -\}$ are the controlled state components of ν , specifying the cell occupied by the vehicle and the direction of travel, respectively, and for each $i \in \{1, \dots, M\}$, $o_i \in \{0, 1\}$ indicates whether the i^{th} cell is occupied by an obstacle. This leads to the total of $2M2^M$ possible states of \mathbb{D} . With the horizon length $N = 12$, it can be shown that based on the Reachability Problem defined in Section V-A, there is a transition $\nu_1 \rightarrow \nu_2$ in \mathbb{D} if the controlled state components of ν_1 and ν_2 correspond

to adjacent cells (i.e., they share an edge in the road network of Fig. 5).

Since the only progress property is to visit the two cells marked by star infinitely often, the set I_g in (6) has two elements, say, $I_g = \{1, 2\}$. We let \mathcal{W}_0^1 be the set of abstract states whose ς component corresponds to one of these two cells and define \mathcal{W}_0^2 similarly for the other cell as shown in Fig. 5. Other \mathcal{W}_j^i is defined such that it includes all the abstract states whose ς component corresponds to cells across the width of the road (cf. Fig. 5).

Next, we define Φ such that it excludes states where the vehicle is not in the travel lane while there is no obstacle blocking the lane and states where the vehicle is in the same cell as an obstacle or none the cells adjacent to the vehicle are obstacle-free. Using this Φ , the specification in (8) is realizable with $\mathcal{F}^i(\mathcal{W}_j^i) = \mathcal{W}_k^i$ where \mathcal{W}_j^i and \mathcal{W}_k^i correspond to dotted (red) rectangles in Fig. 5 that are two cells apart (e.g. $\mathcal{F}^i(\mathcal{W}_{j+1}^i) = \mathcal{W}_{j-1}^i$). The finite transition system \mathbb{T}^i used by the goal planner can then be constructed such that there is a transition $\mathcal{W}_j^i \rightarrow \mathcal{W}_k^i$ for any \mathcal{W}_j^i and \mathcal{W}_k^i that are two cells apart from each other. With this transition relation, for any $i \in I_g$ and $j \in \{0, \dots, p\}$, there exists a path in \mathbb{T}^i from \mathcal{W}_j^i to \mathcal{W}_0^i and the trajectory planner essentially only has to plan one step ahead. Thus, the size of finite state automata synthesized by the trajectory planner to satisfy the specification in (8) is completely independent of M .

Using JTLV [19], each of these automata has less than 900 states and only takes approximately 1.5 seconds to compute on a MacBook with a 2 GHz Intel Core 2 Duo processor and 4 Gb of memory. In addition, with an efficient graph search algorithm, the computation time requires by the goal generator is in the order of milliseconds. Hence, with a real-time implementation of optimization-based control such as NTG [35] at the continuous controller level, our approach can be potentially implemented in real-time.

C. Results and Discussions

A simulation result is shown in Fig. 6(a), illustrating a correct execution of the vehicle even in the presence of exogenous disturbances when all the assumptions on the environment and initial condition are satisfied. Note that the original synthesis problem requires considering 282 cells. Hence, without the receding horizon strategy, the discretized state space may contains as many as 10^{87} states, making this problem impossible to solve. In fact, JTLV fails to solve this problem with 4 Gb of memory due to an out of memory error. With the receding horizon strategy, only 9 of the 282 cells need to be considered in each synthesis problem. Thus, the number of possible states in the discretized state space reduces from 10^{87} to less than 10^4 , enabling JTLV to easily solve the problem as previously discussed.

To illustrate the benefit of the response mechanism, we add a road blockage on R_2 to violate the assumption (A1). The result is shown in Fig. 6(b). Once the vehicle discovers the road blockage, the trajectory planner cannot find the current state of the system in the finite state automaton synthesized from the specification in (8) since the assumption on the environment is violated. The trajectory planner then informs




Fig. 6. Simulation results with (left) no road blockage, (right) a road blockage on R_2 . The corresponding movies can be downloaded from <http://www.cds.caltech.edu/tulip>




Fig. 7. Simulation result with the presence of disturbances not incorporated in the control protocol synthesis.

the goal generator of the failure to satisfy the corresponding specification with the associated pair of \mathcal{W}_j^i and $\mathcal{F}(\mathcal{W}_j^i)$. Subsequently, the goal generator removes the transition from \mathcal{W}_j^i to $\mathcal{F}(\mathcal{W}_j^i)$ in \mathbb{T}^i and re-computes a path to \mathcal{W}_0^i . As a result, the vehicle continues to exhibit a correct behavior by making a U-turn and completing the task using a different path.

The result with exactly the same setup is also shown in Fig. 7 where exogenous disturbances are not incorporated in the control protocol synthesis. Once the vehicle overtakes the obstacles on R_1 , the continuous controller computes the sequence of control inputs that is expected to bring the vehicle back to its travel lane as commanded by the trajectory planner. But due to the disturbance, the vehicle remains in the opposite lane. In the meantime, the disturbance also causes the vehicle to drift slowly to the right. This cycle continues, leading to violation of the desired property that the vehicle has to stay in the travel lane unless there is an obstacle blocking the lane.

