

# Entropy Maximization for Markov Decision Processes Under Temporal Logic Constraints

Yagiz Savas, Melkior Ornik, Murat Cubuktepe, Mustafa O. Karabag, Ufuk Topcu

**Abstract**—We study the problem of synthesizing a policy that maximizes the entropy of a Markov decision process (MDP) subject to a temporal logic constraint. Such a policy minimizes the predictability of the paths it generates, or dually, maximizes the exploration of different paths in an MDP while ensuring the satisfaction of a temporal logic specification. We first show that the maximum entropy of an MDP can be finite, infinite or unbounded. We provide necessary and sufficient conditions under which the maximum entropy of an MDP is finite, infinite or unbounded. We then present an algorithm which is based on a convex optimization problem to synthesize a policy that maximizes the entropy of an MDP. We also show that maximizing the entropy of an MDP is equivalent to maximizing the entropy of the paths that reach a certain set of states in the MDP. Finally, we extend the algorithm to an MDP subject to a temporal logic specification. In numerical examples, we demonstrate the proposed method on different motion planning scenarios and illustrate the relation between the restrictions imposed on the paths by a specification, the maximum entropy, and the predictability of paths.

## I. INTRODUCTION

Markov decision processes (MDPs) model sequential decision-making in stochastic systems with nondeterministic choices. A policy, i.e., a decision strategy, resolves the nondeterminism in an MDP and induces a stochastic process. In this regard, an MDP represents a (infinite) family of stochastic processes. In this paper, for a given MDP, we aim to synthesize a policy that induces a process with maximum entropy among the ones whose paths satisfy a temporal logic specification.

Entropy, as an information-theoretic quantity, measures the unpredictability of outcomes in a random variable [1]. Considering a stochastic process as an infinite sequence of (dependent) random variables, we define the entropy of a stochastic process as the joint entropy of these random variables by following [2], [3]. Therefore, intuitively, our objective is to obtain a process whose paths satisfy a temporal logic specification in the most unpredictable way to an observer.

Typically, in an MDP, a decision-maker is interested in satisfying certain properties [4] or accomplishing a task [5]. Linear temporal logic (LTL) is a formal specification language [6] that has been widely used to check the reliability of

software [7], describe tasks for autonomous robots [8], [9] and verify the correctness of communication protocols [10]. For example, in a robot navigation scenario, it allows to specify tasks such as safety (never visit the region A), liveness (eventually visit the region A) and priority (first visit the region A, then B).

The entropy of paths of a (Markovian) stochastic process is introduced in [11] and quantifies the randomness of realizations with fixed initial and final states. We first extend the definition for the entropy of paths to realizations that reach a certain set of states, rather than a fixed final state. Then, we show that the entropy of a stochastic process is equal to the entropy of paths of the process, if the process has a finite entropy. The established relation provides a mathematical basis to the intuitive idea that maximizing the entropy of an MDP minimizes the predictability of paths.

We observe that the maximum entropy of an MDP under stationary policies may not exist, i.e., for any given level of entropy, using stationary policies, one can induce a process whose entropy is greater than that level. In this case, we say that the maximum entropy of the MDP is unbounded. Additionally, if there exists a process with the maximum entropy, the entropy of such a process can be finite or infinite. Hence, before attempting to synthesize a policy that maximizes the entropy of an MDP, we first verify whether there exists a policy that attains the maximum entropy.

The contributions of this paper are fourfold. First, we provide necessary and sufficient conditions on the structure of the MDP under which the maximum entropy of the MDP is finite, infinite or unbounded. We also present a polynomial-time algorithm to check whether the maximum entropy of an MDP is finite, infinite or unbounded. Second, we present a polynomial-time algorithm based on a convex optimization problem to synthesize a policy that maximizes the entropy of an MDP. Third, we show that maximizing the entropy of an MDP with non-infinite maximum entropy is equivalent to maximizing the entropy of paths of the MDP. Lastly, we provide a procedure to obtain a policy that maximizes the entropy of an MDP subject to a general LTL specification.

The applications of this theoretical framework range from motion planning and stochastic traffic assignments to software security. In a motion planning scenario, for security purposes, an autonomous robot might need to randomize its paths while carrying out a mission [12], [13]. In such a scenario, a policy synthesized by the proposed methods both provides probabilistic guarantees on the completion of the mission and minimizes the predictability of the robot's paths through the use of online randomization mechanisms. Additionally, such

This work was supported in part by the grants DARPA # D19AP00004 and AFRL # FA9550-19-1-0169.

Y. Savas, M. Cubuktepe and U. Topcu are with the Department of Aerospace Engineering, University of Texas at Austin, TX, USA. E-mail: {yagiz.savas, mcubuktepe, utopcu}@utexas.edu

M. Ornik is with the Department of Aerospace Engineering and the Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, IL, USA. E-mail: mornik@illinois.edu

M. O. Karabag is with the Department of Electrical and Computer Engineering, University of Texas at Austin, TX, USA. E-mail: karabag@utexas.edu

a policy allows the robot to explore different parts of the environment [14], and behave robustly against uncertainties in the environment [15]. The proposed methods can also be used to distribute traffic assignments over a network, which is known as stochastic traffic assignments [16], as it promotes the use of different paths. Finally, as it is shown in [2], the maximum information that an adversary can leak from a (deterministic) software, which is modeled as an MDP, can be quantified by computing the maximum entropy of the MDP.

**Related Work.** A preliminary version [17] of this paper considered entropy maximization problem for MDPs subject to expected reward constraints. This considerably extended version includes an additional section establishing the relation between the maximum entropy of an MDP and the entropy of paths of the MDP, detailed proofs for all theoretical results, and additional numerical examples.

The computation of the maximum entropy of an MDP is first considered in [3], where the authors present a robust optimization problem to compute the maximum entropy for an MDP with finite maximum entropy. However, their approach does not allow to incorporate additional constraints due to the formulation of the problem. References [2] and [18] compute the maximum entropy of an MDP for special cases without providing a general algorithm.

The work [2] provides the necessary and sufficient conditions for an interval Markov chain (MC) to have a finite maximum entropy. Therefore, some of the results provided in this paper, e.g., the necessary and sufficient conditions for an MDP to have finite, unbounded or infinite maximum entropy, can be seen as an extension of the results given in [2].

In [19], [20], the authors study the problem of synthesizing a transition matrix with maximum entropy for an irreducible MC subject to graph constraints. The problem studied in this paper is considerably different from that problem since MDPs represent a more general model than MCs, and an MC induced from an MDP by a policy is not necessarily irreducible.

In [12], the authors maximize the entropy of a *policy* while keeping the expected total reward above a threshold. They claim that the entropy maximization problem is not convex. Their formulation is a special case of the convex optimization problem that we provide in this paper. Therefore, here, we also prove the convexity of their formulation.

The entropy of paths of absorbing MCs is discussed in [11], [16], [21]. The reference [14] establishes the equivalence between the entropy of paths and the entropy of an absorbing MC. We establish this relation for a general MC and show the connections to the maximum entropy of an MDP.

We also note that none of the above work discusses the unbounded and infinite maximum entropy for an MDP or considers LTL to specify desired system properties.

**Organization.** We provide the preliminary definitions and formal problem statement in Sections II and III, respectively. We analyze the properties of the maximum entropy of an MDP and present an algorithm to synthesize a policy that maximizes the entropy of an MDP in Section IV. The relation between the maximum entropy of an MDP and the entropy of paths is established in Section V. We present a procedure to synthesize a policy that maximizes the entropy of an MDP subject to

an LTL specification in Section VI. We provide numerical examples in Section VII and conclude with suggestions for future work in Section VIII. Proofs for all results are provided in Appendix A, and a procedure to synthesize a policy that maximizes the entropy of an MDP with infinite maximum entropy is presented in Appendix B.

## II. PRELIMINARIES

**Notation:** For a set  $S$ , we denote its power set and cardinality by  $2^S$  and  $|S|$ , respectively. For a matrix  $P \in \mathbb{R}^{n \times n}$ , we use  $P^k$  and  $P_{i,j}^k$  to denote the  $k$ -th power of  $P$  and the  $(i, j)$ -th component of the  $k$ -th power of  $P$ , respectively. All logarithms are to the base 2 and the set  $\mathbb{N}$  denotes  $\{0, 1, 2, \dots\}$ .

### A. Markov chains and Markov decision processes

**Definition 1:** A *Markov decision process* (MDP) is a tuple  $\mathcal{M} = (S, s_0, \mathcal{A}, \mathbb{P}, \mathcal{AP}, \mathcal{L})$  where  $S$  is a finite set of states,  $s_0 \in S$  is the initial state,  $\mathcal{A}$  is a finite set of actions,  $\mathbb{P}: S \times \mathcal{A} \times S \rightarrow [0, 1]$  is a transition function such that  $\sum_{t \in S} \mathbb{P}(s, a, t) = 1$  for all  $s \in S$  and  $a \in \mathcal{A}$ ,  $\mathcal{AP}$  is a set of atomic propositions, and  $\mathcal{L}: S \rightarrow 2^{\mathcal{AP}}$  is a function that labels each state with a subset of atomic propositions.

We denote the transition probability  $\mathbb{P}(s, a, t)$  by  $\mathbb{P}_{s,a,t}$ , and all available actions in a state  $s \in S$  by  $\mathcal{A}(s)$ . The set of successor states for a state action pair  $(s, a)$  is defined as  $Succ(s, a) := \{t \in S \mid \mathbb{P}_{s,a,t} > 0, a \in \mathcal{A}(s)\}$ . The *size of an MDP* is the number of triples  $(s, a, t) \in S \times \mathcal{A} \times S$  such that  $\mathbb{P}_{s,a,t} > 0$ .

A *Markov chain* (MC)  $\mathcal{C}$  is an MDP such that  $|\mathcal{A}| = 1$ . We denote the transition function (matrix) for an MC by  $\mathcal{P}$ , and the set of successor states for a state  $s \in S$  by  $Succ(s) := \{t \in S \mid \mathcal{P}_{s,t} > 0\}$ . The *expected residence time* in a state  $s \in S$  for an MC  $\mathcal{C}$  is defined as

$$\xi_s := \sum_{k=0}^{\infty} \mathcal{P}_{s_0,s}^k. \quad (1)$$

The expected residence time  $\xi_s$  represents the expected number of visits to state  $s$  starting from the initial state [22]. A state  $s \in S$  is *recurrent* for an MC if and only if  $\xi_s = \infty$ , and is *transient* otherwise; it is *stochastic* if and only if it satisfies  $|Succ(s)| > 1$ , and is *deterministic* otherwise; and it is *reachable* if and only if  $\xi_s > 0$ , and is *unreachable* otherwise.

