Receding Horizon Temporal Logic Control for Finite Deterministic Systems

Xuchu Ding, Mircea Lazar and Calin Belta

Abstract— This paper considers receding horizon control of finite deterministic systems, which must satisfy a high level, rich specification expressed as a linear temporal logic formula. Under the assumption that time-varying rewards are associated with states of the system and they can be observed in real-time, the control objective is to maximize the collected reward while satisfying the high level task specification. In order to properly react to the changing rewards, a controller synthesis framework inspired by model predictive control is proposed, where the rewards are locally optimized at each time-step over a finite horizon, and the immediate optimal control is applied. By enforcing appropriate constraints, the infinite trajectory produced by the controller is guaranteed to satisfy the desired temporal logic formula. Simulation results demonstrate the effectiveness of the approach.

I. INTRODUCTION

This paper considers the problem of controlling a deterministic discrete-time system with a finite state-space, which is also referred to as a finite transition system. Such systems can be effectively used to capture behaviors of more complex dynamical systems, and as a result, greatly reduce the complexity of control design.

A finite transition system can be constructed from a continuous system via an "abstraction" process. For example, for an autonomous robotic vehicle moving in an environment, the motion of the vehicle can be abstracted to a finite system through a partition of the environment. The set of states can be seen as a set of labels for the regions in the partition, and each transition corresponds to a controller driving the vehicle between two adjacent regions. By partitioning the environment into simplicial, rectangular or polyhedral regions, continuous feedback controllers that drive a robotic system from any point inside a region to a desired facet of an adjacent region have been developed for linear [1], multi-affine [2], piecewise-affine [3]–[5], and non-holonomic (unicycle) [6], [7] dynamical models. By relating the initial continuous dynamical system and the abstract discrete finite system with simulation or bisimulation relations [8], the abstraction process allows one to solve a control problem for the more complex continuous system with the "equivalent" abstract system.

It has been proposed by several authors [1], [5], [9]–[11] to use temporal logics, such as linear temporal logic (LTL) and computation tree logic (CTL) [12], as specification

languages for finite transition systems due to their well defined syntax and semantics. These logics can be easily used to specify complex behavior, and in particular with LTL, persistent mission task such as "pick up items at the region `pickup`, and then drop them off at the region `dropoff`, infinitely often, while always avoiding `unsafe` regions". The applications of these temporal logics in computer science in the area of model checking [12] and temporal logic game [13] has resulted in off-the-shelf tools and algorithms that can be readily adapted to synthesize provably correct control strategies [1], [5], [9]–[11].

While the works mentioned above address the temporal logic controller synthesis problem, several problems and questions remain to be answered. In particular, the problem of combining temporal logic controller synthesis with optimality with respect to a suitable cost function remains to be solved. This problem becomes even more difficult if the optimization problem depends on time-varying parameters, *e.g.,* dynamic events that occur during the operation of the plant. For traditional control problems (without temporal logic constraints) and dynamical systems, this problem can be effectively addressed using a model predictive control (MPC) approach (see *e.g.,* [14]), which has reached a mature level in both academia and industry, with many successful implementations. The basic MPC set-up consists of the following sequence of steps: at each time instant, a cost function of the current state is optimized over a finite horizon, only the first element of the optimal finite sequence of controls is applied and the whole process is repeated at the next time instant for the new measured state. Thus, MPC is also referred to as receding horizon control. Since the finite horizon optimization problem is solved repeatedly at each time instant, real-time dynamical events can be effectively managed.

However, it is not yet well-understood how to combine a receding horizon control approach with a provably correct control strategy satisfying a temporal logic formula. The aim of this paper is to address this issue for a specific system set-up (deterministic systems on a finite state-space) and problem formulation (dynamic optimization of rewards). More specifically, the role of the receding horizon controller is to maximize over a finite horizon the accumulated rewards associated with states of the system, under the assumption that the rewards change dynamically with time and they can only be observed in real-time. The rewards model dynamical events that can be triggered in real-time, which is an often used model in coverage control literature [15].

The key challenge in this controller synthesis framework is to ensure correctness of the produced infinite trajectory

and recursive feasibility of the optimization problem solved at each time-step. For a constrained MPC optimization problem, which is solved recursively on-line, feasible at all times or recursively feasible means that if the optimization problem is feasible (has a solution) for the initial state at initial time, then it remains feasible for all future time instants, when it will be solved with a different initial condition resulting from the generated closed-loop trajectory. A proof that the proposed receding horizon control framework satisfies both properties is provided. Similar to standard MPC, where certain *terminal* constraints must be enforced in the optimization problem in order to guarantee certain properties for the system (*e.g.,* stability), the correctness of produced trajectory and recursive feasibility are also ensured via a set of suitable constraints.

This work can be seen as an extension and generalization of the set-up presented in [16], where a similar control objective was tackled. In [16] an optimization based controller was designed, which consists of repeatedly solving a finite horizon optimal control problem every N steps and implementing the complete sequence of control actions. This procedure is more close to finite-horizon optimal control than true receding horizon control and its main drawback comes from the inability of reacting to dynamical events (*i.e.,* rewards) triggered or varying during the execution of the finite trajectory. This paper removes this limitation by attaining a truly receding horizon controller for deterministic systems on a finite state-space. Another related work is [5], where a provably correct control strategy was obtained for large scale systems by dividing the control synthesis problem into smaller sub-problems in a receding horizon like manner. However, in [5] dynamical events were addressed differently and the specification language was restricted to a fragment of LTL, whereas in this paper full LTL expressivity is allowed.

II. PROBLEM FORMULATION AND APPROACH

In this paper, we consider a discrete-time system with a finite state space, *i.e.,* the system evolves on a graph. Each vertex of the graph produces an output, which is a set of observations. Such a system can be described by a finite deterministic transition system, which can be formally defined as follows.

Definition II.1 (Finite Deterministic Transition System)**.** *A finite (weighted) deterministic transition system (DTS) is a tuple $\mathcal{T} = (Q, q_0, \Delta, \omega, \Pi, h)$, where*

- *Q is a finite set of states;*
- *$q_0 \in Q$ is the initial state;*
- *$\Delta \subseteq Q \times Q$ is the set of transitions;*
- *$\omega : \Delta \to \mathbb{R}^+$ is a weight function that assigns positive values to all transitions;*
- *Π is a set of observations; and*
- *$h : Q \to 2^\Pi$ is the observation map.*

For convenience of notation, we denote $q \to_\mathcal{T} q'$ if $(q, q') \in \Delta$. We assume $\mathcal{T}$ to be non-blocking, i.e., for each $q \in Q$, there exists $q' \in Q$ such that $q \to_\mathcal{T} q'$ (such a system is also called a Kripke structure [17]). A trajectory of a DTS is an infinite sequence $\mathbf{q} = q_0 q_1...$ where $q_k \to_\mathcal{T} q_{k+1}$ for

all $k \geq 0$. A trajectory $\mathbf{q}$ generates an output trajectory $\mathbf{o} = o_0 o_1...$, where $o_k = h(q_k)$ for all $k \geq 0$.