IX. CONCLUSIONS AND FUTURE WORK

Motivated by the DARPA Urban Challenge, we considered the control protocol synthesis problem. Specifically, we proposed an approach to automatically synthesizing a control protocol that ensures system correctness with respect to its specification expressed in linear temporal logic regardless of the environment in which the system operates. A receding-horizon-based framework that allows a computationally complex synthesis problem to be reduced to a set of significantly smaller problems was presented. An implementation of the proposed framework leads to a hierarchical, modular design with a goal generator, a trajectory planner and a continuous controller. A response mechanism that increases the robustness of the system with respect to a mismatch between the system and its model and between the actual behavior of the environment and its assumptions was discussed. By taking into account the presence of exogenous disturbances in the synthesis process, the resulting system is provably robust with respect to bounded exogenous disturbances.

Future work includes further investigation of the robustness of the receding horizon framework. Specifically, we want to formally identify the types of properties and faults/failures

that can be correctly handled using the proposed response mechanism. This type of mechanism has been implemented on the autonomous vehicle built at Caltech for the DARPA Urban Challenge for distributed mission and contingency management [3]. Based on extensive simulations and field tests, it has been shown to handle many types of failures and faults at different levels of the system, including inconsistency of the states of different software modules and hardware and software failures. Another direction of research is to study an asynchronous execution of the goal generator, the trajectory planner and the continuous controller. As described in this paper, these components are to be executed sequentially. However, with certain assumptions on the communication channels, a distributed, asynchronous implementation of these components may still guarantee the correctness of the system. Finally, we want to extend the proposed receding horizon framework to other class of temporal logics that allow better specification of temporal properties.

ACKNOWLEDGMENTS

This work is partially supported by AFOSR under MURI grant FA9550-06-1-0303 and the Boeing Corporation. The authors gratefully acknowledge Hadas Kress-Gazit and Yaniv Sa'ar for inspiring discussions.

REFERENCES

- [1] J. W. Burdick, N. DuToit, A. Howard, C. Looman, J. Ma, R. M. Murray, and T. Wongpiromsarn, "Sensing, navigation and reasoning technologies for the DARPA Urban Challenge," DARPA Urban Challenge Final Report, Tech. Rep., 2007.
- [2] N. E. DuToit, T. Wongpiromsarn, J. W. Burdick, and R. M. Murray, "Situational reasoning for road driving in an urban environment," in *International Workshop on Intelligent Vehicle Control Systems*, 2008.
- [3] T. Wongpiromsarn and R. M. Murray, "Distributed mission and contingency management for the DARPA Urban Challenge," in *International Workshop on Intelligent Vehicle Control Systems*, 2008.
- [4] T. Wongpiromsarn, S. Mitra, R. M. Murray, and A. Lamperski, *Periodically Controlled Hybrid Systems: Verifying A Controller for An Autonomous Vehicle*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2009, vol. 5469, ch. 28, pp. 396–410.
- [5] Z. Manna and A. Pnueli, *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, 1992.
- [6] M. Huth and M. Ryan, *Logic in Computer Science: Modelling and Reasoning about Systems*, 2nd ed. Cambridge University Press, 2004.
- [7] E. A. Emerson, "Temporal and modal logic," in *Handbook of theoretical computer science (vol. B): formal models and semantics*. Cambridge, MA, USA: MIT Press, 1990, pp. 995–1072.
- [8] S. Karaman and E. Frazzoli, "Sampling-based motion planning with deterministic μ -calculus specifications," in *Proc. of IEEE Conference on Decision and Control*, 2009.
- [9] H. Kress-Gazit, G. Fainekos, and G. Pappas, "Where's Waldo? Sensor-based temporal logic motion planning," in *Proc. of IEEE International Conference on Robotics and Automation*, Apr 2007, pp. 3116–3121.
- [10] D. Conner, H. Kress-Gazit, H. Choset, A. Rizzi, and G. Pappas, "Valet parking without a valet," in *Proc. of IEEE/RSS International Conference on Intelligent Robots and Systems*, 2007, pp. 572–577.
- [11] 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, Feb 2008.
- [12] P. Tabuada and G. J. Pappas, "Linear time logic control of discrete-time linear systems," *IEEE Transactions on Automatic Control*, vol. 51, no. 12, pp. 1862–1877, Dec 2006.
- [13] A. Girard and G. J. Pappas, "Hierarchical control system design using approximate simulation," *Automatica*, vol. 45, no. 2, pp. 566–571, Feb 2009.
- [14] T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Receding horizon temporal logic planning for dynamical systems," in *Proc. of IEEE Conference on Decision and Control*, 2009.