**Definition 2:** A *policy* for an MDP  $\mathcal{M}$  is a sequence  $\pi = \{\mu_0, \mu_1, \dots\}$  where each  $\mu_k: S \times \mathcal{A} \rightarrow [0, 1]$  is a function such that  $\sum_{a \in \mathcal{A}(s)} \mu_k(s, a) = 1$  for all  $s \in S$ . A *stationary policy* is a policy of the form  $\pi = \{\mu, \mu, \dots\}$ . For an MDP  $\mathcal{M}$ , we denote the set of all policies and all stationary policies by  $\Pi(\mathcal{M})$  and  $\Pi^S(\mathcal{M})$ , respectively.

We denote the probability of choosing an action  $a \in \mathcal{A}(s)$  in a state  $s \in S$  under a stationary policy  $\pi$  by  $\pi_s(a)$ . For an MDP  $\mathcal{M}$ , a stationary policy  $\pi \in \Pi^S(\mathcal{M})$  induces an MC denoted by  $\mathcal{M}^\pi$ . We refer to  $\mathcal{M}^\pi$  as *induced MC* and specify the transition matrix for  $\mathcal{M}^\pi$  by  $\mathcal{P}^\pi$ , whose  $(s, t)$ -th component is given by

$$\mathcal{P}_{s,t}^\pi = \sum_{a \in \mathcal{A}(s)} \pi_s(a) \mathbb{P}_{s,a,t}. \quad (2)$$

Throughout the paper, we assume that for a given MDP  $\mathcal{M}$ , for any state  $s \in S$  there exists an induced MC  $\mathcal{M}^\pi$  for which the

state  $s$  is reachable. This is a standard assumption for MDPs [9], which ensures that each state in the MDP is reachable under some policy.

An infinite sequence  $\varrho^\pi = s_0 s_1 s_2 \dots$  of states generated in  $\mathcal{M}$  under a policy  $\pi \in \Pi(\mathcal{M})$  is called a *path*, starting from the initial state  $s_0$  and satisfies  $\sum_{a_k \in \mathcal{A}(s_k)} \mu_k(s_k)(a_k) \mathbb{P}_{s_k, a_k, s_{k+1}} > 0$  for all  $k \geq 0$ . Any finite prefix of  $\varrho^\pi$  that ends in a state is a finite path fragment. We define the set of all paths and finite path fragments in  $\mathcal{M}$  under the policy  $\pi$  by  $Paths^\pi(\mathcal{M})$  and  $Paths_{fin}^\pi(\mathcal{M})$ , respectively.

We use the standard probability measure over the outcome set  $Paths^\pi(\mathcal{M})$  [23]. For a path  $\varrho^\pi \in Paths^\pi(\mathcal{M})$ , let the sequence  $s_0 s_1 \dots s_n$  be the finite path fragment of length  $n$ , and let  $Paths^\pi(\mathcal{M})(s_0 s_1 \dots s_n)$  denote the set of all paths in  $Paths^\pi(\mathcal{M})$  starting with the prefix  $s_0 s_1 \dots s_n$ . The probability measure  $\Pr_{\mathcal{M}}^\pi$  defined on the smallest  $\sigma$ -algebra over  $Paths^\pi(\mathcal{M})$  that contains  $Paths^\pi(\mathcal{M})(s_0 s_1 \dots s_n)$  for all  $s_0 s_1 \dots s_n \in Paths_{fin}^\pi(\mathcal{M})$  is the unique measure that satisfies

$$\Pr_{\mathcal{M}}^\pi \{ Paths^\pi(\mathcal{M})(s_0 \dots s_n) \} = \prod_{0 \leq k < n} \sum_{a_k \in \mathcal{A}(s_k)} \mu_k(s_k)(a_k) \mathbb{P}_{s_k, a_k, s_{k+1}}. \quad (3)$$

### B. The entropy of stochastic processes

For a (discrete) random variable  $X$ , its support  $\mathcal{X}$  defines a countable sample space from which  $X$  takes a value  $x \in \mathcal{X}$  according to a probability mass function (pmf)  $p(x) := \Pr(X=x)$ . The *entropy* of a random variable  $X$  with countable support  $\mathcal{X}$  and pmf  $p(x)$  is defined as

$$H(X) := - \sum_{x \in \mathcal{X}} p(x) \log p(x). \quad (4)$$

We use the convention that  $0 \log 0 = 0$ . Let  $(X_0, X_1)$  be a pair of random variables with the joint pmf  $p(x_0, x_1)$  and the support  $\mathcal{X} \times \mathcal{X}$ . The *joint entropy* of  $(X_0, X_1)$  is

$$H(X_0, X_1) := - \sum_{x_0 \in \mathcal{X}} \sum_{x_1 \in \mathcal{X}} p(x_0, x_1) \log p(x_0, x_1), \quad (5)$$

and the *conditional entropy* of  $X_1$  given  $X_0$  is

$$H(X_1|X_0) := - \sum_{x_0 \in \mathcal{X}} \sum_{x_1 \in \mathcal{X}} p(x_0, x_1) \log p(x_1|x_0). \quad (6)$$

The definitions of the joint and conditional entropies extend to collection of  $k$  random variables as it is shown in [1]. A discrete *stochastic process*  $\mathbb{X}$  is a discrete time-indexed sequence of random variables, i.e.,  $\mathbb{X} = \{X_k \in \mathcal{X} : k \in \mathbb{N}\}$ .

**Definition 3:** (Entropy of a stochastic process) [24] The *entropy of a stochastic process*  $\mathbb{X}$  is defined as

$$H(\mathbb{X}) := \lim_{k \rightarrow \infty} H(X_0, X_1, \dots, X_k). \quad (7)$$

Note that this definition is different from the *entropy rate* of a stochastic process, which is defined as  $\lim_{k \rightarrow \infty} \frac{1}{k} H(X_0, X_1, \dots, X_k)$  when the limit exists [1]. The limit in (7) either converges to a non-negative real number or diverges to positive infinity [24].

An MC  $\mathcal{C}$  is equipped with a discrete stochastic process  $\{X_k \in S : k \in \mathbb{N}\}$  where each  $X_k$  is a random variable

over the state space  $S$ . For a given  $k$ -dimensional pmf  $p(s_0, s_1, \dots, s_k)$ , this process respects the *Markov property*, i.e.,  $p(s_k | s_{k-1}, \dots, s_0) = p(s_k | s_{k-1})$  for all  $k \in \mathbb{N}$ . Then, the *entropy of a Markov chain*  $\mathcal{C}$  is given by

$$H(\mathcal{C}) = H(X_0) + \sum_{i=1}^{\infty} H(X_i | X_{i-1}) \quad (8)$$

using (5), (6) and (7). Note that  $H(X_0) = 0$ , since we define an MC with a unique initial state.

For an MDP  $\mathcal{M}$ , a policy  $\pi \in \Pi(\mathcal{M})$  induces a discrete stochastic process  $\{X_k \in S : k \in \mathbb{N}\}$ . We denote the entropy of an MDP  $\mathcal{M}$  under a policy  $\pi \in \Pi(\mathcal{M})$  by  $H(\mathcal{M}, \pi)$ .

Using the next proposition, we restrict our attention to stationary policies for maximizing the entropy of an MDP.

**Proposition 1:** The following equality holds:

$$\sup_{\pi \in \Pi(\mathcal{M})} H(\mathcal{M}, \pi) = \sup_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi). \quad (9)$$

A proof for Proposition 1 is provided in Appendix A.

**Remark 1:** If the supremum in (9) is infinite, the set of stationary policies may not be sufficient to attain the supremum while a non-stationary policy can attain it. In particular, there exists a family of distributions that are defined over a countable support and have infinite entropy (see equation (7) in [25]). It can be shown that for some MDPs, there exists a non-stationary policy that induces a stochastic process with such a probability distribution, and hence, have infinite entropy, while stationary policies can only induce stochastic processes with finite entropies.<sup>1</sup>

**Definition 4:** (Maximum entropy of an MDP) The *maximum entropy of an MDP*  $\mathcal{M}$  is

$$H(\mathcal{M}) := \sup_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi). \quad (10)$$

A policy  $\pi^* \in \Pi^S(\mathcal{M})$  *maximizes* the entropy of an MDP  $\mathcal{M}$  if  $H(\mathcal{M}) = H(\mathcal{M}, \pi^*)$ . Finally, we define the properties of the maximum entropy of an MDP as follows.

**Definition 5:** (The properties of the maximum entropy) The maximum entropy of an MDP  $\mathcal{M}$  is

- *finite*, if and only if

$$H(\mathcal{M}) = \max_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi) < \infty; \quad (11)$$

- *infinite*, if and only if

$$H(\mathcal{M}) = \max_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi) = \infty; \quad (12)$$

- *unbounded*, if and only if the following two conditions hold.

$$(i) \quad H(\mathcal{M}) = \sup_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi) = \infty, \quad (13)$$

$$(ii) \quad H(\mathcal{M}, \pi) < \infty \text{ for all } \pi \in \Pi^S(\mathcal{M}). \quad (14)$$

Although it is not defined here, there is a fourth possible property which is unachievable finite maximum entropy, i.e.,  $\max_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi) < H(\mathcal{M}) < \infty$ . In Theorem 1, we show that it is not possible for the maximum entropy of an MDP to have this property.

<sup>1</sup>A preliminary version [17] of this paper relied on Proposition 36 from [3]. This proposition is not valid in general. Here, we provide the corrected results by defining the maximum entropy of an MDP over stationary policies.

### C. Linear temporal logic

We employ linear temporal logic (LTL) to specify tasks and refer the reader to [23] for the syntax and semantics of LTL.

An LTL formula is built up from a set  $\mathcal{AP}$  of atomic propositions, logical connectives such as conjunction ( $\wedge$ ) and negation ( $\neg$ ), and temporal modal operators such as always ( $\square$ ) and eventually ( $\diamond$ ). An infinite sequence of subsets of  $\mathcal{AP}$  defines an infinite *word*, and an LTL formula is interpreted over infinite words on  $2^{\mathcal{AP}}$ . We denote by  $w \models \varphi$  that a word  $w = w_0 w_1 w_2 \dots$  satisfies an LTL formula  $\varphi$ .

**Definition 6:** A *deterministic Rabin automaton* (DRA) is a tuple  $A = (Q, q_0, \Sigma, \delta, Acc)$  where  $Q$  is a finite set of states,  $q_0 \in Q$  is the initial state,  $\Sigma$  is the alphabet,  $\delta: Q \times \Sigma \rightarrow Q$  is the transition relation, and  $Acc \subseteq 2^Q \times 2^Q$  is the set of accepting state pairs.