Note the absence of the control inputs in the definition of $\mathcal{T}$. This is because $\mathcal{T}$ is deterministic, and one can choose an available transitions at a state. In other words, each transition (q, q') corresponds to a unique control input at state q. This also implies that a trajectory $\mathbf{q} = q_0 q_1 \ldots$ can be used as a control strategy for $\mathcal{T}$, by simply applying the transitions $(q_0, q_1), (q_1, q_2)$, and so on. An example of a DTS is shown in Fig. 1.
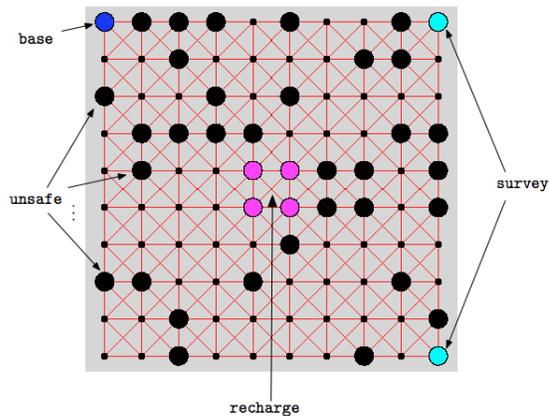

Fig. 1. An example of a finite DTS $\mathcal{T}$ as defined in Def. II.1. In this example, $\mathcal{T}$ has 100 states, which are at the vertices of a rectangular grid with cell size 10. We define the weight function ω to be the Euclidean distance between vertices, and there is a transition between two vertices if the Euclidean distance between them is less than 15. The set of observations is $\Pi = \{$base, survey, recharge, unsafe$\}$. States with no observation are shown with smaller vertices.

The goal of this paper is to synthesize trajectories $\mathbf{q}$ of $\mathcal{T}$ satisfying a behavioral specification, given as a linear temporal logic formula over Π. An LTL formula over Π is interpreted over an (infinite) sequence $\mathbf{o} = o_0 o_1 \ldots$, where $o_k \subseteq \Pi$ for all $k \geq 0$. We say $\mathbf{q}$ satisfies an LTL formula ϕ if it generates an output trajectory $\mathbf{o}$ satisfying ϕ. A detailed description of the syntax and semantics of LTL is beyond the scope of this paper and can be found in [12]. Roughly, an LTL formula is build up from the observations in Π, Boolean operators $\neg$ (negation), $\vee$ (disjunction), $\wedge$ (conjunction), $\longrightarrow$ (implication), and temporal operators X (next), U (until), F (eventually), G (always). For example, the following task command in natural language: *"Reach a* survey *location infinitely often, and always avoid* unsafe *states"* can be translated to the LTL formula: $\phi := \mathsf{G}\,\mathsf{F}\,$survey$\,\wedge\,\mathsf{G}\,\neg$unsafe.

The system is assumed to operate in an environment with dynamical events. In this paper, these events are modelled by a reward process $\mathcal{R} : Q \times \mathbb{N} \to \mathbb{R}^+$, *i.e.,* the reward associated with state $q \in Q$ at time k is $\mathcal{R}(q, k)$. Note that rewards are associated with states in Q in a time varying fashion. We do not make any assumptions on the dynamics governing the rewards, but we make the natural assumption that, at time k, the system can only observe the rewards in a neighborhood $\mathcal{N}(q, k) \subseteq Q$ of the current state q. In this paper, we assume that the reward process $\mathcal{R}$ is unknown and

reward values must be observed and acted upon in real-time. The problem in the case when knowledge of $\mathcal{R}$ is given a-priori is interesting, but will be addressed in future research.

The problem considered in this paper is formally stated next.

Problem II.2. *Given a transition system $\mathcal{T}$ and an LTL formula ϕ over the set of observations of $\mathcal{T}$, design a controller that maximizes the collected reward locally, while it ensures that the produced infinite trajectory satisfies ϕ.*

Since the rewards are time-varying and can only be observed around the current state, inspirations from the area of MPC are drawn (see, e.g. [14]) with the aim of synthesizing a controller such that the rewards are maximized in a receding horizon fashion. At time k with state q_k, the controller generates a finite trajectory $q_{k+1}q_{k+2}\dots q_{k+N}$ by solving an on-line optimization problem maximizing the collected rewards over a horizon N, and the system implements the immediate control action (q_k, q_{k+1}). This process is then repeated at time $k+1$ and state q_{k+1}.

In order to guarantee the satisfaction condition for the LTL formula ϕ, the proposed approach is based on the construction of an automaton that captures all satisfying trajectories of $\mathcal{T}$. This automaton also induces a Lyapunov-like function that can be used to enforce that the trajectory of the system satisfies the desired formula. These steps are formally described in detail in Sec. III. The aforementioned function will be utilized to guarantee recursive feasibility of the developed receding horizon controller, which in turn will yield that the synthesized infinite trajectory satisfies ϕ. The controller synthesis method is presented in Sec. IV.

III. A TOOL FOR ENFORCING THE BÜCHI ACCEPTANCE CONDITION

In this section, we review the definition of Büchi automata and describe the construction of a function that enforces the satisfaction of a Büchi acceptance condition for the trajectories of a DTS.

Definition III.1 (Büchi Automaton). *A (nondeterministic) Büchi automaton is a tuple $\mathcal{B} = (S_{\mathcal{B}}, S_{\mathcal{B}0}, \Sigma, \delta, F_{\mathcal{B}})$, where*

- *$S_{\mathcal{B}}$ is a finite set of states;*
- *$S_{\mathcal{B}0} \subseteq S_{\mathcal{B}}$ is the set of initial states;*
- *Σ is the input alphabet;*
- *$\delta : S_{\mathcal{B}} \times \Sigma \to 2^{S_{\mathcal{B}}}$ is the transition function;*
- *$F_{\mathcal{B}} \subseteq S$ is the set of accepting states.*

We denote $s \xrightarrow{\sigma}_{\mathcal{B}} s'$ if $s' \in \delta(s, \sigma)$. An infinite sequence $\sigma_0\sigma_1 \dots$ over Σ generates trajectories $s_0 s_1 \dots$ where $s_0 \in S_{\mathcal{B}0}$ and $s_k \xrightarrow{\sigma_k}_{\mathcal{B}} s_{k+1}$ for all $k \geq 0$. $\mathcal{B}$ accepts an infinite sequence over Σ if it generates at least one trajectory on $\mathcal{B}$, which intersects the set $F_{\mathcal{B}}$ infinitely many times.

For any LTL formula ϕ over Π, one can construct a Büchi automaton with input alphabet $\Sigma = 2^{\Pi}$ accepting all and only sequences over 2^{Π} that satisfy ϕ [12]. We refer readers to [18] for efficient algorithms and implementations to translate an LTL formula over Π to a corresponding Büchi automaton $\mathcal{B}$.