- [15] —, “Automatic synthesis of robust embedded control software,” in *AAAI Spring Symposium on Embedded Reasoning: Intelligence in Embedded Systems*, 2010, pp. 104–111.
- [16] —, “Receding horizon control for temporal logic specifications,” in *Proc. of the 13th International Conference on Hybrid Systems: Computation and Control*, 2010.
- [17] R. Alur, T. A. Henzinger, G. Lafferriere, and G. J. Pappas, “Discrete abstractions of hybrid systems,” *Proceedings of the IEEE*, vol. 88, no. 7, pp. 971–984, Jul 2000.
- [18] A. Girard, A. A. Julius, and G. J. Pappas, “Approximate simulation relations for hybrid systems,” *Discrete Event Dynamic Systems*, vol. 18, no. 2, pp. 163–179, Jun 2008.
- [19] N. Piterman, A. Pnueli, and Y. Sa’ar, *Synthesis of Reactive(1) Designs*, ser. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2005, vol. 3855, ch. 24, pp. 364–380, software available at <http://jtlv.sourceforge.net/>.
- [20] G. Pola and P. Tabuada, “Symbolic models for nonlinear control systems: Alternating approximate bisimulations,” *SIAM Journal on Control and Optimization*, vol. 48, no. 2, pp. 719–733, 2009.
- [21] D. Limon, T. Alamo, and E. Camacho, “Enlarging the domain of attraction of MPC controllers,” *Automatica*, vol. 41, no. 4, pp. 629–635, Apr 2005.
- [22] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, “Constrained model predictive control: Stability and optimality,” *Automatica*, vol. 36, pp. 789–814, 2000.
- [23] H. Tanner and G. J. Pappas, “Simulation relations for discrete-time linear systems,” in *Proc. of the IFAC World Congress on Automatic Control*, 2002, pp. 1302–1307.
- [24] M. Kvasnica, P. Grieder, and M. Baotić, “Multi-Parametric Toolbox (MPT),” 2004, software available at <http://control.ee.ethz.ch/~mpt>.
- [25] J. Löfberg, “YALMIP : A toolbox for modeling and optimization in MATLAB,” in *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004, software available at <http://control.ee.ethz.ch/~jloef/yalmip.php>.
- [26] R. M. Murray, J. Hauser, A. Jadbabaie, M. B. Milam, N. Petit, W. B. Dunbar, and R. Franz, “Online control customization via optimization-based control,” in *Software-Enabled Control: Information Technology for Dynamical Systems*. Wiley-Interscience, 2002, pp. 149–174, software available at http://www.cds.caltech.edu/~murray/software/2002a_ntg.html.
- [27] F. Borrelli, *Constrained Optimal Control of Linear and Hybrid Systems*, ser. Lecture Notes in Control and Information Sciences. Springer-Verlag, 2003, vol. 290.
- [28] D. Peled and T. Wilke, “Stutter-invariant temporal properties are expressible without the next-time operator,” *Information Processing Letters*, vol. 63, no. 5, pp. 243–246, Sep 1997.
- [29] A. Jadbabaie, “Nonlinear receding horizon control: A control Lyapunov function approach,” Ph.D. dissertation, California Institute of Technology, 2000.
- [30] G. C. Goodwin, M. M. Seron, and J. A. D. Doná, *Constrained Control and Estimation: An Optimisation Approach*. Springer, 2004.
- [31] T. Wongpiromsarn, U. Topcu, N. Ozay, H. Xu, and R. M. Murray, “TuLiP: A software toolbox for receding horizon temporal logic planning,” in *International Conference on Hybrid Systems: Computation and Control*, 2011.
- [32] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Prentice-Hall, 2003.
- [33] H. Kress-Gazit and G. J. Pappas, “Automatically synthesizing a planning and control subsystem for the DARPA Urban Challenge,” in *IEEE International Conference on Automation Science and Engineering*, 2008, pp. 766–771.
- [34] T. Wongpiromsarn, “Formal methods for embedded control systems: Application to an autonomous vehicle,” Ph.D. dissertation, California Institute of Technology, 2010.
- [35] M. B. Milam, K. Mushambi, and R. M. Murray, “A new computational approach to real-time trajectory generation for constrained mechanical systems,” in *Proc. of IEEE Conference on Decision and Control*, 2000, pp. 845–851.



Tichakorn Wongpiromsarn received the B.S. degree in mechanical engineering from Cornell University and the M.S. and Ph.D. degrees in mechanical engineering from California Institute of Technology. She is currently a Postdoctoral Associate at the Singapore-MIT Alliance for Research and Technology. Her research interests span hybrid systems, distributed control systems, formal methods, transportation networks and situational reasoning and decision making in complex, dynamic and uncertain environments.



Ufuk Topcu received his Ph.D. degree in 2008 from the University of California, Berkeley. He currently is a postdoctoral scholar of Control and Dynamical Systems at the California Institute of Technology. His research is on the analysis, design, and verification of networked, information-based systems with current projects in autonomy, advanced air vehicle architectures, and energy networks.



Richard M. Murray received the B.S. degree in Electrical Engineering from California Institute of Technology in 1985 and the M.S. and Ph.D. degrees in Electrical Engineering and Computer Sciences from the University of California, Berkeley, in 1988 and 1991, respectively. He is currently the Thomas E. and Doris Everhart Professor of Control & Dynamical Systems and Bioengineering at Caltech. Murray’s research is in the application of feedback and control to networked systems, with applications in biology and autonomy. Current projects include verification and validation of distributed embedded systems, analysis of insect flight control systems, and biological circuit design.