A *run* of a DRA  $A$ , denoted by  $\sigma = q_0 q_1 q_2 \dots$ , is an infinite sequence of states in  $A$  such that for each  $i \geq 0$ ,  $q_{i+1} \in \delta(q_i, p)$  for some  $p \in \Sigma$ . A run  $\sigma$  is *accepting* if there exists a pair  $(J, K) \in Acc$  and an  $n \geq 0$  such that (i) for all  $m \geq n$  we have  $q_m \notin J$ , and (ii) there exists infinitely many  $k$  such that  $q_k \in K$ .

For any LTL formula  $\varphi$  built up from  $\mathcal{AP}$ , a DRA  $A_\varphi$  can be constructed with input alphabet  $2^{\mathcal{AP}}$  that accepts all and only words over  $\mathcal{AP}$  that satisfy  $\varphi$  [23].

For an MDP  $\mathcal{M}$  under a policy  $\pi$ , a path  $\varrho^\pi = s_0 s_1 \dots$  generates a word  $w = w_0 w_1 \dots$  where  $w_k = \mathcal{L}(s_k)$  for all  $k \geq 0$ . With a slight abuse of notation, we use  $\mathcal{L}(\varrho^\pi)$  to denote the word generated by  $\varrho^\pi$ . For an LTL formula  $\varphi$ , the set  $\{\varrho^\pi \in Paths^\pi(\mathcal{M}) : \mathcal{L}(\varrho^\pi) \models \varphi\}$  is measurable [23]. We define

$$\Pr_{\mathcal{M}}^\pi(s_0 \models \varphi) := \Pr_{\mathcal{M}}^\pi\{\varrho^\pi \in Paths^\pi(\mathcal{M}) : \mathcal{L}(\varrho^\pi) \models \varphi\}$$

as the probability of satisfying the LTL formula  $\varphi$  for an MDP  $\mathcal{M}$  under the policy  $\pi \in \Pi(\mathcal{M})$ .

### III. PROBLEM STATEMENT

The first problem we study concerns the synthesis of a policy that maximizes the entropy of an MDP.

**Problem 1: (Entropy Maximization)** For a given MDP  $\mathcal{M}$ , provide an algorithm to verify whether there exists a policy  $\pi^* \in \Pi^S(\mathcal{M})$  such that  $H(\mathcal{M}) = H(\mathcal{M}, \pi^*)$ . If such a policy exists, provide an algorithm to synthesize it. If it does not exist, provide a procedure to synthesize a policy  $\pi' \in \Pi^S(\mathcal{M})$  such that  $H(\mathcal{M}, \pi') \geq \ell$  for a given constant  $\ell$ .

For an MDP  $\mathcal{M}$ , the synthesis of a policy  $\pi' \in \Pi^S(\mathcal{M})$  such that  $H(\mathcal{M}, \pi') \geq \ell$  allows one to induce a stochastic process with the desired level of entropy, even if there exists no stationary policy that maximizes the entropy of  $\mathcal{M}$ .

In the second problem, we introduce linear temporal logic (LTL) specifications to the framework. In particular, we consider the problem of synthesizing a policy that induces a stochastic process with maximum entropy whose paths satisfy a given LTL formula with desired probability. The formal statement of the second problem is deferred to Section VI since it requires the introduction of additional notations.

### IV. ENTROPY MAXIMIZATION FOR MDPs

In this section, we focus on the entropy maximization problem. We refer to a policy as an *optimal* policy for an MDP if it maximizes the entropy of the MDP.

### A. The entropy of MCs versus MDPs

For an MC, the *local entropy* of a state  $s \in S$  is defined as

$$L(s) := - \sum_{t \in S} \mathcal{P}_{s,t} \log \mathcal{P}_{s,t}. \quad (15)$$

The following proposition characterizes the relationship between the local entropy of states and the entropy of an MC.

**Proposition 2:** (Theorem 1 in [2]) For an MC  $\mathcal{C}$ ,

$$H(\mathcal{C}) = \sum_{s \in S} L(s) \xi_s. \quad \triangleleft \quad (16)$$

An MC  $\mathcal{C}$  has a finite entropy if and only if all of its recurrent states have zero local entropy [2]. That is,  $H(\mathcal{C}) < \infty$  if and only if for all states  $s \in S$ ,  $\xi_s = \infty$  implies  $L(s) = 0$ . If the entropy of an MC is finite, each recurrent state  $s \in S$  has a single successor state, i.e.,  $|Succ(s)| = 1$ . Consequently, recurrent states have no contribution to the sum in (8). In this case, we take the sum in (16) only over the transient states.

For an MDP, different policies may induce stochastic processes with different entropies. For example, consider the MDP given in Fig. 1a and suppose that the action  $a_1$  at state  $s_0$  is taken with probability  $\varepsilon$ . If we let  $\varepsilon$  range over  $[0, \frac{1}{2}]$ , then the entropy of the resulting stochastic processes ranges over  $[0, 1]$ . The optimal policy for this MDP is  $\pi_{s_0}(a_1) = \pi_{s_0}(a_2) = 1/2$ , which uniformly randomizes actions.

Unlike the MDP given in Fig. 1a, the maximum entropy of an MDP is not generally achieved by a policy that chooses available actions at each state uniformly. For example, consider the MDP given in Fig. 1b. The optimal policy for this MDP is  $\pi_{s_0}(a_1) = 2/3$ ,  $\pi_{s_0}(a_2) = 1/3$ .

Examples given in Fig. 1 show that finding an optimal policy for an MDP may not be trivial. To analyze the maximum entropy of an MDP, we first obtain a compact representation of the maximum entropy as follows. For an MC  $\mathcal{M}^\pi$  induced from an MDP  $\mathcal{M}$  by a policy  $\pi \in \Pi^S(\mathcal{M})$ , let the expected residence time in a state  $s \in S$  be

$$\xi_s^\pi := \sum_{k=0}^{\infty} (\mathcal{P}^\pi)^k_{s_0, s}. \quad (17)$$

Additionally, let the local entropy of a state  $s \in S$  in  $\mathcal{M}^\pi$  be  $L^\pi(s) := - \sum_{t \in S} \mathcal{P}_{s,t}^\pi \log \mathcal{P}_{s,t}^\pi$ . Then, the maximum entropy of  $\mathcal{M}$  can be written as

$$H(\mathcal{M}) = \sup_{\pi \in \Pi^S(\mathcal{M})} \left[ \sum_{s \in S} \xi_s^\pi L^\pi(s) \right]. \quad (18)$$

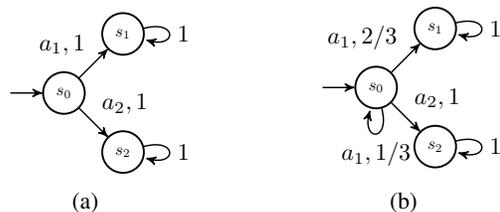


Fig. 1: Randomizing actions uniformly at each state may or may not achieve the maximum entropy. The optimal policy for the MDP given in (a) is  $\pi_{s_0}(a_1) = \pi_{s_0}(a_2) = 1/2$ , and for the MDP given in (b) is  $\pi_{s_0}(a_1) = 2/3$ ,  $\pi_{s_0}(a_2) = 1/3$ .

Note that the right hand side of (18) can still be infinite or unbounded. We analyze the properties of the maximum entropy of MDPs in the next section.

### B. Properties of the maximum entropy of MDPs

The maximum entropy of an MDP can be infinite or unbounded even for simple cases. For example, consider MDPs given in Fig. 2. For the MDP shown in Fig. 2a, let the action  $a_2$  be taken with probability  $\delta \in (0, 1]$  in state  $s_0$ . Then, the expected residence time  $\xi_{s_0}^\pi$  in state  $s_0$  is equal to  $\frac{1}{\delta}$ , and the entropy of the induced MC  $\mathcal{M}^\pi$  is given by

$$H(\mathcal{M}, \pi) = -\frac{(1-\delta)\log(1-\delta) + \delta\log(\delta)}{\delta}, \quad (19)$$

which satisfies  $H(\mathcal{M}, \pi) \rightarrow \infty$  as  $\delta \rightarrow 0$ . Note also that if  $\delta=0$ , the entropy of the induced MC is zero due to (16). Hence, the maximum entropy is unbounded, and there is no optimal stationary policy for this MDP.

For the MDP given in Fig. 2b, choosing a policy such that  $\pi_i(a_j) > 0$  for  $i=1, 2, j=1, 2$  yields  $\xi_{s_0}^\pi = \xi_{s_1}^\pi = \infty$  and  $L^\pi(s_0) > 0, L^\pi(s_1) > 0$ . Then, the maximum entropy of this MDP is infinite, and the maximum can be attained by any randomized policy.

Examples in Fig. 2 show that we should first verify the existence of optimal policies before attempting to synthesize them. We need the following definitions about the structure of MDPs to state the conditions that cause an MDP to have finite, infinite or unbounded maximum entropy.

A *directed graph* (digraph) is a tuple  $G=(V, E)$  where  $V$  is a set of vertices and  $E$  is a set of ordered pairs of vertices  $V$ . For a digraph  $G$ , a path  $v_1 v_2 \dots v_n$  from vertex  $v_1$  to  $v_n$  is a sequence of vertices such that  $(v_k, v_{k+1}) \in E$  for all  $1 \leq k < n$ . A digraph  $G$  is *strongly connected* if for every pair of vertices  $u, v \in V$ , there is a path from  $u$  to  $v$ , and  $v$  to  $u$ .

A *sub-MDP* of an MDP is a pair  $(C, D)$  where  $\emptyset \neq C \subseteq S$  and  $D: C \rightarrow 2^{\mathcal{A}}$  is a function such that (i)  $D(s) \subseteq \mathcal{A}(s)$  is non-empty for all  $s \in C$ , and (ii)  $s \in C$  and  $a \in D(s)$  imply that  $Succ(s, a) \subseteq C$ . An *end component* is a sub-MDP  $(C, D)$  such that the digraph induced by  $(C, D)$  is strongly connected.

**Definition 7:** A *maximal end component* (MEC)  $(C, D)$  in an MDP is an end component such that there is no end component  $(C', D')$  with  $(C, D) \neq (C', D')$ , and  $C \subseteq C'$  and  $D(s) \subseteq D'(s)$  for all  $s \in C$ .