Definition III.2 (Weighted Product Automaton). *Given a weighted DTS $\mathcal{T} = (Q, q_0, \Delta, \omega, \Pi, h)$ and a Büchi automaton $\mathcal{B} = (S_{\mathcal{B}}, S_{\mathcal{B}0}, 2^{\Pi}, \delta_{\mathcal{B}}, F_{\mathcal{B}})$, their product automaton, denoted by $\mathcal{P} = \mathcal{T} \times \mathcal{B}$, is a tuple $\mathcal{P} = (S_{\mathcal{P}}, S_{\mathcal{P}0}, \Delta_{\mathcal{P}}, \omega_{\mathcal{P}}, F_{\mathcal{P}})$ where*

- *$S_{\mathcal{P}} = Q \times S_{\mathcal{B}}$;*
- *$S_{\mathcal{P}0} = \{q_0\} \times S_{\mathcal{B}0}$;*
- *$\Delta_{\mathcal{P}} \subseteq S_{\mathcal{P}} \times S_{\mathcal{P}}$ is the set of transitions, defined by: $((q,s),(q',s')) \in \Delta_{\mathcal{P}}$ iff $q \to_{\mathcal{T}} q'$ and $s \xrightarrow{h(q)}_{\mathcal{B}} s'$;*
- *$\omega_{\mathcal{P}} : \Delta_{\mathcal{P}} \to \mathbb{R}^+$ is the weight function defined by: $\omega_{\mathcal{P}}((q,s),(q',s')) = \omega((q,q'))$*
- *$F_{\mathcal{P}} = Q \times F_{\mathcal{B}}$.*

We denote $(q,s) \to_{\mathcal{P}} (q',s')$ if $((q,s),(q',s')) \in \Delta_{\mathcal{P}}$. A trajectory $\boldsymbol{p} = (q_0, s_0)(q_1, s_1) \dots$ of $\mathcal{P}$ is an infinite sequence such that $(q_0, s_0) \in S_{\mathcal{P}0}$ and $(q_k, s_k) \to_{\mathcal{P}} (q_{k+1}, s_{k+1})$ for all $k \geq 0$. Trajectory $\boldsymbol{p}$ is called accepting if and only if it intersects $F_{\mathcal{P}}$ infinitely many times.

We define the projection $\gamma_{\mathcal{T}}$ of $\boldsymbol{p}$ onto $\mathcal{T}$ as simply removing the automaton states, *i.e.*,

$$\gamma_{\mathcal{T}}(\boldsymbol{p}) = \boldsymbol{q} = q_0 q_1 \dots, \text{ if } \boldsymbol{p} = (q_0, s_0)(q_1, s_1) \dots. \quad (1)$$

We also use the projection operator $\gamma_{\mathcal{T}}$ for finite trajectories (subsequences of $\boldsymbol{p}$). Note that a trajectory $\boldsymbol{p}$ on $\mathcal{P}$ is uniquely projected to a trajectory $\gamma_{\mathcal{T}}(\boldsymbol{p})$ on $\mathcal{T}$. By the construction of $\mathcal{P}$ from $\mathcal{T}$ and $\mathcal{B}$, $\boldsymbol{p}$ is accepted if and only if $\boldsymbol{q} = \gamma_{\mathcal{T}}(\boldsymbol{p})$ satisfies the LTL formula corresponding to $\mathcal{B}$ [12].

In [16], we introduced a real positive function V on the states of the product automaton $\mathcal{P}$ that uses the weights $\omega_{\mathcal{P}}$ to enforce the acceptance condition of the automaton. Conceptually, this function resembles a Lyapunov, or energy function. While in Lyapunov theory energy functions are used to enforce that the trajectories of a dynamical system converge to an equilibrium, this "energy" function enforces that the trajectories of $\mathcal{T}$ satisfy the acceptance condition of a Büchi automaton.

To define the energy function, we first denote a set $A \subseteq S_{\mathcal{P}}$ to be *self-reachable* if and only if all states in A can reach at least one state in A.

Definition III.3 (Energy function of a state in $\mathcal{P}$). *We define $F_{\mathcal{P}}^{\star}$ to be the largest self-reachable subset of $F_{\mathcal{P}}$. The energy function $V(p)$, $p \in S_{\mathcal{P}}$ is defined as the graph distance of p to the set $F_{\mathcal{P}}^{\star}$, i.e., the accumulated weight of the shortest path from p to any states in $F_{\mathcal{P}}^{\star}$.*

Fig. 2 shows an example of $\mathcal{T}$, $\mathcal{B}$, and their product $\mathcal{P}$, as well as the induced energy function defined on states of $\mathcal{P}$. In [16], we showed the following properties for V.

Theorem III.4 (Properties of the energy function). *V satisfies the following:*

(i) If a trajectory $\boldsymbol{p}$ on $\mathcal{P}$ is accepting, then it cannot contain a state p where $V(p) = \infty$.

(ii) All accepting states in an accepting trajectory $\boldsymbol{p}$ are in the set $F_{\mathcal{P}}^{\star}$ and have energy equal to 0; all accepting states that are not in $F_{\mathcal{P}}^{\star}$ have energy equal to ∞.

(iii) For each state $p \in S_{\mathcal{P}}$, if $V(p) > 0$ and $V(p) \neq \infty$, then there exists a state p' where $p \to_{\mathcal{P}} p'$ such that $V(p') < V(p)$.
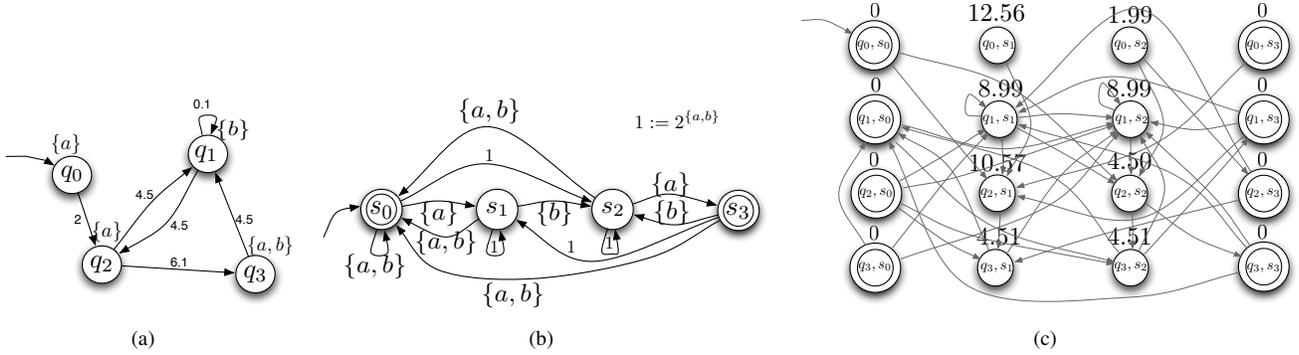
Fig. 2. The construction of the product automaton and the energy function on its states. In this example, the set of observations is $\Pi = \{a, b\}$. The initial states are indicated by incoming arrows. The accepting states are marked by double-strokes. **(a)**: A weighted DTS $\mathcal{T}$. The label atop each state indicates the set of associated observations. (*i.e.,* $\{a, b\}$ means both a and b are observed). The labels on the transitions indicate the weights. **(b)**: The Büchi automaton $\mathcal{B}$ corresponds to LTL formula $\mathsf{G}\,(\mathsf{F}\,(a \wedge \mathsf{F}\,b))$, translated by the tool LTL2BA [18]. **(c)**: The product automaton $\mathcal{P} = \mathcal{T} \times \mathcal{B}$ constructed with Def. III.2 (the weights are inherited from $\mathcal{T}$ and not shown). The number above a state $p \in S_{\mathcal{P}}$ is the energy function $V(p)$. Note that in this example, the set $F_{\mathcal{P}}^{\star} = F_{\mathcal{P}}$, thus $V(p)$ is the graph distance from p to any accepting states.

