Information-guided persistent monitoring under temporal logic constraints

Austin Jones, Mac Schwager, and Calin Belta

Abstract—We study the problem of planning the motion of an agent such that it maintains indefinitely a high-quality estimate of some *a priori* unknown feature, such as traffic levels in an urban environment. Persistent operation requires that the agent satisfy motion constraints, such as visiting charging stations infinitely often, which are readily described by rich linear temporal logic (LTL) specifications. We propose and evaluate via simulation a two-level dynamic programming algorithm that is guaranteed to satisfy given LTL constraints. The low-level path planner implements a receding horizon algorithm that maximizes the local information gathering rate. The high-level planner selects inputs to the low-level planner based on global performance considerations.

I. Motivation

In this paper, we address the problem of planning the path of a mobile robot such that it persistently maintains a high-quality estimate of some unknown feature, i.e. persistent monitoring. This describes many real-world applications in which human decision-makers require real-time information about large environments, such as forest rangers monitoring wildfires. In order to provide up-to-date information reliably, the robot has to plan its motion such that relevant information is gained, e.g. the agent visits locations that have not recently been visited. This can be formalized as maximizing the expected mutual information rate (MIR) between the environment and the agent's sensor measurements. In addition, the agent must continue to function indefinitely, i.e. it must satisfy constraints such as "Regularly visit a recharging station and always avoid obstacles". In this work, we use linear temporal logic (LTL), an extension of Boolean logic that is capable of describing how a system may change over time, to model such constraints. We present a hierarchal stochastic optimal control algorithm that is guaranteed to satisfy the given LTL constraints while attempting to maximize the MIR.

Austin Jones is with the Division of Systems Engineering, Mac Schwager and Calin Belta are with the Division of Systems Engineering and the Department of Mechanical Engineering at Boston University, Boston, MA 02115. Email: {austinmj,schwager,cbelta}@bu.edu

Persistent monitoring strategies have recently been considered in the literature. In [16], the authors demonstrate how to regulate the speed of an agent moving along a fixed path such that the uncertainty about the state of the environment remains below a certain threshold indefinitely. This work is extended to planning trajectories in [14]. The problem of multi-agent persistent monitoring is investigated in [6]. The above works consider a measurement model that is independent of the agent's location in the environment and either completely unconstrained motion or motion along predefined paths. In contrast, we assume that the agent's sensing capability depends on its position in the environment and we incorporate linear temporal logic (LTL) motion constraints.

Linear temporal logic [2], [17] can be used to formulate liveness ("Visit a charging station infinitely often"), fairness ("Visit a data upload station before visiting a charging station"), and safety ("Always avoid obstacles") properties. Off-the-shelf formal synthesis tools exist that guarantee a robot's trajectory will satisfy an LTL formula [1], [5], [18]. These tools were applied to the problem of persistent monitoring in [8] in which the goal was to minimize the amount of time between visiting pre-specified surveillance regions. This work did not explicitly represent the agent's sensing model nor the environment model.

In [9], [10], we considered planning under syntactically co-safe linear temporal logic (scLTL) constraints on the motion of the robot, a finite-time subset of LTL. scLTL can be used to describe reachability and finite-time safety properties, among others. By considering full LTL formulae, the approach in this paper allows us to enforce infinite-horizon properties such as visiting different sequences of regions in the environment infinitely often. The algorithm in this paper extends the receding-horizon planner developed in [10] to guarantee satisfaction of LTL constraints. In addition, to avoid myopia inherent in receding horizon implementations, we develop a high-level sample-based planner that selects inputs to the receding horizon algorithm that attempt to maximize the MIR over the infinite horizon. An implementation of our procedure is applied to a simulation of a target tracking case study.

II. MATHEMATICAL PRELIMINARIES

We assume the reader is familiar with standard set theoretic notation and information-theoretic concepts such as mutual information and entropy [15]. We use the shorthand notation $x^{1:t}$ for a time-indexed sequence $x^1 \ldots x^t$. The set of all finite and set of all infinite words over alphabet Σ are denoted by Σ^* and Σ^∞, respectively.