A MEC  $(C, D)$  in an MDP is *bottom strongly connected* (BSC) if for all  $s \in C, \mathcal{A}(s) \setminus D(s) = \emptyset$ . For a given state  $s \in C$ , we define the set of all actions under which the MDP can leave the MEC  $(C, D)$  as  $D_0(s) := \{a \in \mathcal{A}(s) \mid Succ(s, a) \not\subseteq C\}$ . Note that in a BSC MEC  $(C, D), D_0(s) = \emptyset$  for all  $s \in C$ .

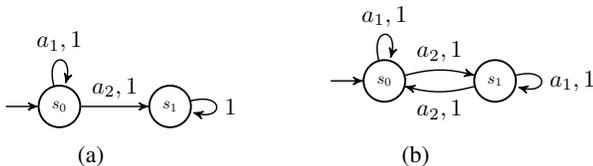


Fig. 2: Examples of MDPs with (a) unbounded maximum entropy and (b) infinite maximum entropy.

**Lemma 1:** For an MDP  $\mathcal{M}$  with MECs  $(C_i, D_i)$   $i=1, 2, \dots, n$ , let  $C := \cup_{i=1}^n C_i$  and  $D := \cup_{i=1}^n D_i$ . Then, there exists an induced MC  $\mathcal{M}^\pi$  for which a state  $s \in C$  is both stochastic and recurrent if and only if  $|\cup_{a \in D(s)} Succ(s, a)| > 1$ .  $\triangleleft$

**Theorem 1:** For an MDP  $\mathcal{M}$  with MECs  $(C_i, D_i)$   $i=1, 2, \dots, n$ , let  $C := \cup_{i=1}^n C_i$  and  $D := \cup_{i=1}^n D_i$ . Then, the following statements hold.

- (i)  $H(\mathcal{M})$  is infinite if and only if there exists an induced MC for which a state  $s \in C$  is both stochastic and recurrent.
- (ii)  $H(\mathcal{M})$  is unbounded if and only if  $|\cup_{a \in D(s)} Succ(s, a)| = 1$  for all  $s \in C$ , and there exists a MEC that is not bottom strongly connected.
- (iii)  $H(\mathcal{M})$  is finite if and only if it is not infinite and not unbounded.  $\triangleleft$

Proofs for above results can be found in Appendix A. Informally, Theorem 1 states that for an MDP to have finite maximum entropy, all recurrent states of all MCs that are induced from the MDP by a stationary policy should be deterministic. Although necessary conditions for the finiteness of the maximum entropy is quite restrictive, there are some special cases, such as stochastic shortest path (SSP) problems [26], where MDP structures actually satisfy the necessary conditions. Specifically, since all *proper* policies in SSP problems are guaranteed to reach an absorbing target state within finite time steps with probability 1, the problem of synthesizing a proper policy with maximum entropy has a finite solution.

The following corollary is due to Proposition 1, Theorem 1, and the definition of finite maximum entropy (11).

**Corollary 1:** If  $\sup_{\pi \in \Pi(\mathcal{M})} H(\mathcal{M}, \pi) < \infty$ , then we have

$$\sup_{\pi \in \Pi(\mathcal{M})} H(\mathcal{M}, \pi) = \max_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi). \quad (20)$$

We present Algorithm 1 which, for an MDP  $\mathcal{M}$ , verifies whether  $H(\mathcal{M})$  is finite, infinite or unbounded by checking the necessary conditions in Theorem 1. For  $\mathcal{M}$ , its MECs can be found in  $\mathcal{O}(|S|^2)$  time [23],  $Succ(s, a)$  can be found in  $\mathcal{O}(|S|^2|\mathcal{A}|)$  time, and the necessary conditions can be verified in  $\mathcal{O}(|S|)$  time since no state can belong to more than one MEC. Hence, Algorithm 1 runs in polynomial-time in the size of  $\mathcal{M}$ .

---

**Algorithm 1** Verify the properties of the maximum entropy.

---

**Require:**  $\mathcal{M}=(S, s_0, \mathcal{A}, \mathbb{P}, \mathcal{AP}, \mathcal{L})$

**Return:** R

Find: MECs  $(C_i, D_i), i = 1, \dots, n$

Find:  $Succ(s, a)$  for all  $s \in S, a \in \mathcal{A}(s)$

R :=  $\emptyset$ ;

**for**  $i=1, 2, \dots, n$  **do**

**for**  $s$  **in**  $C_i$  **do**

**if**  $|\cup_{a \in D_i(s)} Succ(s, a)| > 1$  **then**

      R := R  $\cup$  {infinite} ;

**if**  $\mathcal{A}(s) \setminus D_i(s) \neq \emptyset$  **then**

      R := R  $\cup$  {unbounded} ;

**if** infinite  $\in$  R **then** R=infinite

**else if** unbounded  $\in$  R **then** R=unbounded

**else** R=finite

---

### C. Policy synthesis

We now provide algorithms to synthesize policies that solve the entropy maximization problem.

1) *Finite maximum entropy*: We first modify a given MDP by making all states in its MECs absorbing.

**Proposition 3:** Let  $\mathcal{M}$  be an MDP such that  $H(\mathcal{M}) < \infty$ ,  $(C_i, D_i)$   $i=1, 2, \dots, n$  be MECs in  $\mathcal{M}$ ,  $C := \cup_{i=1}^n C_i$ , and  $\mathcal{M}'$  be the modified MDP that is obtained from  $\mathcal{M}$  by making all states  $s \in C$  absorbing, i.e., if  $s \in C$ , then  $\mathbb{P}_{s,a,s} = 1$  for all  $a \in \mathcal{A}(s)$  in  $\mathcal{M}'$ . Then, we have  $H(\mathcal{M}) = H(\mathcal{M}')$ .  $\triangleleft$

There is a one-to-one correspondence between the paths of  $\mathcal{M}$  and  $\mathcal{M}'$  since all states in the set  $C$  must have a single successor state in an MDP with finite maximum entropy due to Theorem 1. Moreover, for a given policy  $\pi' \in \Pi^S(\mathcal{M}')$  on  $\mathcal{M}'$ , the policy  $\pi \in \Pi^S(\mathcal{M})$  induced by  $\pi'$  on  $\mathcal{M}$  is the same policy with  $\pi'$ , i.e.  $\pi' = \pi$ . Therefore, we synthesize an optimal policy for  $\mathcal{M}$  by synthesizing an optimal policy for  $\mathcal{M}'$ .

We use the nonlinear programming problem in (21a)-(21g) to synthesize an optimal policy for  $\mathcal{M}'$ .

$$\text{maximize}_{\lambda(s,a), \lambda(s)} - \sum_{s \in S \setminus C} \sum_{t \in S} \eta(s, t) \log \left( \frac{\eta(s, t)}{\nu(s)} \right) \quad (21a)$$

subject to:

$$\nu(s) - \sum_{t \in S \setminus C} \eta(t, s) = \alpha(s) \quad \forall s \in S \setminus C \quad (21b)$$

$$\lambda(s) - \sum_{t \in S \setminus C} \eta(t, s) = \alpha(s) \quad \forall s \in C \quad (21c)$$

$$\eta(s, t) = \sum_{a \in \mathcal{A}(s)} \lambda(s, a) \mathbb{P}_{s,a,t} \quad \forall t \in S, \forall s \in S \setminus C \quad (21d)$$

$$\nu(s) = \sum_{a \in \mathcal{A}(s)} \lambda(s, a) \quad \forall s \in S \setminus C \quad (21e)$$

$$\lambda(s, a) \geq 0 \quad \forall a \in \mathcal{A}(s), \forall s \in S \setminus C \quad (21f)$$

$$\lambda(s) \geq 0 \quad \forall s \in C \quad (21g)$$

The decision variables in (21a)-(21c) are  $\lambda(s)$  for each  $s \in C$ , and  $\lambda(s, a)$  for each  $s \in S \setminus C$  and each  $a \in \mathcal{A}(s)$ . The function  $\alpha: S \rightarrow \{0, 1\}$  satisfies  $\alpha(s_0) = 1$  and  $\alpha(s) = 0$  for all  $s \in S \setminus \{s_0\}$ . Variables  $\eta(s, t)$  and  $\nu(s)$  are functions of  $\lambda(s, a)$ , and used just to simplify the notation.

The constraints (21b)-(21c) represent the balance between the ‘‘inflow’’ to and ‘‘outflow’’ from states. The constraints (21d) and (21e) are used to simplify the notation and define the variables  $\eta(s, t)$  and  $\nu(s)$ , respectively. The constraints (21f) and (21g) ensure that the expected residence time in the state-action pair  $(s, a)$  and the probability of reaching the state  $s$  is non-negative, respectively. We refer the reader to [22], [27] for further details about the constraints.

**Proposition 4:** The nonlinear program in (21a)-(21g) is convex.  $\triangleleft$

The above result indicates that a global maximum for the problem in (21a)-(21g) can be computed efficiently. We now introduce Algorithm 2 to synthesize an optimal policy for a given MDP with finite maximum entropy.

**Theorem 2:** Let  $\mathcal{M}$  be an MDP such that  $H(\mathcal{M}) < \infty$ ,  $(C_i, D_i)$   $i=1, 2, \dots, n$  be MECs in  $\mathcal{M}$ , and  $C := \cup_{i=1}^n C_i$ .

---

### Algorithm 2 Synthesize the maximum entropy policy

---

**Require:**  $\mathcal{M} = (S, s_0, \mathcal{A}, \mathbb{P}, \mathcal{AP}, \mathcal{L})$  and  $C$ .

**Return:** An optimal policy  $\pi^*$  for  $\mathcal{M}$

- 1: Form the modified MDP  $\mathcal{M}'$ .
  - 2: Solve (21a)-(21g) for  $(\mathcal{M}', C)$ , and obtain  $\lambda^*(s, a)$ .
  - 3: **for**  $s \in S$  **do**
    - if**  $s \notin C$  **then**
      - if**  $\sum_{a \in \mathcal{A}(s)} \lambda^*(s, a) > 0$  **then**

$$\pi_s^*(a) := \frac{\lambda^*(s, a)}{\sum_{a \in \mathcal{A}(s)} \lambda^*(s, a)}$$
      - else**

$$\pi_s^*(a) := 1 \text{ for an arbitrary } a \in \mathcal{A}(s),$$
    - else**

$$\pi_s^*(a) := 1 \text{ for an arbitrary } a \in \mathcal{A}(s).$$
- 

For the input  $(\mathcal{M}, C)$ , Algorithm 2 returns an optimal policy  $\pi^* \in \Pi^S(\mathcal{M})$  for  $\mathcal{M}$ , i.e.  $H(\mathcal{M}, \pi^*) = H(\mathcal{M})$ .  $\triangleleft$

Proofs for above results can be found in Appendix A. Computationally, the most expensive step of Algorithm 2 is to solve the convex optimization problem (21a)-(21g). A solution whose objective value is arbitrarily close to the optimal value of (21a) can be computed in time polynomial in the size of  $\mathcal{M}$  via interior-point methods [28], [29]. Hence, the time complexity of Algorithm 2 is polynomial in the size of  $\mathcal{M}$ .