We see from the above facts that $V(p)$ resembles an energy-like function, which justifies the name we use. We refer to the value of $V(p)$ at a state $p \in S_{\mathcal{P}}$ as the "energy of the state". Note that satisfying the LTL formula is equivalent to reaching states where $V(p) = 0$ for infinitely many times. Therefore, for each state $p \in S_{\mathcal{P}}$, $V(p)$ provides a measure of progress towards satisfying the LTL formula. An algorithm generating $V(p)$ for an arbitrary product automaton can be found in [16].

IV. MAIN RESULTS: RECEDING HORIZON CONTROLLER DESIGN

In this section, we present a solution to Prob. II.2. The central component of our control design is a state-feedback controller operating on the product automaton that optimizes finite trajectories over a pre-determined, fixed horizon N, subject to certain constraints. These constraints ensure that the energy of states on the product automaton decreases in finite time, thus guaranteeing that progress is made towards the satisfaction of the LTL formula. Note that the proposed controller does not enforce the energy to decrease at each time-step, but rather that it eventually decreases. The finite trajectory returned by the receding horizon controller is projected onto $\mathcal{T}$, the controller applies the first transition, and this process is repeated again at the next time-step.

In this section, we first describe the receding horizon controller and show that it is feasible (a solution exists) at all time-steps $k \in \mathbb{N}$. Then, we present the general control algorithm and show that it always produces (infinite) trajectories satisfying the given LTL formula.

A. Receding horizon controller

In order to explain the working principle of the controller, we first define a finite *predicted trajectory* on $\mathcal{P}$ at time k. Denote the current state at time k as p_k. A predicted trajectory of horizon N at time k is a finite sequence $\mathbf{p}_k := p_{1|k} \ldots p_{N|k}$, where $p_{i|k} \in S_{\mathcal{P}}$ for all $i = 1, \ldots, N$, $p_{i|k} \rightarrow_{\mathcal{P}} p_{i+1|k}$ for all $i = 1, \ldots N-1$ and $p_k \rightarrow_{\mathcal{P}} p_{1|k}$. Here, $p_{i|k}$ is a

notation used frequently in MPC, which denotes the ith state of the predicted trajectory at time k. Moreover, we denote the set $\mathbf{P}(p_k, N)$ as the set of all finite trajectories of horizon N from a state $p_k \in S_{\mathcal{P}}$. Note that the finite predicted trajectory $\mathbf{p}_k$ of $\mathcal{P}$ uniquely projects to a finite trajectory $\mathbf{q}_k := \gamma_{\mathcal{T}}(\mathbf{p}_k)$ of $\mathcal{T}$.

For the current state q_k at time k, we denote the observed reward at any state $q \in Q$ as $R_k(q)$, and we have that

$$R_k(q) = \begin{cases} \mathcal{R}(q, k) & \text{if } q \in \mathcal{N}(q_k, k) \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Note that $\mathcal{R}(q, k) = 0$ if $q \notin \mathcal{N}(q_k, k)$ because the rewards outside of the neighbourhood cannot be observed. We can now define the *predicted reward* associated with a predicted trajectory $\mathbf{p}_k \in \mathbf{P}(p_k, N)$ at time k. The predicted reward of $\mathbf{p}_k$, denoted as $\Re_k(\mathbf{p}_k)$, is simply the amount of accumulated rewards by $\gamma_{\mathcal{T}}(\mathbf{p}_k)$ of $\mathcal{T}$:

$$\Re_k(\mathbf{p}_k) = \sum_{i=1}^{N} R_k\left(\gamma_{\mathcal{T}}(p_{i|k})\right). \quad (3)$$

The receding horizon controller executed at the initial state at time $k = 0$ is described next. This is a special case because the initial state of $\mathcal{P}$ is not unique, and as a result we can pick any initial state of $\mathcal{P}$ from the set $S_{\mathcal{P}0} = \{q_0\} \times S_{\mathcal{B}0}$. We denote the controller executed at the initial state as $\mathrm{RH}^0(S_{\mathcal{P}0})$, and we define it as follows

$$\begin{aligned} \mathbf{p}_0^{\star} &= \mathrm{RH}^0(S_{\mathcal{P}0}) \\ &:= \underset{\mathbf{p}_0 \in \{\mathbf{P}(p_0, N) \,|\, V(p_0) < \infty\}}{\arg\max} \Re_0(\mathbf{p}_0). \end{aligned} \quad (4)$$

The controller maximizes the predicted cumulative rewards over all possible projected trajectories over horizon N initiated from a state $p_0 \in S_{\mathcal{P}0}$ where the energy is finite, and returns the optimal projected trajectory $\mathbf{p}_0^{\star}$. The requirement that $V(p_0) < \infty$ is critical because otherwise, the trajectory starting from p_0 cannot be accepting. If there does not exist

p_0 such that $V(p_0) < \infty$, then an accepting trajectory does not exist and there is no trajectory of $\mathcal{T}$ satisfying the LTL formula (*i.e.,* Prob. II.2 has no solution).

Lemma IV.1 (Feasiblity of (4)). *Optimization problem* (4) *always has at least one solution if there exists p_0 such that $V(p_0) < \infty$.*

Proof. The proof follows from the fact that $\mathcal{T}$ is non-blocking, and thus the set $\mathbf{P}(p_0, N)$ is not empty. ∎

Next, the receding horizon control algorithm for any time instant $k = 1, 2, \ldots$ and corresponding state $p_k \in S_{\mathcal{P}}$ is presented. This controller is of the form

$$\mathbf{p}_k^\star = \text{RH}(p_k, \mathbf{p}_{k-1}^\star) \tag{5}$$

i.e., it depends both on the current state p_k and the optimal predicted trajectory $\mathbf{p}_{k-1}^\star = p_{1|k-1}^\star \ldots p_{N|k-1}^\star$ obtained at the previous time-step. Note that, by the nature of a receding horizon control scheme, the first control of the previous predicted trajectory is always applied. Therefore, we have the following equality

$$p_k = p_{1|k-1}^\star, \quad k = 1, 2, \ldots. \tag{6}$$

As it will become clear in the text below, $\mathbf{p}_{k-1}^\star$ is used to enforce repeated executions of this controller to eventually reduce the energy of the state on $\mathcal{P}$ to 0.