A *deterministic transition system* [2] is a tuple $TS = (Q, q_0, Act, Trans, AP, L)$, where Q is a set of states, $q_0 \in Q$ is the initial state, Act is a set of actions, $Trans \subseteq Q \times Act \times Q$ is a transition relation, AP is a set of atomic propositions, and $L : Q \to 2^{AP}$ is a labeling function of states to atomic propositions.

A *discrete time Markov Chain* (MC) is a tuple $MC = (S, s^0, P)$ where S is a discrete set of states, $P : S \times S \to [0, 1]$ is a probabilistic transition relation such that the chain moves from state s to state s' with probability $P(s, s')$, and s^0 is an initial state of the system. A *discrete time Markov decision process* (MDP) is a tuple $MDP = (S, s^0, P, Act)$, where S, s^0 are as defined for a MC, Act is a set of actions, and $P : S \times Act \times S \to [0, 1]$ is a probabilistic transition relation such that taking action a drives MDP from state s to state s' with probability $P(s, a, s')$. In this paper, we assume that the reader is familiar with stochastic optimal control over MDPs [4].

A *linear temporal logic* (LTL) formula is inductively defined as follows [12]:

$$\phi := p | \neg \phi | \phi \vee \phi | \phi \wedge \phi | \phi \; \mathcal{U} \; \phi | \bigcirc \phi | \Diamond \phi | \square \phi, \quad (1)$$

where p is an atomic proposition, $\neg$ (negation), $\vee$ (disjunction), and $\wedge$ (conjunction) are Boolean operators, and $\bigcirc$ ("next"), $\mathcal{U}$ ("until"), $\Diamond$ ("eventually"), and $\square$ ("always") are temporal operators. LTL allows for the specification of persistent performance requirements, such as "Data is uploaded infinitely often".

A *deterministic Buchi automaton* (DBA) [2] is a tuple $\mathcal{A} = (\Sigma, \Pi, \delta, \sigma_0, F)$ where Σ is a finite set of states, Π is a finite alphabet, $\delta \subseteq \Sigma \times \Pi \times \Sigma$ is a deterministic transition relation, $\sigma_0 \in \Sigma$ is the initial state of the automaton, and $F \subseteq \Sigma$ is a set of final (accepting) states. An *accepting run* on a DBA is a sequence of states in Σ^∞ that intersects with F infinitely often. The *language* of a DBA (written $\mathcal{L}(\mathcal{A})$) is the set of all accepting words in Σ^∞. For a large subset of LTL formulae ϕ, there exist algorithmic procedures to construct a DBA $\mathcal{A}_\phi$ with alphabet 2^{AP} such that the language of all words satisfying ϕ, $\mathcal{L}(\phi)$, is equal to $\mathcal{L}(\mathcal{A}_\phi)$. [17].

The *product automaton* of a transition system $TS = (Q, q_0, Act, Trans, AP, L)$ and a DBA $\mathcal{A}_\phi = (\Sigma, 2^{AP}, \delta, \sigma_0, F)$ is the Buchi automaton $\mathcal{P} = TS \times$ $\mathcal{A}_\phi = (\Sigma_{\mathcal{P}}, \chi^0, Act, F_{\mathcal{P}}, \Delta_{\mathcal{P}})$ [2]. $\Sigma_{\mathcal{P}} \subseteq Q \times \Sigma$ is the state space of the automaton, $\chi^0 = (q_0, \sigma_0)$ is the initial state, and $F_{\mathcal{P}} \subseteq Q \times F$ is the set of accepting states. The transition relation is defined as $\Delta_{\mathcal{P}} = \{(q, \sigma), p, (q', \sigma') | (q, p, q') \in Trans, (\sigma, L(q), \sigma') \in \delta\}$. The state of the automaton at time k, (q^k, σ^k) is denoted as χ^k for short. If $\chi^{0:k}$ satisfies the acceptance condition on $\mathcal{P}$, then the trajectory $q^{0:k}$ satisfies ϕ.