2) *Unbounded maximum entropy*: There is no optimal policy for this case due to (13)-(14). Therefore, for a given MDP  $\mathcal{M}$  and a constant  $\ell$ , we synthesize a policy  $\pi \in \Pi^S(\mathcal{M})$  such that  $H(\mathcal{M}, \pi) \geq \ell$ . Let  $S_B$  be the union of all states in BSC MECs of  $\mathcal{M}$ , which can be found by using Algorithm 1. We modify the MDP  $\mathcal{M}$  by making all states  $s \in S_B$  absorbing and denote the modified MDP by  $\mathcal{M}'$ . It can be shown that  $H(\mathcal{M}') = H(\mathcal{M})$  by using arguments similar to the ones used in the proof of Proposition 3. As the first approach, we solve a convex feasibility problem. Specifically, we remove the objective in (21a) and add the constraint

$$- \sum_{s \in S \setminus S_B} \sum_{t \in S} \eta(s, t) \log \left( \frac{\eta(s, t)}{\nu(s)} \right) \geq \ell \quad (22)$$

to the constraints in (21b)-(21g). Then, we solve the resulting convex feasibility problem for  $(\mathcal{M}', S_B, \ell)$  and obtain the desired policy  $\pi$  by using the step 3 of Algorithm 2.

Recall from Theorem 1 that the unboundedness of the maximum entropy is caused by the existence of non-BSC MECs in  $\mathcal{M}'$ . In particular, we can induce MCs with arbitrarily large entropy by making the expected residence time in states contained in non-BSC MECs arbitrarily large. As the second approach, we bound the expected residence time in states  $s \in S \setminus S_B$  in  $\mathcal{M}'$  and relax this bound according to the desired level of entropy. Specifically, we add the constraint

$$\sum_{s \in S \setminus S_B} \sum_{a \in \mathcal{A}(s)} \lambda(s, a) \leq \Gamma \quad (23)$$

to the problem in (21a)-(21g). For the constraint (23),  $\Gamma \geq 0$  is a predefined value and limits the expected residence time

in states  $s \in S \setminus S_B$ . Let  $H_\Gamma(\mathcal{M}')$  denote the maximum entropy  $H(\mathcal{M}')$  of  $\mathcal{M}'$  subject to the constraint (23). Then, we have

$$H_\Gamma(\mathcal{M}') \geq H_{\Gamma'}(\mathcal{M}') \quad (24)$$

for  $\Gamma \geq \Gamma'$ , and  $H_\Gamma(\mathcal{M}') = \infty$  for  $\Gamma = \infty$ . Therefore, by choosing an arbitrarily large  $\Gamma$  value, we can synthesize a policy that induces an MC with arbitrarily large entropy.

3) *Infinite maximum entropy*: The procedure to synthesize an optimal policy for MDPs with infinite maximum entropy is very similar to the ones described in Sections IV-C1 and IV-C2. Therefore, we provide it in Appendix B.

## V. RELATING THE MAXIMUM ENTROPY OF AN MDP WITH THE PROBABILITY DISTRIBUTION OF PATHS

In this section, we establish a link between the maximum entropy of an MDP  $\mathcal{M}$  and the entropy of paths in an MC  $\mathcal{M}^\pi$  induced from  $\mathcal{M}$  by a stationary policy  $\pi \in \Pi^S(\mathcal{M})$ .

We begin with an example demonstrating the probability distribution of paths in an MC induced by a policy that maximizes the entropy of an MDP. Consider the MDP shown in Fig. 3a which is used in [2]. The policy that maximizes the entropy of the MDP is given by  $\pi_{s_0}(a_1) = 2/3$ ,  $\pi_{s_0}(a_2) = 1/3$ ,  $\pi_{s_1}(a_1) = \pi_{s_1}(a_2) = 1/2$ . The MC induced by this policy is shown in Fig. 3b. There are three paths that reach the MECs, i.e.,  $(\{s_3\}, \{a_1\})$  and  $(\{s_4\}, \{a_1\})$ , of the MDP, each of which is followed with probability  $1/3$  in the induced MC, i.e., the probability distribution of paths is uniform.

Note that for the example given in Fig. 3a, the optimal policy that maximizes the entropy of the MDP is randomized, and action-selection at each state is performed in an online manner. In particular, an agent that follows the optimal policy chooses its action at each stage according to the outcomes of an online randomization mechanism. Therefore, it does not commit to follow a specific path at any state.

To rigorously establish the relation, illustrated in Fig. 3a, between the maximum entropy of an MDP and the entropy of paths in an induced MC, we need the following definitions.

A *strongly connected component* (SCC)  $V \subseteq S$  in an MC  $\mathcal{M}^\pi$  induced by a policy  $\pi \in \Pi^S(\mathcal{M})$  is a maximal set of states in  $\mathcal{M}^\pi$  such that for any  $s, t \in V$ ,  $(\mathcal{P}^\pi)_{s,t}^n > 0$  for some  $n \in \mathbb{N}$ . A *bottom strongly connected component* (BSCC)  $S_B$  in  $\mathcal{M}^\pi$  is an SCC such that for all  $s \in S_B$ ,  $(\mathcal{P}^\pi)_{s,t}^n = 0$  for all  $n \in \mathbb{N}$  and for all  $t \in S \setminus S_B$ .

In this section, for an induced MC  $\mathcal{M}^\pi$ , we denote the probability of a path with the finite path fragment  $s_0 \dots s_n$  by

$$\mathcal{P}^\pi(s_0 \dots s_n) := \prod_{0 \leq k < n} \mathcal{P}_{s_k, s_{k+1}}^\pi, \quad (25)$$

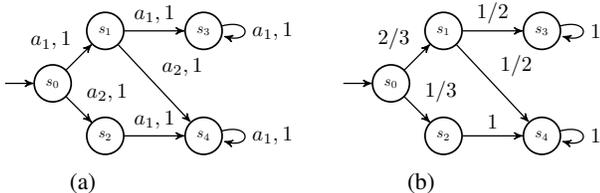


Fig. 3: (a) An MDP example [2]. (b) The MC induced by the policy that maximizes the entropy of the MDP.

which agrees with the probability measure introduced in Section II. Additionally, if the finite path fragment  $s_0 \dots s_n$  in  $\mathcal{M}^\pi$  satisfies  $s_0, \dots, s_{n-1} \notin S_B$  and  $s_n \in S_B$  for some  $S_B \subseteq S$ , we write  $s_0 \dots s_n \in (S \setminus S_B)^* S_B$ .

**Definition 8:** (Entropy of paths) Let  $\mathcal{M}^\pi$  be an MC induced from an MDP  $\mathcal{M}$  by a stationary policy  $\pi \in \Pi^S(\mathcal{M})$  and  $S_B \subseteq S$  be the union of all BSCCs in  $\mathcal{M}^\pi$ . For  $\mathcal{M}^\pi$ , the *entropy of the paths* that start from the initial state and reach a state in a BSCC in  $\mathcal{M}^\pi$  is defined as

$$H(\text{Paths}^\pi(\mathcal{M})) := - \sum_{s_0 \dots s_n \in T} \mathcal{P}^\pi(s_0 \dots s_n) \log \mathcal{P}^\pi(s_0 \dots s_n) \quad (26)$$

where  $T := \text{Paths}_{fin}^\pi(\mathcal{M}) \cap (S \setminus S_B)^* S_B$ .

A similar definition for the entropy of paths with fixed initial and final states can be found in [11], [21]. We note that

$$\sum_{s_0 \dots s_n \in T} \mathcal{P}^\pi(s_0 \dots s_n) = 1, \quad (27)$$

since any finite-state MC eventually reaches a BSCC [23]. The following lemma establishes a relation between the entropy of paths and the entropy of an induced MC.

**Lemma 2:** Let  $\mathcal{M}$  be an MDP such that  $H(\mathcal{M}, \pi) < \infty$  for any  $\pi \in \Pi^S(\mathcal{M})$ . Then, for any  $\pi \in \Pi^S(\mathcal{M})$ , we have

$$H(\text{Paths}^\pi(\mathcal{M})) = H(\mathcal{M}, \pi). \quad \triangleleft \quad (28)$$

A proof for Lemma 2 can be found in Appendix A. Finally, from the definition of the properties of the maximum entropy, Proposition 1 and Lemma 2, we conclude that, if an MDP  $\mathcal{M}$  has non-infinite maximum entropy, then we have

$$H(\mathcal{M}) = \sup_{\pi \in \Pi^S(\mathcal{M})} H(\text{Paths}^\pi(\mathcal{M})). \quad (29)$$

The equality in (29) states that, for an MDP with non-infinite maximum entropy, a policy that maximizes the entropy of the MDP induces an MC with maximum entropy of paths among all MCs that can be induced from the MDP. Moreover, considering (27), such a policy maximizes the randomness of all paths with non-zero probability in an induced MC.

## VI. CONSTRAINED ENTROPY MAXIMIZATION FOR MDPs

In this section, we consider the problem of maximizing the entropy of an MDP subject to an LTL constraint. We note that stationary policies are not sufficient to satisfy LTL constraints in general [23]. Therefore, to be consistent with our definition of maximum entropy (10), we first introduce the product MDP, over which LTL constraints are transformed into reachability constraints for which stationary policies are sufficient.

### A. Product MDP

We construct an MDP that captures all paths of an MDP  $\mathcal{M}$  satisfying an LTL specification  $\varphi$  by taking the product of  $\mathcal{M}$  and the DRA  $A_\varphi$  corresponding to the specification  $\varphi$ .