We define controller (5) with the following three cases:

1) Case 1. $V(p_k) > 0$ and $V(p_{i|k-1}^\star) \neq 0$ for all $i = 1, \ldots, N$: In this case, the receding horizon controller is defined as follows.

$$
\begin{aligned}
\mathbf{p}_k^\star &= \text{RH}(p_k, \mathbf{p}_{k-1}^\star) \\
&:= \underset{\mathbf{p}_k \in \mathbf{P}(p_k, N)}{\arg\max} \, \Re_k(\mathbf{p}_k), \\
&\quad \text{subject to: } V(p_{N|k}) < V(p_{N|k-1}^\star).
\end{aligned} \tag{7}
$$

The key to guarantee that the energy of the states on $\mathcal{P}$ eventually decreases is the terminal constraint $V(p_{N|k}) < V(p_{N|k-1}^\star)$, *i.e.,* the optimal finite predicted trajectory $p_k^\star$ must end at a state with lower energy than that of the previous predicted trajectory $p_{k-1}^\star$. This terminal constraint mechanism is graphically illustrated in Fig. 3.
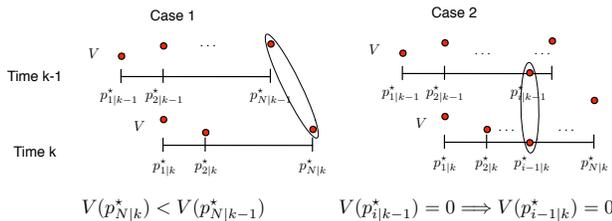

Fig. 3. Constraints enforced for the receding horizon control law $\mathbf{p}_k^\star = \text{RH}(p_k, \mathbf{p}_{k-1}^\star)$ for Cases 1 and 2.

To verify the feasibility of the optimization problem under this constraint, we make use of the third property of V in Thm. III.4. Namely, each state with positive finite energy can make a transition to a state with strictly lower energy.

Lemma IV.2 (Feasibility of (7)). *Optimization problem* (7) *always has at least one solution if $V(p_k) < \infty$.*

Proof. Given $\mathbf{p}_{k-1}^\star = p_{1|k-1}^\star \ldots p_{N|k-1}^\star$, since $p_k = p_{1|k-1}^\star$, we have $p_k \to_{\mathcal{P}} p_{2|k-1}^\star$. Therefore, we can construct a finite predicted trajectory $\mathbf{p}_k = p_{1|k} \ldots p_{N|k}$ where $p_{i|k} = p_{i+1|k-1}^\star$ for all $i = 1, \ldots, N-1$. Using Thm. III.4 (iii), there exists a state p where $p_{N-1|k} \to_{\mathcal{P}} p$ such that $V(p) < V(p_{N-1|k})$. Setting $p_{N|k} = p$, the finite trajectory $\mathbf{p}_k = p_{1|k} \ldots p_{N|k} \in \mathbf{P}(p_k, N)$ satisfies the constraint $V(p_{N|k}) < V(p_{N|k-1}^\star)$, and therefore (7) has at least one solution. ∎

2) Case 2. $V(p_k) > 0$ and there exists $i \in \{1, \ldots, N\}$ with $V(p_{i|k-1}^\star) = 0$: We denote $i^0(\mathbf{p}_{k-1}^\star)$ as the index of the first occurrence in $\mathbf{p}_{k-1}^\star$ where the energy is 0, *i.e.,* $V(p_{i^0(\mathbf{p}_{k-1}^\star)|k-1}^\star) = 0$. We then propose the following controller.

$$
\begin{aligned}
\mathbf{p}_k^\star &= \text{RH}(p_k, \mathbf{p}_{k-1}^\star) \\
&:= \underset{\mathbf{p}_k \in \mathbf{P}(p_k, N)}{\arg\max} \, \Re_k(\mathbf{p}_k), \\
&\quad \text{subject to: } V(p_{i^0(\mathbf{p}_{k-1}^\star)-1|k}^\star) = 0.
\end{aligned} \tag{8}
$$

Namely, this controller enforces a state in the optimal predicted trajectory to have 0 energy if the previous predicted trajectory contains such a state. This constraint is illustrated in Fig. 3. Note that, if $i^0(\mathbf{p}_{k-1}^\star) = 1$, then from (6), the current state p_k is such that $V(p_k) = 0$, and Case 2 does not apply but Case 3 (described below) applies instead.

Lemma IV.3 (Feasibility of (8)). *Optimization problem* (8) *always has at least one solution if $V(p_k) < \infty$.*

Proof. Given $\mathbf{p}_{k-1}^\star = p_{1|k-1}^\star \ldots p_{N|k-1}^\star$, since $p_k = p_{1|k-1}^\star$, we have $p_k \to_{\mathcal{P}} p_{2|k-1}^\star$. Therefore, we can construct a finite predicted trajectory $\mathbf{p}_k = p_{1|k} \ldots p_{N|k}$ where $p_{i|k} = p_{i+1|k-1}^\star$ for all $i = 1, \ldots, N-1$. If we let $p_{N|k}$ to be any state where $p_{N-1|k} \to_{\mathcal{P}} p_{N|k}$ and $V(p_{N|k}) < \infty$, then $\mathbf{p}_k = p_{1|k} \ldots p_{N|k} \in \mathbf{P}(p_k, N)$ satisfies the constraint. Thm. III.4 (iii) gurantees that such a state $p_{N|k}$ exists. ∎

3) Case 3, $V(p_k) = 0$: In this case, the terminal constraint is that energy value of the terminal state is finite. The controller is defined as follows.

$$
\begin{aligned}
\mathbf{p}_k^\star &= \text{RH}(p_k, \mathbf{p}_{k-1}^\star) \\
&:= \underset{\mathbf{p}_k \in \mathbf{P}(p_k, N)}{\arg\max} \, \Re_k(\mathbf{p}_k). \\
&\quad \text{subject to: } V(p_{N|k}) < \infty.
\end{aligned} \tag{9}
$$

Lemma IV.4 (Feasiblity of (9)). *Optimization problem* (9) *always has at least one solution.*

Proof. If $V(p_k) = 0$, then there exists $p_{1|k}$ such that $p_k \to_{\mathcal{P}} p_{1|k}$ and $V(p_{1|k}) < \infty$ (if not, then $V(p_k)$ must equal to ∞). From Thm. III.4 (iii), we have that there exists $p_{2|k}$ such that $p_{1|k} \to_{\mathcal{P}} p_{2|k}$ and $V(p_{2|k}) < V(p_{1|k}) < \infty$. By induction, there exists $\mathbf{p}_k \in \mathbf{P}(p_k, N)$ such that $V(p_{N|k}) < \infty$. ∎

Remark IV.5. *The proposed receding horizon control law is designed using an extension of the terminal constraint*

approach in model predictive control [14] to finite deterministic systems. The particular setting of the Büchi acceptance condition, combined with the energy function V, makes it possible to obtain a non-conservative analogy of the terminal constraint approach, via either a terminal inequality condition (7) *or a terminal equality condition* (8).