We define the *distance to acceptance* [1] as a function $W : \Sigma_{\mathcal{P}} \to \mathbb{Z}^+$ such that $W(\chi)$ is the minimal number of actions that can be taken to drive $\mathcal{P}$ from χ to an accepting state in $F_{\mathcal{P}}$. The *k-step boundary* about a state χ, denoted $\partial N(\chi, k)$ is the sets of states that can be reached by applying exactly k inputs to the product automaton.

III. MODELS

A. Robot motion model

We consider a single robot with perfect localization and known dynamics operating in a continuous, partitioned environment. We abstract the motion of the robot to a transition system $Robot = (Q, q_0, Act, Trans, AP, L)$ where Q is a partition of the environment, q_0 is the region in which the agent is initially located, Act is a set of control policies, AP is a set of known regional properties of the environment, and $(q, a, q') \in Trans$ if a can drive the robot from q to q'. Tools have been developed for low-level control synthesis in partitioned systems with linear [11] and piecewise affine [18] dynamics. With some conservativism, these can be extended to more realistic dynamics such as the ones modeling unicycles and car-like vehicles [3].

Example 1. Consider a robot R_p operating in a grid environment as shown in Figure 1(a) to track a target robot R_t. In this case, $Q = \{1, \ldots 4\}^2 \times \{N, S, E, W\}$ (north, south, east, west) where a state $q = (i, j, dir)$ means that R_p is in grid cell (i, j) and facing direction dir. The set of actions is $Act = \{straight, CW, CCW\}$ which mean go straight, rotate 90° clockwise, and rotate 90° counterclockwise, respectively. $AP = \{\pi_{recharge}(\text{green}), \pi_{data}(\text{magenta}), \pi_{alarm}(\text{cyan}), \pi_{reset}(\text{blue}), \pi_{obs}(\text{red})\}$. A subset of the transition system is shown in Figure 1(b). $\square$

B. Robot sensing model

We associate with the environment a feature that evolves in time synchronously with the robot according to the Markov chain $Env = (S, s^0, P)$. We denote the state of Env at time k as s^k. The initial state s^0 is *a priori* unknown. At each time k, when the robot moves to state q^k, it measures s^k according
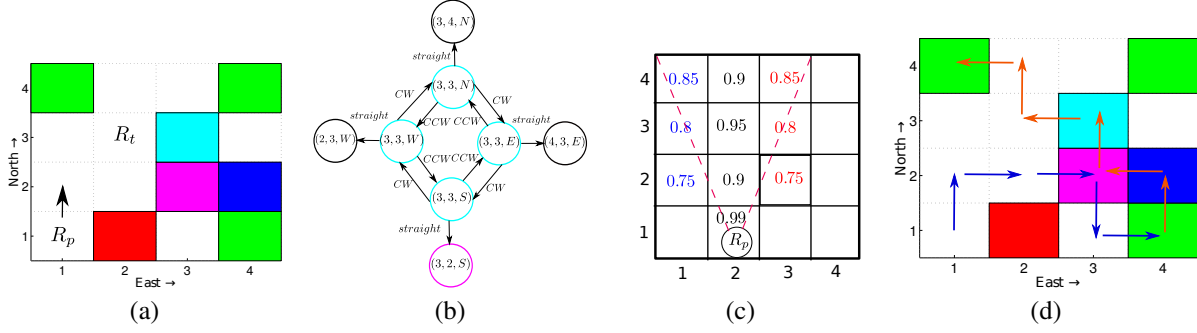
Fig. 1. (a) Environment used in Example 1. (b) Part of transition system corresponding to (a). (c) Part of the measurement likelihood function. Numbers in blue, black, and red text indicate the probability of $l, c,$ or r being measured by R_p if R_t is in the indicated cell. (d) Example of a prefix (blue) and a suffix cycle (orange) from formula 5.

to a noisy measurement $y^k \in R_Y$. y^k is a realization of a discrete random variable Y^k. We denote the conditional measurement distribution as $h(y, s, q) = Pr[$ measurement is $y|Env$ in $s, Robot$ in $q]$.