**Definition 9:** (Product MDP) Let  $\mathcal{M} = (S, s_0, \mathcal{A}, \mathbb{P}, \mathcal{AP}, \mathcal{L})$  be an MDP and  $A_\varphi = (Q, q_0, 2^{AP}, \delta, Acc)$  be a DRA. The product MDP  $\mathcal{M}_p = \mathcal{M} \otimes A_\varphi = (S_p, s_{0,p}, \mathcal{A}, \mathbb{P}_p, \mathcal{L}_p, Acc_p)$  is a tuple where

- $S_p = S \times Q$ ,
- $s_{0_p} = (s_0, q)$  such that  $q = \delta(q_0, \mathcal{L}(s_0))$ ,
- $\mathbb{P}_p((s, q), a, (s', q')) = \begin{cases} \mathbb{P}_{s, a, s'} & \text{if } q' = \delta(q, \mathcal{L}(s')) \\ 0 & \text{otherwise,} \end{cases}$
- $\mathcal{L}_p((s, q)) = \{q\}$ ,
- $Acc_p = \{(J_1^p, K_1^p), \dots, (J_k^p, K_k^p)\}$  where  $J_i^p = S \times J_i$  and  $K_i^p = S \times K_i$  for all  $(J_i, K_i) \in Acc$  and for all  $i=1, \dots, k$ .

The product MDP  $\mathcal{M}_p$  may contain unreachable states which can be found in time polynomial in the size of  $\mathcal{M}_p$  by graph search algorithms, e.g., breadth-first search. Such states have no effect in the analysis of MDPs, and hence, can be removed from the MDP. We hereafter assume that there is no unreachable state in  $\mathcal{M}_p$ .

There is a one-to-one correspondence between the paths of  $\mathcal{M}_p$  and  $\mathcal{M}$  [23]. Moreover, a similar one-to-one correspondence exists between policies on  $\mathcal{M}_p$  and  $\mathcal{M}$ . More precisely, for a given policy  $\pi^p = \{\mu_0^p, \mu_1^p, \dots\}$  on  $\mathcal{M}_p$ , we can construct a policy  $\pi = \{\mu_0, \mu_1, \dots\}$  on  $\mathcal{M}$  by setting  $\mu_i(s_i) = \mu_i^p((s_i, q_i))$ . For a given policy  $\pi^p \in \Pi^S(\mathcal{M}_p)$  on  $\mathcal{M}_p$ , the policy  $\pi \in \Pi(\mathcal{M})$  constructed in this way is a non-stationary policy [23].

Let  $\pi^p \in \Pi^S(\mathcal{M}_p)$  be a policy on  $\mathcal{M}_p$  and  $\pi \in \Pi(\mathcal{M})$  be the policy on  $\mathcal{M}$  constructed from  $\pi^p$  through the procedure explained above. The paths of the MDP  $\mathcal{M}$  under the policy  $\pi$  satisfies the LTL specification  $\varphi$  with probability of at least  $\beta$ , i.e.,  $\Pr_{\mathcal{M}}^{\pi}(s_0 \models \varphi) \geq \beta$ , if and only if the paths of the product MDP  $\mathcal{M}_p$  under the policy  $\pi^p$  reaches accepting MECs in  $\mathcal{M}_p$  with probability of at least  $\beta$  and stays there forever [23].

**Definition 10:** (Accepting MEC) A MEC  $(C, D)$  in a product MDP  $\mathcal{M}_p$  with the set of accepting state pairs  $Acc_p$  is an *accepting MEC* if for some  $(J^p, K^p) \in Acc_p$ ,  $J^p \not\subseteq C$  and  $K^p \subseteq C$ .

Informally, accepting MECs are sets of states where the system can remain forever, and where the set  $K_p$  is visited infinitely often and the set  $J_p$  is visited finitely often.

### B. Constrained Problem

In this section, we formally state the constrained entropy maximization problem. Recall that, for an MDP  $\mathcal{M}$ , the problem of synthesizing a policy  $\pi \in \Pi(\mathcal{M})$  that satisfies an LTL formula  $\varphi$  with probability of at least  $\beta$ , i.e.,  $\Pr_{\mathcal{M}}^{\pi}(s_0 \models \varphi) \geq \beta$ , is equivalent to the problem of synthesizing a policy  $\pi \in \Pi^S(\mathcal{M}_p)$  that reaches the accepting MECs in  $\mathcal{M}_p$  with probability of at least  $\beta$  and stays there forever.

Our objective is to synthesize a policy that induces a stochastic process with maximum entropy whose paths satisfy the given LTL specification with desired probability. To this end, we synthesize a policy  $\pi \in \Pi^S(\mathcal{M}_p)$  on  $\mathcal{M}_p$  that reaches the accepting MECs in  $\mathcal{M}_p$  with probability of at least  $\beta$  and stays there forever.

We first partition the set  $S_p$  of states of  $\mathcal{M}_p$  into three disjoint sets as follows. We let  $B$  be the set of all states in accepting MECs in  $\mathcal{M}_p$ , and  $S_0$  be the set of all states that have zero probability of reaching the set  $B$ . Finally, we let  $S_r = S_p \setminus \{B \cup S_0\}$  be the set of all states that are not in an accepting MEC in  $\mathcal{M}_p$  and have nonzero probability of reaching the set  $B$ . These sets can be found in time polynomial in the size of  $\mathcal{M}_p$  by graph search algorithms [23].

**Problem 2: (Constrained Entropy Maximization)** For a product MDP  $\mathcal{M}_p$ , verify whether there exists a policy  $\pi^* \in \Pi^S(\mathcal{M}_p)$  that solves the following problem:

$$\underset{\pi \in \Pi^S(\mathcal{M}_p)}{\text{maximize}} \quad H(\mathcal{M}_p, \pi) \quad (30a)$$

$$\text{subject to:} \quad \Pr_{\mathcal{M}_p}^{\pi}(s_0 \models \diamond B) \geq \beta \quad (30b)$$

where  $\Pr_{\mathcal{M}_p}^{\pi}(s_0 \models \diamond B)$  denotes the probability of reaching the set  $B$  from the initial state in  $\mathcal{M}_p$  under the policy  $\pi$ . If such a policy exists, provide an algorithm to synthesize it. If it does not exist, provide a procedure to synthesize a policy  $\pi' \in \Pi^S(\mathcal{M}_p)$  such that  $\Pr_{\mathcal{M}_p}^{\pi'}(s_0 \models \diamond B) \geq \beta$  and  $H(\mathcal{M}_p, \pi') \leq \ell$  for a given constant  $\ell$ .

Note that if a policy that solves the problem in (30a)-(30b) chooses the actions in states  $s \in B$  such that they form a BSCC in the induced MC, then the resulting policy ensures that the paths of the induced MC visit the states inside the set  $B$  infinitely often and thus satisfies  $\varphi$  [23].

### C. Policy synthesis

In this section, for a product MDP  $\mathcal{M}_p$  and its state partition  $S_p = B \cup S_0 \cup S_r$ , we assume that  $0 < \beta \leq \max_{\pi \in \Pi(\mathcal{M})} \Pr_{\mathcal{M}_p}^{\pi}(s_0 \models \diamond B)$ , which can be verified in polynomial time by solving a linear optimization problem as shown in [22], [23]. We refer to a policy  $\pi \in \Pi^S(\mathcal{M}_p)$  as an *optimal policy* if it is a solution to the problem in (30a)-(30b) and chooses the actions in states  $s \in B$  such that they form a BSCC in the induced MC.

For the synthesis of an optimal policy, we consider three cases according to the maximum entropy  $H(\mathcal{M}_p)$  of  $\mathcal{M}_p$ , namely, finite, unbounded and infinite.

1) *Finite maximum entropy:* Let  $(C_i, D_i)$   $i=1, 2, \dots, n$  be the MECs in  $\mathcal{M}_p$ ,  $C := \bigcup_{i=1}^n C_i$ , and  $D := \bigcup_{i=1}^n D_i$ . We form the modified product MDP  $\mathcal{M}'_p$  by making all states  $s \in C$  absorbing in  $\mathcal{M}_p$ . We have  $H(\mathcal{M}'_p) = H(\mathcal{M}_p)$  due to Proposition 3. Recall that for a state  $s \in C$ , the variable  $\lambda(s)$  in (21a)-(21g) represents the probability of reaching the state  $s$  from the initial state [22]. Hence, we append the constraint

$$\sum_{s \in B} \lambda(s) \geq \beta \quad (31)$$

to the problem in (21a)-(21g) in order to obtain a policy that induces an MC whose paths satisfy  $\varphi$  with probability of at least  $\beta$ . Noting that  $\beta \leq \sum_{s \in B} \lambda(s) \leq \max_{\pi \in \Pi(\mathcal{M})} \Pr_{\mathcal{M}_p}^{\pi}(s_0 \models \diamond B)$ , the resulting optimization problem always has a solution since its feasible set constitutes a closed compact set when the product MDP has finite maximum entropy.

The procedure to obtain a policy  $\pi_p^* \in \Pi^S(\mathcal{M}_p)$  that solves the problem in (30a)-(30b) for  $\mathcal{M}_p$  with finite maximum entropy is as follows. First, we find MECs  $(C, D)$  in  $\mathcal{M}_p$  and form the modified MDP  $\mathcal{M}'_p$  by making all states  $s \in C$  absorbing. Second, we solve the problem in (21a)-(21g) for  $(\mathcal{M}'_p, C, \beta)$  with the additional constraint (31). Finally, we use step 3 of Algorithm 2 to synthesize the policy  $\pi_p^* \in \Pi(\mathcal{M}_p)$ . Note that the constructed policy ensures that, once reached, the system stays in the set  $B$  forever, since all MECs in  $\mathcal{M}_p$  with finite maximum entropy are bottom strongly connected.

2) *Unbounded maximum entropy*: In this case, the product MDP  $\mathcal{M}_p$  contains a non-BSC MEC due to Theorem 1. We assume that there is only one non-BSC MEC in  $\mathcal{M}_p$ , and it is contained in  $S_r$ . We first form the modified product MDP  $\mathcal{M}'_p$  by making all states in BSC MECs in  $\mathcal{M}_p$  absorbing. Note that  $H(\mathcal{M}'_p)=H(\mathcal{M}_p)$ . Let  $S_B$  denote the union of all absorbing states in  $\mathcal{M}'_p$ . We verify the existence of a solution to the problem in (30a)-(30b) by solving the problem in (21a)-(21g) for  $(\mathcal{M}'_p, S_B, \beta)$  with the additional constraint (31). If the optimum value of the resulting problem is bounded, then we synthesize an optimal policy through step 3 of Algorithm 2. If it is not bounded, then there exists no optimal policy, in which case for a given constant  $\ell$ , we synthesize a policy  $\pi^* \in \Pi^S(\mathcal{M}_p)$  such that  $H(\mathcal{M}_p, \pi^*) \geq \ell$  and  $\Pr_{\mathcal{M}_p}^{\pi^*}(s_0 \models \diamond B) \geq \beta$  by employing two different approaches.