B. Control algorithm and its correctness

The overall control strategy for the transition system $\mathcal{T}$ is given in Alg. 1. After the off-line computation of the product automaton and the energy function, the algorithm applies the receding horizon controller $\texttt{RH}^0(S_{\mathcal{P}0})$ at time $k = 0$, or $\texttt{RH}(p_k, \mathbf{p}_{k-1}^\star))$ at time $k > 0$. At each iteration of the algorithm, the receding horizon controller returns the optimal predicted trajectory $\mathbf{p}_k^\star$. The immediate transition $(p_k, p_{1|k}^\star)$ is applied on $\mathcal{P}$ and the corresponding transition $(q_k, \gamma_\mathcal{T}(p_{1|k}^\star))$ is applied on $\mathcal{T}$. This process is then repeated at time $k + 1$.

Algorithm 1 Receding horizon control algorithm for $\mathcal{T} = (Q, q_0, \Delta, \omega, \Pi, h)$, given an LTL formula ϕ over Π

Executed Off-line:
1: Construct a Büchi automaton $\mathcal{B} = (S_\mathcal{B}, S_{\mathcal{B}0}, 2^\Pi, \delta_\mathcal{B}, F_\mathcal{B})$ corresponding to ϕ.
2: Construct the product automaton $\mathcal{P} = \mathcal{T} \times \mathcal{B} = (S_\mathcal{P}, S_{\mathcal{P}0}, \Delta_\mathcal{P}, \omega_\mathcal{P}, F_\mathcal{P})$. Find $V(p)$ for all $p \in S_\mathcal{P}$ [16].

Executed On-line:
1: **if** there exists $p_0 \in S_{\mathcal{P}0}$ such that $V(p_0) \neq \infty$ **then**
2:      Set $k = 0$.
3:      Observe rewards for all $q \in \mathcal{N}(q_0, k)$ and obtain $R_0(q)$.
4:      Obtain $\mathbf{p}_0^\star = \texttt{RH}^0(S_{\mathcal{P}0})$.
5:      Implement transition $(p_0, p_{1|0}^\star)$ on $\mathcal{P}$ and transition $(q_0, \gamma_\mathcal{T}(p_{1|0}^\star))$ on $\mathcal{T}$.
6:      Set $k = 1$
7:      **loop**
8:          Observe rewards for all $q \in \mathcal{N}(q_k, k)$ and obtain $R_k(q)$.
9:          Obtain $\mathbf{p}_k^\star = \texttt{RH}(p_k, \mathbf{p}_{k-1}^\star)$.
10:         Implement transition $(p_k, p_{1|k}^\star)$ on $\mathcal{P}$ and transition $(q_k, \gamma_\mathcal{T}(p_{1|k}^\star))$ on $\mathcal{T}$.
11:         Set $k \leftarrow k + 1$
12:      **end loop**
13: **else**
14:      There is no run originating from q_0 that satisfies ϕ.
15: **end if**

First, we show that the receding horizon controllers used in Alg. 1 are always feasible. We use a recursive argument, which shows that if the problem is feasible for the initial state, or at time $k = 0$, then it remains feasible for all future time-steps $k = 1, 2, \ldots$.

Theorem IV.6 (Recursive Feasiblity)**.** *If there exists $p_0 \in S_{\mathcal{P}0}$ such that $V(p_0) \neq \infty$, then $\texttt{RH}^0(S_{\mathcal{P}0})$ is feasible and $\texttt{RH}(p_k, \mathbf{p}_{k-1}^\star))$ is feasible for all $k = 1, 2, \ldots$.*

Proof. From Lemma IV.1, $\texttt{RH}^0(S_{\mathcal{P}0})$ is feasible. From the definition of $V(p)$, for all $p \in S_\mathcal{P}$, if $p \to_\mathcal{P} p'$, then $V(p') <$

∞ if and only if $V(p) < \infty$. Since $\texttt{RH}^0(S_{\mathcal{P}0})$ is feasible, we have $p_1 = p_{1|0}^\star$ and thus $V(p_1) < \infty$. At each time $k > 0$, if $V(p_k) < \infty$, from Lemmas IV.2, IV.3 and IV.4, we have that controller $\texttt{RH}(p_k, \mathbf{p}_{k-1}^\star)$ is feasible. Since $p_{k+1} = p_{1|k}^\star$, we have $V(p_{k+1}) < \infty$. Using induction we have that $\texttt{RH}(p_k, \mathbf{p}_{k-1}^\star)$ is feasible for all $k = 1, 2, \ldots$. ∎

Finally, we show that Alg. 1 always produces an infinite trajectory satisfying the given LTL formula ϕ, giving a solution to Prob. II.2.

Theorem IV.7 (Correctness of Alg. 1)**.** *Assume that there exists a satisfying run originating from q_0 for a transition system $\mathcal{T}$ and an LTL formula ϕ. Then, Alg. 1 produces an (infinite) trajectory $\mathbf{q} = q_0 q_1 \ldots$ satisfying ϕ.*

Proof. If there exits a satisfying run originating from q_0, then there exists a state $p_0 \in S_{\mathcal{P}0}$ such that $V(p_0) < \infty$. Therefore, from Thm. IV.6, the receding horizon controller is feasible for all $k > 0$, and Alg. 1 will always produce an infinite trajectory $\mathbf{q}$.

At each state p_k at time $k > 0$, if $V(p_k) > 0$, then either Case 1 or Case 2 of the controller $\texttt{RH}(p_k)$ applies. If Case 1 applies, since $V(p_{k|N}^\star) > V(p_{k+1|N}^\star > V(p_{k+2|N}^\star) \ldots$, there exists $j > k$ such that $V(p_{j|N}^\star) = 0$. This is because the state-space $S_\mathcal{P}$ is finite, and therefore, there is only a finite number of possible values for the energy function $V(p)$. At time j, Case 2 of the proposed controller becomes active until time $l = j + i^0(p_j^\star)$, where $V(p_l) = 0$. Therefore, for each time k, if $V(p_k) > 0$, there exists $l > k$ such that $V(p_l) = 0$ by repeatedly applying the receding horizon controller. If $V(p_k) = 0$, then Case 3 of the proposed controller applies, in which case either $V(p_{k+1}) = 0$ or $V(p_{k+1}) > 0$. In either case, using the previous argument, there exists $j > k$ where $V(p_j) = 0$.