Example 1 (continued)**.** In the above scenario, $S = \{1, \ldots, 4\}^2$ is the set of possible locations of R_T. s^0 is chosen randomly. Env can transition to an adjacent state with probability p_{move}. R_p has a noisy camera that can produce measurements in $R_Y = \{0, l, c, r\}$ which correspond to R_t not being observed and R_t being observed in the left, center, or right portion of R_p's field of view. Part of the measurement likelihood function is summarized in Figure 1(c). □

The robot's estimate of s^k is given via the estimate pmf b^k, called the *belief* , where $b^k(s) = Pr[s^k = s|y^{1:k}, q^{0:k}]$. The initial belief b^0 reflects prior knowledge about s^0. b^k is maintained via the recursive Bayes filter

$$b^k(s) = \frac{h(y^k, s, q^k) \sum_{s' \in S} P(s', s) b^{k-1}(s')}{\sum_{\sigma \in S} h(y^k, \sigma, q^k) \sum_{s' \in S} P(s', \sigma) b^{k-1}(s')}. \tag{2}$$

The belief b^k evolves over time according to the MDP $Est = (B, b^0, P_{est}, Q)$. Q and b^0 are as defined previously. P_{est} is the probabilistic transition relation such that if b' is the result of applying (2) with measurement y collected in state q, then $P_{est}(b, q, b')$ is the probability of observing y. B is the countably infinite set of all possible beliefs that can be outputs of computing the Bayes filter with initial belief b^0 with an infinite run of $Robot$ along with the resulting measurements. The MDP Est is referred to as the *belief tree* in the partially observable MDP (POMDP) literature [13].

IV. PROBLEM STATEMENT

Our goal is to plan an infinite horizon trajectory $q^{0:\infty}$ for $Robot$ that maximizes the quality of the estimate pmf

over time while satisfying a linear temporal logic specification. In [9], [10], we used the expected conditional entropy $E_{Y^{1:t}}[H(b^t)]$ to quantify the expected quality of the estimate that would result from the robot traveling along the finite path $q^{0:t}$. Here, we use a related quantity called the mutual information rate [7], [15]

$$MIR = \lim_{t \to \infty} \frac{I(s^{0:t}; Y^{0:t})}{t}. \tag{3}$$

MIR is the average rate of information gain about the state of Env when a new measurement is taken. We wish to maximize this quantity, as we want to increase the rate at which measurements help to identify the state of Env. Maximizing (3) over the set of actions $Robot$ can take is equivalent to minimizing the *entropy rate*

$$ER = \lim_{t \to \infty} \frac{H(b^t)}{t}. \tag{4}$$

Any trajectory that satisfies an LTL formula can be segmented into a finite length prefix path and an infinite sequence of finite length suffix cycles.

Definition 1 (Prefix Path and Suffix Cycle)**.** Let $q^{0:\infty}$ be an infinite horizon trajectory over a transition system TS and ϕ be an LTL formula over the properties AP in TS. Let $\mathcal{P} = TS \times \mathcal{A}_\phi$ and let $\chi^{0:\infty}$ be the infinite run over $\mathcal{P}$ induced by $q^{0:\infty}$. Let l_n be the nth time at which $\chi^{0:\infty}$ intersects F_P. The *prefix path* is the finite sequence $q^{0:l_1}$. A *suffix cycle* is a finite sequence of states $q^{l_n+1:l_{n+1}}$.

Example 1 (continued)**.** The constraints on R_p can be given as the LTL formula

$$\phi = \quad \square \ \lozenge \ \pi_{recharge} \wedge \ \square \ \lozenge \ \pi_{data} \wedge \ \square \ (\neg \pi_{obs})$$
$$\wedge \ \square \ (\pi_{alarm} \Rightarrow (\neg \pi_{recharge} \ \mathcal{U} \ \pi_{reset})), \tag{5}$$

which in plain English is "Visit recharging and data upload stations infinitely often while avoiding obstacles.