As the first approach, we solve a convex feasibility problem. Specifically, for the problem in (21b)-(21g), we remove the objective (21a) and append the constraints (22) and (31) to the problem. Then, we solve the resulting convex feasibility problem for  $(\mathcal{M}'_p, S_B, \ell, \beta)$ , and using step 3 of Algorithm 2, obtain a policy  $\pi_p^* \in \Pi^S(\mathcal{M}_p)$  such that  $H(\mathcal{M}_p, \pi_p^*) \geq \ell$  and  $\Pr_{\mathcal{M}_p}^{\pi_p^*}(s_0 \models \diamond B) \geq \beta$ .

The second approach to obtain an induced MC with arbitrarily large entropy, whose paths satisfy the LTL specification with desired probability, is to bound the expected residence time in states  $s \in S_p \setminus S_B$  and relax this bound according to the desired level of entropy. Specifically, we solve the problem in (21a)-(21g) for  $(\mathcal{M}'_p, S_B, \beta, \Gamma)$  together with the constraints (23) and (31), where  $\Gamma$  is as defined in Section IV-C2. Then, by choosing an arbitrarily large  $\Gamma$  value, we obtain an induced MC with the desired level of entropy.

Finally, to ensure that the paths of the MC that is induced by the synthesized policy satisfies the LTL specification  $\varphi$  with desired probability, we choose actions in states  $s \in B$  such that  $\text{Succ}(s) \subseteq B$ .

3) *Infinite maximum entropy*: For product MDPs with infinite maximum entropy, the verification of the existence and the synthesis of an optimal policy are achieved by procedures that are very similar to the ones presented in Sections VI-C1 and VI-C2. Hence, we provide the analysis for product MDPs with infinite maximum entropy in Appendix B.

## VII. EXAMPLES

In this section, we illustrate the proposed methods on different motion planning scenarios. All computations are run on a 2.2 GHz dual core desktop with 8 GB RAM. All optimization problems are solved by using the splitting conic solver (SCS) [30] in CVXPY [31]. For all LTL specifications, we construct deterministic Rabin automata using ltl2dstar [32].

In most motion planning scenarios, an agent can return to its current position by following different paths. Therefore, in general, the maximum entropy of an MDP that models the motion of an agent is either unbounded or infinite. However, as explained in Section IV-C and shown in the following examples, a policy that induces a stochastic process with an arbitrarily large entropy can easily be obtained by introducing constraints on the expected residence time in certain states. Additional motion planning examples are provided in [17].

### A. Relation between entropy and exploration

Randomizing an agent's paths while ensuring the completion of a task is important for achieving a better exploration of the environment [14] and obtaining a robust behavior against transition perturbations [15]. In this example, we demonstrate how the proposed method randomizes the agent's paths depending on the expected time until the completion of the task.

*Environment*: We consider the grid world shown in Fig. 4 (left). The agent starts from the brown state. The red and green states are absorbing, i.e., once entered those states cannot be left. The agent has four actions in all other states, namely left, right, up and down. At each state, a transition to the chosen direction occurs with probability (w.p.) 0.7, and the agent slips to each adjacent state in the chosen direction w.p. 0.15. If the adjacent state in the chosen direction is a wall, e.g. up in brown state, a transition to the chosen direction occurs w.p. 0.85. If the state in the chosen direction is a wall, e.g., left in brown state, the agent stays in the same state w.p. 0.7 and moves to each adjacent state w.p. 0.15.

*Task*: The agent's task is to reach and stay in the green state, labeled as  $T$ , while avoiding the red states, labeled as  $B$ . Formally, the task is  $\varphi = \square \neg B \wedge \diamond \square T$ .

We form the product MDP for the given task. It has 484 states, 1196 transitions, 10 MECs, and the average number of states in each MEC is 23. We require the agent to complete the task w.p. 1, i.e.,  $\Pr_{\mathcal{M}}^{\pi}(s_0 \models \varphi) = 1$ . The maximum entropy of the product MDP subject to the LTL constraint is unbounded. The minimum expected time  $\Gamma$  required to complete the task  $\varphi$  is roughly 14 time steps, which can be computed by replacing the objective in (21a) with "minimize  $\sum_{s \in S_r} \sum_{a \in \mathcal{A}(s)} \lambda(s, a)$ " and appending (31) to the constraints in (21b)-(21g).

We synthesize two policies for two different expected times until the completion of the task. First, we synthesize a policy by requiring the agent to complete the task as fast as possible, i.e.,  $\Gamma = 14$  time steps. Then, we synthesize a policy by allowing the agent to spend more time in the environment until the completion of the task, i.e.,  $\Gamma = 60$  time steps. Solving the convex optimization problems take 122 and 166 seconds for  $\Gamma = 14$  and  $\Gamma = 60$  time steps, respectively.

The expected residence time in states for the induced MCs are shown in Fig. 5. We remind the reader that the environment is given in Fig. 4 (left). When the agent is given the minimum time  $\Gamma = 14$  time steps (left) to complete the task, it follows only the shortest paths, and therefore, cannot explore the

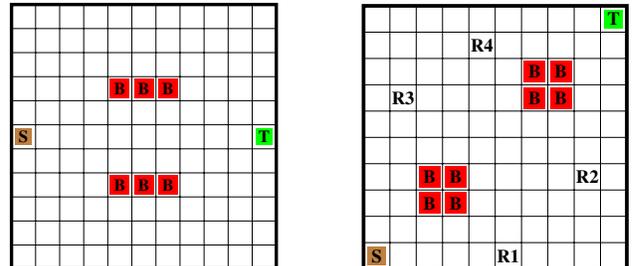


Fig. 4: Grid world environments. The brown (S) and green (T) states are the initial and target states, respectively. The red (B) states are absorbing.

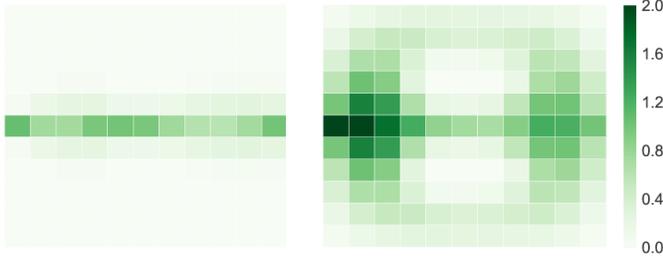


Fig. 5: The expected residence time in states for different expected times  $\Gamma$  until the completion of the task *in the same environment*. (Left)  $\Gamma=14$  time steps, i.e., the minimum time required to complete the task with probability 1. (Right)  $\Gamma=60$ .

environment. On the other hand, as it is allowed to spend more time, i.e.,  $\Gamma=60$  time steps (right), in the environment, it visits different states more often and utilizes different paths to complete the task. Consequently, the synthesized policy enables the continual exploration of the environment while ensuring the completion of the task.

### B. Relation between entropy and predictability

In this example, we consider an agent whose aim is to complete a task while leaking minimum information about its paths to an observer. We illustrate how the restrictions applied to the agent’s paths by the task affect the predictability.

*Environment:* We consider the grid world shown in Fig. 4 (right). The agent starts from the brown (S) state. The red (B) states and green (T) state are absorbing. The agent has four actions in all other states, namely left, right, up and down. A transition to the chosen direction occurs w.p. 1 if the state in that direction is not a wall. If it is a wall, e.g., left direction in brown state, the agent stays in the same state w.p. 1.

*Tasks:* We consider five increasingly restrictive task specifications for the agent which are listed in Table I. The first task  $\varphi_1$  is to reach and stay in the T state while avoiding all red states. The second task  $\varphi_2$  requires the agent to visit R4 state before completing the first task. The third task  $\varphi_3$  requires the agent to visit R3 state before completing the second task and so on.

TABLE I: The agent’s tasks.

$\varphi_1 = \square \neg Red \wedge \diamond \square T$
$\varphi_2 = \square \neg Red \wedge \diamond R4 \wedge \diamond \square T$
$\varphi_3 = \square \neg Red \wedge \diamond (R3 \wedge \diamond R4) \wedge \diamond \square T$
$\varphi_4 = \square \neg Red \wedge \diamond (R2 \wedge \diamond (R3 \wedge \diamond R4)) \wedge \diamond \square T$
$\varphi_5 = \square \neg Red \wedge \diamond (R1 \wedge \diamond (R2 \wedge \diamond (R3 \wedge \diamond R4))) \wedge \diamond \square T$

*Observer:* There is an observer that aims to predict the agent’s paths in the environment. The observer is aware of the agent’s task, knows the transition probabilities exactly, and runs yes-no probes in each state to determine the successor state of the agent, i.e., probes that return an answer yes if the agent moves to the predicted successor state and no otherwise. The average number of yes-no probes run in a state is the expected number of observations needed by the observer to determine the correct successor state in that state [12]. The observer uses the Huffman procedure [33] to minimize the required number of probes. Let  $\mathcal{P}_s = (\mathcal{P}_{s,1}, \mathcal{P}_{s,2}, \dots, \mathcal{P}_{s,n})$  be the

transition probabilities from state  $s$  to successor states sorted in decreasing order. The number of yes-no probes run in state  $s$  is denoted by  $\Upsilon_s = \mathcal{P}_{s,1} + \dots + (n-1)\mathcal{P}_{s,n-1} + (n-1)\mathcal{P}_{s,n}$ . The expected number of observations required to determine the agent’s path is given by  $O_{avg} = \sum_s \xi_s \Upsilon_s$ , which weighs the required number of probes in each state with the expected residence time in the state. We refer the reader to [12] for further details about the observer model.

We form product MDPs for all tasks. The product MDP with the maximum number of states and transitions is the one constructed for the task  $\varphi_5$ . It has 800 states, 2138 transitions, 12 MECs, and the average number of states in each MEC is 29. For each task, we require the agent to complete the task w.p. 1. The maximum entropy of all product MDPs subject to corresponding LTL constraints are unbounded. We bound the expected time until the completion of any task by taking  $\Gamma=33$  time steps, which is the minimum expected time required to complete the task  $\varphi_5$ , i.e., the most restrictive task. For each task, we synthesize a policy for the agent using the procedure explained in Section VI. The longest computation time, which is for  $\varphi_5$ , is 15.2 seconds.

The entropy of Markov chains induced by the synthesized policies and the average number of observations required to predict the agent’s paths are shown in Fig. 6. As the task imposes more restrictions on the agent’s paths, the entropy of the induced MC decreases and the prediction requires less observations in average. Additionally, as can be seen in Fig. 6, the required numbers of observation for  $\varphi_4$  and  $\varphi_5$  are significantly smaller than those for  $\varphi_1$ ,  $\varphi_2$  and  $\varphi_3$ . This decrease is due to that these tasks force the agent to follow an “S-shaped” path in a restricted time, i.e.  $\Gamma = 33$  time steps. For these tasks, although the agent still randomizes its paths to some extent, better predictability results cannot be obtained due to time restrictions.