Therefore, at any time k, there exists $j > k$ where $V(p_j) = 0$. Furthermore, since j is finite, we can conclude that the number of times where $V(p_k) = 0$ is infinite. By the definition of $V(p), p \in S_\mathcal{P}$, $V_k = 0$ is equivalent to that $p_k \in F_\mathcal{P}^\star \subseteq F_\mathcal{P}$. Therefore, the trajectory $\mathbf{p}$ is accepting. The trajectory produced on $\mathcal{T}$ is exactly the projection $\mathbf{q} = \gamma_\mathcal{T}(\mathbf{p})$, and thus, it can be concluded that $\mathbf{q}$ satisfies ϕ, which completes the proof. ∎

C. Discussions

It is possible to extend the optimization problem of maximizing rewards to other meaningful cost functions. For example, it is possible to assign penalties or costs on states of the system and minimize the accumulated cost of trajectories in the horizon. It is also possible to define costs on state transitions and minimize the control effort (or the combination of this cost function with the one above).

The complexity of the off-line portion of Alg. 1 depends on the size of $\mathcal{P}$. Denoting $|S|$ as the cardinality of a set S, from [18], a Büchi automaton translated from an LTL formula over Π contains at most $|\Pi| \times 2^{|\Pi|}$ states[1].

---

[1]In practice, this upper limit is almost never reached (see [1]).

Therefore, the size of $S_\mathcal{P}$ is bounded by $|Q| \times |\Pi| \times 2^{|\Pi|}$. From [16], the complexity of generating the energy function is $O(|S_\mathcal{P}|^2 + |F_\mathcal{P}|^3)$. The complexity of the on-line portion of Alg. 1 is highly dependent on the horizon N. If the maximal number of transitions at each state of $\mathcal{P}$ is $\Delta_\mathcal{P}^{\max}$, then the complexity at each iteration of the receding horizon controller is bounded by $(\Delta_\mathcal{P}^{\max})^N$, assuming a depth first search algorithm is used to find the optimal trajectory. It may be possible to reduce this complexity from exponential to polynomial if one applies a more efficient graph search algorithm using Dynamic Programming. This will be studied in future research.

V. SOFTWARE IMPLEMENTATION AND CASE STUDY

The control framework presented in this paper was implemented in a user friendly software package, available on `http://hyness.bu.edu/LTL_MPC.html`. To utilize this software, a user needs to input the finite transition system $\mathcal{T}$, an LTL formula ϕ, the horizon N, and a function $\mathcal{R}(q, k)$ that generates the time-varying rewards defined on the states of $\mathcal{T}$. The software executes the control algorithm outlined in Alg. 1, and produces a trajectory in $\mathcal{T}$ that satisfies ϕ and maximizes the rewards collected locally with the proposed receding horizon control laws. This software uses the LTL2BA [18] tool for the translation of an LTL formula to a Büchi automaton.

We now present a case study applying the software package. In this case study, we use the transition system defined as vertices of a rectangular grid as shown in Fig. 1. We consider the following LTL formula, which expresses a robotic surveillance task:

$$
\begin{aligned}
\phi \quad := \quad & \mathsf{G\,F\,base} \\
& \wedge \mathsf{G}\,(\mathsf{base} \longrightarrow \mathsf{X}\,\neg\mathsf{base}\,\mathsf{U}\,\mathsf{survey}) \\
& \wedge \mathsf{G}\,(\mathsf{survey} \longrightarrow \mathsf{X}\,\neg\mathsf{survey}\,\mathsf{U}\,\mathsf{recharge}) \\
& \wedge \mathsf{G}\,\neg\mathsf{unsafe}. \quad\quad (10)
\end{aligned}
$$

The first line of ϕ, $\mathsf{G\,F\,base}$, enforces that the state with observation base is repeatedly visited (possibly for uploading data). The second line ensures that after base is reached, the system is driven to a state with observation survey, before going back to base. Similarly, the third line ensures that after reaching survey, the system is driven to a state with observation $\mathsf{recharge}$, before going back to survey. The last line ensures that, at any time, the states with observation unsafe should be avoided.

We assume that at each state $q \in Q$, the rewards at state q' can be observed if the Euclidean distance between q and q' is less than or equal to 25. In this case study, we define $\mathcal{R}(q, k)$ as follows. At time $k = 0$, the reward value $\mathcal{R}(q, 0)$ at each state q is generated randomly by a uniform sampling in the range of $[10, 25]$. At each subsequent time $k > 0$, if the reward value at a state is positive, then it decays with a specific rate. Otherwise, there is a probability that a reward is assigned to this state with a value chosen by a uniform sampling in the range of $[10, 25]$. In this case study, the states with rewards can be seen as "targets", and the reward values

can be seen as the "amount of interest" associated with each target. The control objective of maximizing the collected rewards can be interpreted as maximizing the information gathered from surveying states with high interest.
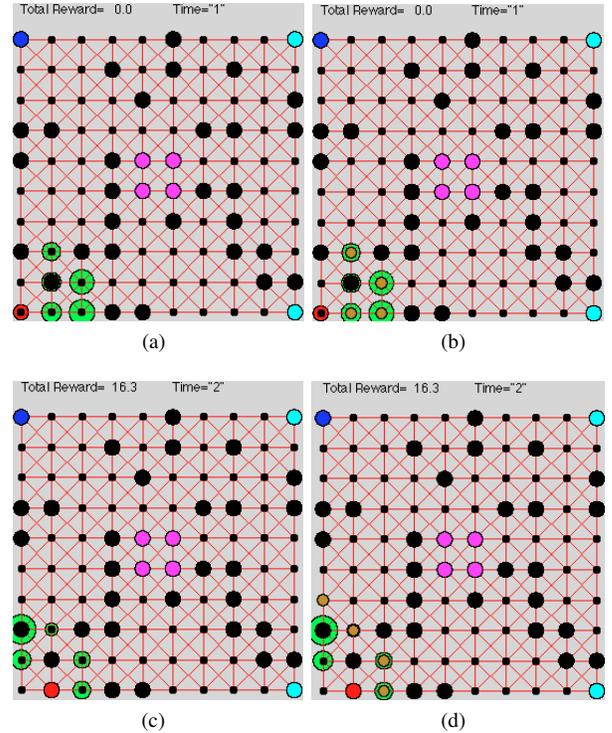

Fig. 4. Snapshots of the system trajectory under the proposed receding horizon control laws. In all snapshots, the state with rewards are marked in green, where the size of the state is proportional with the associated reward. **a)** At time $k = 0$, the initial state of the system is marked in red (in the lower left corner). **b)** The controller $\mathbf{p}_0^\star = \mathrm{RH}^0(S_{\mathcal{P}0})$ is computed at the initial state. The optimal predicted trajectory $\mathbf{p}_0^\star$ is marked by a sequence of states in brown. **c)** The first transition $q_0 \to_\mathcal{T} q_1$ is applied on $\mathcal{T}$ and transition $p_0 \to_\mathcal{P} p_1$ is applied on $\mathcal{P}$. The current state (q_1) of the system is marked in red. **d)** The controller $\mathbf{p}_1^\star = \mathrm{RH}(p_1, \mathbf{p}_0^\star)$ is computed at p_1. The optimal predicted trajectory $\mathbf{p}_1^\star$ is marked by a sequence of states in brown.