3

If an alarm is triggered, visit the shutdown switch before visiting the recharging station." Figure 1(d) shows a prefix cycle and a suffix cycle for ϕ. □

We wish to minimize the average entropy rate per cycle, given as

$$AERPC = \lim_{n \to \infty} \frac{\sum_{i=0}^{n} \frac{H(b^{l_n}) - H(b^{l_{n-1}})}{l_n - l_{n-1}}}{n} \qquad (6)$$

This formulation is similar to the average cost-per-stage problem [8], [4]. With this new objective, we define the constrained persistent monitoring problem as

Problem 1. *Consider an agent that is estimating the state of the environment. Solve*

$$\min_{a^0 \ldots \in Act^\infty} E_{Y^{0:l_\infty}}[AERPC]$$
$$s.t. \qquad \phi, \qquad l_{n+1} - l_n \le \ell_{max} \forall n, \qquad (7)$$

where ϕ is an LTL formula and ℓ_{max} is a per-cycle budget.

The budget ℓ_{max} describes energy constraints on the agent's motion, e.g. the maximum amount of time that the agent can be active between recharging.

V. SOLUTION

In [10], we defined a Markov decision process that encapsulated the effect of the robot's action on it's location in the environment, progress towards satisfying the given specification, and it's expected gain in information.

Definition 2. Full Model MDP The MDP $FullModel = (\Sigma_\mathcal{P} \times B, P_{tot}, Act, (\chi_0, b^0))$, where $\mathcal{P} = Robot \times \mathcal{A}_\phi$ and B, b^0 are derived from Est, describes the synchronous evolution of the robot's position and belief state. The probabilistic transition relationship P_{tot} is defined as

$$P_{tot}((\chi, b), a, (\chi', b')) =$$
$$P_{est}(b, q', b') I((\chi, a, \chi') \in \Delta_\mathcal{P}) I(W(\chi') < \infty) \qquad (8)$$

where I is the indicator function ($I(x \in X)$ is 1 if $x \in X$ and 0 if $x \notin X$) and $\chi' = (q', \sigma')$.

We map Problem 1 to the constrained stochastic optimal control problem

$$\min_{a^0 \ldots \in Act^\infty} E_{Y^{0:l_\infty}}[AERPC]$$
$$\text{subject to} \qquad (9)$$
$$l_{n+1} - l_n \le \ell_{max} \forall n \in \mathbb{N},$$
$$(\chi^{i+1}, b^{i+1}) \sim P_{tot}(\cdot, a^i, (\chi^i, b^i)) \ \forall i \in \mathbb{N}.$$

We propose a hybrid dynamic programming-based method to circumvent these difficulties. A local receding-horizon algorithm (Section V-A) chooses actions that

locally improve the entropy rate and drive the agent to an accepting state in $\mathcal{P}$. The budget of actions given to the receding horizon algorithm is selected according to a policy that is calculated off-line via value iteration [4] to optimize the infinite-horizon AERPC (Section V-B).

A. Optimization of a single cycle

In [10], we presented a receding horizon algorithm that locally minimized entropy and was guaranteed to satisfy the given scLTL constraint within the budget. Algorithm 1 extends that procedure to handle LTL constraints.

Algorithm 1 Receding horizon planner.

---

  function RHP($(\chi, b), \ell, m, n$)
  $k = 0$; $\chi[k] = \chi$;
  while $\chi[k] \notin F_\mathcal{P}$ **do**
    $\Sigma_p :=$ PossibleStates($\chi[k]$,m,$\ell - k$)
    $(X_p, P_{tot}, Act_p) :=$ ConstructMDP($\chi[k]$,b,Σ_p)
    $\mu_l :=$ BellmanIteration(X_p, P_{tot}, Act_p)
    if $k \ge \ell - m$ **then**
      $n := \ell - k$
    for $i := 1$ to n **do**
      $(\chi[k+1], b) :=$ result from applying $\mu(i, (\chi, b))$
      $k++$
  return $(\chi[k], b, k)$

At the kth time-step after the invocation of Algorithm 1, the procedure PossibleStates constructs m sets $\Sigma_p[i] \subseteq \Sigma_\mathcal{P}$ where that contains all states reachable from χ^k in i actions and from which an accepting state can be reached under the remaining budget. From the sets Σ_p, the procedure ConstructMDP constructs a subset of the full-model MDP such that at the ith level, only pairs (χ, b) such that $\chi \in \Sigma_p[i]$ appear. Next, the algorithm BellmanIteration performs standard Bellman iteration over the constructed MDP where the terminal cost is given as $\frac{H(b[m]) - H(b[0])}{m}$, where $b[m]$ is the terminal state and $b[0]$ was the initial belief state input to Algorithm 1. The algorithm is executed in a receding horizon fashion until an accepting state is reached. Applying Algorithm 1 infinitely often provably satisfies the given specification and a single application of Algorithm 1 is tractable.

Theorem 1. *Let $\ell' = \max_{\chi \in F_\mathcal{P}} \min_{\chi' \in \partial N(\chi, 1)} W(\chi') + 1$. If $\ell' \le \ell_{max}$, then sequentially applying Algorithm 1 infinitely often with budget at least ℓ' is guaranteed to drive the robot to satisfy the given LTL specification.*

Proof. Applying Algorithm 1 one time is guaranteed to drive the system to an accepting state [10]. Applying Algorithm 1 infinitely often guarantees that the system will

4

Fig. 2. A block diagram illustrating the hierarchal algorithm.

be in an accepting state infinitely often, thus satisfying the Buchi acceptance condition and therefore satisfying the given LTL specification. □

Proposition 1. *The time complexity of Algorithm 1 is* $O(|Act|^m |R_Y|^m |S| \lceil \frac{\ell}{n} \rceil)$.

B. Choosing cycle budgets

Although applying Algorithm 1 infinitely often with a fixed budget ℓ is guaranteed to satisfy the LTL specification, there is no guarantee that the local information gathering performs well over an infinite time horizon. We propose to pair the local information gathering with long-term planning by finding a policy $\mu_g : F_{\mathcal{P}} \cup \{\chi_0\} \times B \to \mathbb{N}$ that maps an initial or accepting state in the automaton and a belief state, e.g. the configuration of the robot at the end of a cycle, to the budget ℓ that should be given to the receding horizon planner in the next cycle. The hierarchal structure of the algorithm is illustrated in Figure 2. The policy construction is performed off-line before the agent is deployed.

Algorithm 2 details how μ_g is calculated. We use simulation to characterize the per-cycle performance of the receding horizon algorithm. We consider a finite set of states $T \subset F_{\mathcal{P}} \times B$ in our optimization. Algorithm 2 begins by performing j_{max} simulations of Algorithm 1 from the initial state (χ^0, b^0) for each value of ℓ from $\min_{\chi'' \in \partial N(\chi^0, 1)} W(\chi'') + 1$ to ℓ_{max}. The result of each simulation is a state (χ', b') that is reachable from the initial state. In our algorithm, if two states are ϵ close in the 1-norm, we consider them identical. The cost of going from (χ, b) to (χ', b') under budget ℓ, denoted $g((\chi, b), \ell, (\chi', b')) = \frac{H(b') - H(b)}{\ell}$, is calculated and recorded. Similarly, the frequency of transitioning from state (χ, b) to (χ', b'), under budget ℓ, denoted $P((\chi, b), \ell, (\chi', b'))$, is calculated from the set of simulations and recorded. The reachable state (χ', b') is then added to the set the set of states still to be checked C.

It may seem that the longer the budget ℓ handed to the receding horizon planner, the better we expect the performance to be. As ℓ increases, on average more states are included in the MDPs constructed by Algorithm 1. However, the agent can take an action that locally performs better but which may cause the long-term performance to deviate far from optimal behavior.

After this set of simulations is completed, the states of C are iterated through and more simulations occur.