By applying the method described in the paper, our software package first translates ϕ to a Büchi automaton $\mathcal{B}$, which has 12 states. This procedure took 0.5 second on a Macbook Pro with a 2.2GHz Quad-core CPU. Since $\mathcal{T}$ contains 100 states, $|S_\mathcal{P}|$ is 1200. The generation of the product automaton $\mathcal{P}$ and the computation of the energy function V took 4 seconds. In this case study, we chose the horizon N to be 4. By applying Alg. 1, some snapshots of the system trajectory are shown in Fig. 4. Each iteration of Alg. 1 took $1 - 3$ seconds (due to different numbers of graph searches needed, the computation time varies for each iteration).

We applies the control algorithm for 100 time-steps. We plotted the results after 100 time-steps in Fig. 5. At the top, we plot the energy $V(p)$ at the each time-step. We see that after 55 time-steps, the energy is 0, meaning that an accepting state is reached. Note that, each time an accepting state is reached, the system visits the base, survey and $\mathsf{recharge}$ states at least once *i.e.,* one cycle of the surveillance mission

task (base − survey − recharge) is completed. We also compare the receding horizon controller with the controller proposed in [16] at the bottom of Fig. 5. We clearly see that the receding horizon controller proposed in this paper performs better in terms of rewards collection, since it reacts much quicker to the time varying rewards. An example video of the evolution of the system trajectory is also available at `http://hyness.bu.edu/LTL_MPC.html`.
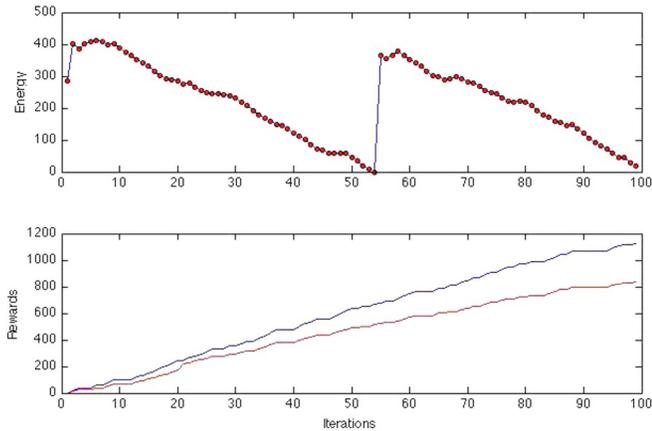

Fig. 5.   Upper figure: plot of energy $V(p)$ at the current state for 100 time-steps. Bottom figure: in blue, plot of the cumulative rewards collected in 100 time-steps by the proposed receding horizon controller; in red, plot of the cumulative rewards collected by the controller in [16] using the same reward function $\mathcal{R}(q, k)$.

VI. CONCLUSION AND FINAL REMARKS

In this paper, a receding horizon control framework that optimizes the trajectory of a finite deterministic system locally, while guaranteeing that the infinite trajectory satisfies a given linear temporal logic formula, was proposed. The optimization criterion was defined as maximization of time-varying rewards associated with the states of the system. A control strategy that makes real-time control decisions in terms of maximizing the reward while ensuring satisfaction of the LTL specification was developed. The proposed framework is a step toward synergy of model predictive control and formal controller synthesis, which is beneficial for both areas.

Future research deals with the extension of the proposed framework to finite probabilistic systems, such as Markov decision processes or partially observed Markov decision processes, where the specifications are given as formulas of probabilistic temporal logic.

REFERENCES

[1] M. Kloetzer and C. Belta, "A fully automated framework for control of linear systems from temporal logic specifications," *IEEE Transactions on Automatic Control*, vol. 53, no. 1, pp. 287–297, 2008.

[2] L. Habets, P. Collins, and J. van Schuppen, "Reachability and control synthesis for piecewise-affine hybrid systems on simplices," *IEEE Transactions on Automatic Control*, vol. 51, pp. 938–948, 2006.

[3] J. Desai, J. Ostrowski, and V. Kumar, "Controlling formations of multiple mobile robots," in *IEEE International Conference on Robotics and Automation*, 1998, pp. 2864–2869.

[4] L. Habets, M. Kloetzer, and C. Belta, "Control of rectangular multi-affine hybrid systems," in *IEEE Conference on Decision and Control*, 2006, pp. 2619–2624.

[5] T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Receding horizon temporal logic planning for dynamical systems," in *IEEE Conference on Decision and Control and Chinese Control Conference*, Shanghai, China, December 2009, pp. 5997–6004.

[6] C. Belta, V. Isler, and G. Pappas, "Discrete abstractions for robot motion planning and control in polygonal environments," *Robotics, IEEE Transactions on*, vol. 21, no. 5, pp. 864–874, 2005.

[7] S. Lindemann, I. Hussein, and S. LaValle, "Real time feedback control for nonholonomic mobile robots with obstacles," in *IEEE Conference on Decision and Control*, New Orleans, LA, December 2007, pp. 2406–2411.

[8] R. Milner, *Communication and Concurrency*.   Prentice-Hall, 1989.

[9] H. Kress-Gazit, G. Fainekos, and G. J. Pappas, "Where's Waldo? Sensor-based temporal logic motion planning," in *IEEE International Conference on Robotics and Automation*, 2007, pp. 3116–3121.

[10] S. Karaman and E. Frazzoli, "Sampling-based motion planning with deterministic μ-calculus specifications," in *IEEE Conference on Decision and Control*, Shanghai, China, 2009, pp. 2222 – 2229.

[11] S. G. Loizou and K. J. Kyriakopoulos, "Automatic synthesis of multi-agent motion tasks based on LTL specifications," in *IEEE Conference on Decision and Control*, Paradise Islands, The Bahamas, December 2004, pp. 153–158.

[12] E. M. Clarke, D. Peled, and O. Grumberg, *Model checking*.   MIT Press, 1999.

[13] N. Piterman, A. Pnueli, and Y. Saar, "Synthesis of reactive(1) designs," in *International Conference on Verification, Model Checking, and Abstract Interpretation*, Charleston, SC, January 2006, pp. 364–380.

[14] J. B. Rawlings and D. Q. Mayne, *Model Predictive Control: Theory and Design*.   Nob Hill Publishing, 2009.

[15] W. Li and C. Cassandras, "A cooperative receding horizon controller for multivehicle uncertain environments," *Automatic Control, IEEE Transactions on*, vol. 51, no. 2, pp. 242–257, 2006.

[16] X. C. Ding, C. Belta, and C. G. Cassandras, "Receding horizon surveillance with temporal logic specifications," in *IEEE Conference on Decision and Control*, December 2010, pp. 256–261.

[17] M. Browne, E. Clarke, and O. Grumberg, "Characterizing finite kripke structures in propositional temporal logic," *Theoretical Computer Science*, vol. 59, no. 1-2, pp. 115–131, 1988.

[18] P. Gastin and D. Oddoux, "Fast LTL to Buchi automata translation," *Lecture Notes in Computer Science*, pp. 53–65, 2001.