T, P, and g are populated until $|T| \geq N$ and all of the accepting states in $F_{\mathcal{P}}$ have been visited or there are no states left to be checked We invoke a procedure called MakeRecurrent to ensure every state in T is connected via P to another state in T. Finally, Algorithm 2 uses value iteration to find the optimal policy

$$\mu_g(\chi, b) = \arg\min_\ell \sum_{(\chi', b') \in T} P((\chi, b), \ell(\chi', b'))$$
$$[g((\chi, b), \ell(\chi', b') + J^\infty((\chi', b'))], \quad (10)$$

for all $(\chi, b) \in T$ where $J^\infty : T \to \mathbb{R}$ is the infinite-horizon cost-to-go function. [4] This optimization minimizes the sum of the expected entropy rates per cycle rather than directly minimize the AERPC. However, since we are minimizing the entropy rate on a cycle-by-cycle basis, the average rate will also be minimized. Algorithm 3 summarizes our approach.

Algorithm 2 Constructs optimal budget policy.

1: **function** OptimalBudget($\chi_0, b^0, \ell_{max}, m, n$)
2: $T = \{(\chi_0, b^0)\}$; $C = \{(\chi_0, b_0)\}$; $Acc = F_{\mathcal{P}}$
3: **while** $|T| \leq N$ and $C \neq \emptyset$ **do**
4: $\quad (\chi, b) = C.\text{pop}()$
5: $\quad$ **for** $\ell = \min_{\chi'' \in \partial Nin} W(\chi'') + 1$ to ℓ_{max} **do**
6: $\quad\quad$ **for** $j = 1$ to j_{max} **do**
7: $\quad\quad\quad (\chi', b', k) = \text{RHP}((\chi, b), \ell, m, n)$
8: $\quad\quad\quad$ **if** $\exists (\chi', b'') \in T$ s.t. $||b'' - b'||_1 < \epsilon$ **then**
9: $\quad\quad\quad\quad b' = b''$
10: $\quad\quad\quad$ **else**
11: $\quad\quad\quad\quad T.\text{add}((\chi', b'))$; $C.\text{add}((\chi', b'))$
12: $\quad\quad\quad g((\chi, b), \ell, (\chi', b')) = \frac{H(b') - H(b)}{k}$
13: $\quad\quad\quad P((\chi, b), \ell, (\chi', b')) \mathrel{+}= \frac{1}{j_{max}}$
14: $\quad\quad\quad Acc = Acc \setminus \{\chi'\}$
15: $\quad\quad\quad$ **if** $|T| = N$ and $Acc \neq \emptyset$ **then**
16: $\quad\quad\quad\quad N = N + 1$
17: $(P, g, T) = \text{MakeRecurrent}(P, g, T)$
18: $\mu_g = \text{ValueIteration}(P, g, T)$
19: **return** μ_g, T

Algorithm 3 Constrained Persistent Monitoring.

$\mu_g, T = \text{OptimalBudget}(\chi_0, b^0, \ell_{max}, m, n)$;
$\chi = \chi_0$; $b = b^0$;
while TRUE **do**
$\quad (\chi, b, k) = \text{RHP}((\chi, b), \mu_g((\chi, b)), m, n)$

VI. CASE STUDY

We implemented Algorithm 3 in software and applied it to a simulation of the scenario described in the running

5

Fig. 3. (a) random walk policy (b) Algorithm 3.

example Example 1. The probability that the tracked target R_t moves to an adjacent cell is $p_{move} = 0.15$. The optimal policy μ_g was computed with parameters $\epsilon = 0.01$, $N = 1500$, $j_{max} = 100$, $\ell_{max} = 15$, $m = 2$, and $n = 1$. The computation required approximately 6 hrs. of processor time. We performed 500 Monte Carlo trials of the system in which the system was simulated until 5 suffix cycles were completed and compared the results to 500 "random walk" simulations which were guaranteed to satisfy the constraints. The results from the simulations are summarized in the histograms in Figure 3. The random walk policy resulted in an average entropy rate of 0.0420 bits/action, while the average entropy rate when using Algorithm 3 was -0.0090 bits/action. A two-sample t-test confirmed that this difference in means is statistically significant with p-value less than 10^{-95}. Further, 3.5 % of the random walk trials resulted in negative entropy rates while 64.8% of the trials with Algorithm 3 had negative rates, i.e. on average gain information about the location of R_t with each action taken, i.e. Algorithm 3 has better performance more often.

VII. Conclusions

We have extended our receding horizon planner for constrained informative planning to handle infinite-horizon constraints modeled as LTL formulae. In order to mitigate myopia, we developed and implemented an off-line algorithm to select the number of actions that the receding horizon planner can take between subsequent cycles of the robot's trajectory. Initial results indicate the receding horizon implementation significantly outperforms a random walk simulation. We will investigate more sophisticated reinforcement learning algorithms and extend these results to multi-agent systems.

REFERENCES

[1] Ebru Aydin Gol, Mircea Lazar, and Calin Belta. Language-guided controller synthesis for discrete-time linear systems. In *Proceedings of the 15th ACM international conference on Hybrid Systems: Computation and Control*, pages 95–104, New York, NY, USA, 2012.

[2] Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking (Representation and Mind Series)*. The MIT Press, 2008.

[3] C. Belta, V. Isler, and G. J. Pappas. Discrete abstractions for robot planning and control in polygonal environments. *IEEE Trans. on Robotics*, 21(5):864–874, 2005.

[4] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2nd edition, 2000.

[5] A. Bhatia, L.E. Kavraki, and M.Y. Vardi. Motion planning with hybrid dynamics and temporal goals. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 1108 –1115, Dec. 2010.

[6] Christos G. Cassandras and Xuchao Lin. Optimal control of multi-agent persistent monitoring systems with performance constraints. In Danielle C. Tarraf, editor, *Control of Cyber-Physical Systems*, volume 449 of *Lecture Notes in Control and Information Sciences*, pages 281–299. Springer International Publishing, 2013.

[7] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, 2nd edition, 2006.

[8] Xu Chu Ding, C. Belta, and C.G. Cassandras. Receding horizon surveillance with temporal logic specifications. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 256 –261, Dec. 2010.

[9] Austin Jones, Mac Schwager, and Calin Belta. A receding horizon algorithm for informative path planning with temporal logic constraints. In *International Conference on Robotics and Automation (ICRA)*, 2013.

[10] Austin Jones, Mac Schwager, and Calin Belta. Formal synthesis of optimal information-gathering policies. *IEEE Trans. on Robotics*, Submitted.

[11] M. Kloetzer and C. Belta. A fully automated framework for control of linear systems from temporal logic specifications. *Automatic Control, IEEE Transactions on*, 53(1):287 –297, feb. 2008.

[12] Orna Kupferman and Moshe Y. Vardi. Model checking of safety properties. *Formal Methods in System Design*, 2001. volume 19, pages 291-314.

[13] Hanna Kurniawati, Yanzhu Du, David Hsu, and Wee Sun Lee. Motion planning under uncertainty for robotic tasks with long time horizons. *The International Journal of Robotics Research*, 2010.

[14] Xiaodong Lan and M. Schwager. Planning periodic persistent monitoring trajectories for sensing robots in gaussian random fields. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 2415–2420, May 2013.

[15] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423,623–656, 1948.

[16] S. L. Smith, M. Schwager, and D. Rus. Persistent robotic tasks: Monitoring and sweeping in changing environments. *IEEE Transactions on Robotics*, 28(2):410–426, April 2012.

[17] Moshe Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Logics for Concurrency: Structure versus Automata, volume 1043 of Lecture Notes in Computer Science*, pages 238–266. Springer-Verlag, 1996.

[18] Boyan Yordanov, Jana Tumova, Ivana Cerna, Jiri Barnat, and Calin Belta. Temporal logic control of discrete-time piece-wise affine systems. *IEEE Transactions on Automatic Control*, 57:1491–1504, 2012.

6