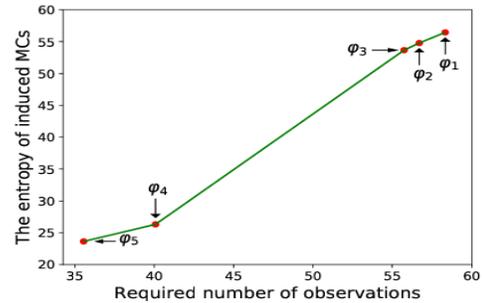


Fig. 6: The relation between the maximum entropy of an MDP subject to an LTL constraint and the required number of observations to predict the agent’s paths.

### C. Predictability in a randomly generated MDP

In this example, we investigate the relation between the probability of completing a task and the predictability of paths. We also evaluate the proposed algorithm against the algorithms introduced in [12].

*Environment:* We generate an MDP with 200 states, where each state has 8 randomly selected successor states. We choose four states, make them absorbing, and label three of them as “unsafe” states and the remaining one as the “target” state.

The agent has 5 actions at each state, for which the transition probabilities to successor states are assigned randomly.

*Task:* The agent’s task is to reach the target state while avoiding the unsafe states, i.e.,  $\varphi = \square \neg \text{unsafe} \wedge \diamond \square \text{target}$ .

*Observer:* We use the same observer model introduced in Section VII-B.

*Policies:* We compare the proposed method with weighted maximum entropy (*WME*) and binary search for randomization linear programming (*BRLP*) algorithms which are introduced in [12] for randomizing an agent’s policy to minimize predictability. We note that in [12], the authors claim that *WME* algorithm is non-convex and cannot be solved in polynomial time. However, its convexity can be proven by Proposition 4 since it solves a special case of the convex optimization problem given in (21), i.e., it is equivalent to problem in (21) when transition probabilities are either 0 or 1. We refer the reader to [12] for further details about the *WME* and *BRLP* algorithms.

We form the product MDP. It has 800 states, 2172 transitions and 5 MECs each of which contains a single state. The maximum probability of completing the task  $\varphi$  is obtained as  $\beta=0.75$  by solving a linear programming problem introduced in [22]. The maximum entropy of the product MDP subject to the LTL constraint  $\Pr_{\mathcal{M}}^{\pi}(s_0 \models \varphi) \geq \beta$  is unbounded for all  $\beta > 0$ . We fix the expected time until the completion of the task to  $\Gamma=200$  time steps, and synthesize policies for different values of  $\beta$ . Solving the optimization problems takes at most 150, 155, and 92 seconds for the proposed method, *WME* and *BRLP* algorithms, respectively.

The required number of observations to predict the agent’s paths for different  $\beta$  values are shown in Fig. 7. As the probability of completing the task decreases, the randomness of the agent’s paths increases and the prediction requires more observations in average. Therefore, there is a trade-off between the probability of satisfying the task and the randomness of the paths. Additionally, the proposed method (green) requires two times more observations than the *BRLP* algorithm (red) when  $\beta=0.5$ . Note also that the *WME* algorithm (blue) cannot achieve better predictability results than the proposed method because it does not exploit the inherent stochasticity in the environment and rely solely on the randomization of the agent’s actions to generate unpredictable paths.

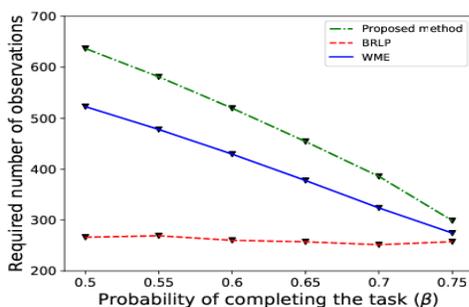


Fig. 7: The trade-off between the probability of completing the task and the predictability of paths. *BRLP* and *WME* are algorithms proposed in [12] to randomize the agent’s actions.

## VIII. CONCLUSIONS AND FUTURE WORK

We showed that the maximum entropy of an MDP can be either finite, infinite or unbounded, and presented an algorithm to verify the property of the maximum entropy for a given MDP. We presented an algorithm, based on a convex optimization problem, to synthesize a policy that maximizes the entropy of an MDP. For MDPs with non-infinite maximum entropy, we established the equivalence between the maximum entropy of an MDP and the maximum entropy of paths in the MDP. Finally, we provided a procedure to obtain a policy that maximizes the entropy of an MDP while ensuring the satisfaction of a temporal logic specification with desired probability.

An interesting future direction is to include adversaries to the framework by modeling the problem as a two-player game. Being informed about the aims and capabilities of rational/irrational adversaries in the environment, an agent may want to explore its environment while avoiding the threats caused by adversaries. Another future direction may be to extend this work to multi-agent scenarios by describing the tasks, and communication and coordination constraints between the agents as temporal logic specifications.

## REFERENCES

- [1] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, Inc., 2006.
- [2] F. Biondi, A. Legay, B. F. Nielsen, and A. Wasowski, “Maximizing entropy over Markov processes,” *Journal of Logical and Algebraic Methods in Programming*, vol. 83, no. 5, pp. 384 – 399, 2014.
- [3] T. Chen and T. Han, “On the complexity of computing maximum entropy for Markovian models,” in *International Conference on Foundation of Software Technology and Theoretical Computer Science*, vol. 29, 2014, pp. 571–583.
- [4] M. E. Chamie, Y. Yu, B. Acikmese, and M. Ono, “Controlled Markov processes with safety state constraints,” *IEEE Transactions on Automatic Control*, 2018.
- [5] J. Fu and U. Topcu, “Probably approximately correct MDP learning and control with temporal logic constraints,” *arXiv:1404.7073 [cs.SY]*, 2014.
- [6] A. Pnueli, “The temporal semantics of concurrent programs,” in *Proceedings of the International Symposium on Semantics of Concurrent Computation*, London, UK, UK, 1979, pp. 1–20.
- [7] L. Tan, O. Sokolsky, and I. Lee, “Specification-based testing with linear temporal logic,” in *Proceedings of the 2004 IEEE International Conference on Information Reuse and Integration, 2004. IRI 2004.*, 2004, pp. 493–498.
- [8] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, “Where’s Waldo? sensor-based temporal logic motion planning,” in *International Conference on Robotics and Automation*, 2007, pp. 3116–3121.
- [9] X. Ding, S. L. Smith, C. Belta, and D. Rus, “Optimal control of Markov decision processes with linear temporal logic constraints,” *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1244–1257, 2014.
- [10] Arpit and A. Kumar, “Verification of elgamal algorithm cryptographic protocol using linear temporal logic,” in *2011 International Conference on Multimedia Technology*, 2011, pp. 6662–6665.
- [11] L. Ekroot and T. M. Cover, “The entropy of Markov trajectories,” *IEEE Transactions on Information Theory*, vol. 39, no. 4, pp. 1418–1421, 1993.
- [12] P. Paruchuri, M. Tambe, F. Ordóñez, and S. Kraus, “Security in multiagent systems by policy randomization,” in *Joint Conference on Autonomous Agents and Multiagent Systems*, 2006, pp. 273–280.
- [13] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordóñez, and S. Kraus, “Coordinating randomized policies for increasing security of agent systems,” *Information Technology and Management*, vol. 10, no. 1, pp. 67–79, 2009.
- [14] M. Saerens, Y. Achbany, F. Fous, and L. Yen, “Randomized shortest-path problems: Two related models,” *Neural Computation*, vol. 21, no. 8, pp. 2363–2404, 2009.

- [15] T. Haarnoja, H. Tang, P. Abbeel, and S. Levine, "Reinforcement learning with deep energy-based policies," in *International Conference on Machine Learning*, vol. 70, 2017, pp. 1352–1361.
- [16] T. Akamatsu, "Cyclic flows, Markov process and stochastic traffic assignment," *Transportation Research Part B: Methodological*, vol. 30, no. 5, pp. 369 – 386, 1996.
- [17] Y. Savas, M. Ornik, M. Cubuktepe, and U. Topcu, "Entropy maximization for constrained Markov decision processes," in *Allerton Conference on Communication, Control, and Computing*, 2018, pp. 911–918.
- [18] F. Biondi, A. Legay, P. Malacaria, and A. Wasowski, "Quantifying information leakage of randomized protocols," *Theoretical Computer Science*, vol. 597, no. C, pp. 62–87, 2015.
- [19] M. George, S. Jafarpour, and F. Bullo, "Markov chains with maximum entropy for robotic surveillance," *IEEE Transactions on Automatic Control*, 2018.
- [20] X. Duan, M. George, and F. Bullo, "Markov chains with maximum return time entropy for robotic surveillance," *arXiv:1803.07705 [math.OC]*, 2018.
- [21] M. Kafsi, M. Grossglauser, and P. Thiran, "The entropy of conditional Markov trajectories," *IEEE Transactions on Information Theory*, vol. 59, no. 9, pp. 5577–5583, 2013.
- [22] K. Etessami, M. Kwiatkowska, M. Y. Vardi, and M. Yannakakis, "Multi-objective model checking of Markov decision processes," in *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, 2007, pp. 50–65.
- [23] C. Baier and J.-P. Katoen, *Principles of Model Checking*. MIT Press, 2008.
- [24] F. Biondi, *Markovian Processes for Quantitative Information Leakage*. PhD thesis, IT University of Copenhagen, 2014.
- [25] V. Baccetti and M. Visser, "Infinite Shannon entropy," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2013, no. 04, p. P04010, 2013.
- [26] D. P. Bertsekas and J. N. Tsitsiklis, *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- [27] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, 1st ed. New York, NY, USA: John Wiley & Sons, Inc., 1994.
- [28] S. A. Serrano, *Algorithms for unsymmetric cone optimization and implementation for problems with the exponential cone*. PhD thesis, Stanford University, 2015.
- [29] Y. Nesterov and A. Nemirovski, *Interior-Point Polynomial Algorithms in Convex Programming*. Society for Industrial and Applied Mathematics, 1994.
- [30] B. O'Donoghue, E. Chu, N. Parikh, and S. Boyd, "Conic optimization via operator splitting and homogeneous self-dual embedding," *Journal of Optimization Theory and Applications*, vol. 169, no. 3, pp. 1042–1068, June 2016.
- [31] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *Journal of Machine Learning Research*, vol. 17, no. 83, pp. 1–5, 2016.
- [32] J. Klein, "ltl2dstar- LTL to deterministic Streett and Rabin automata," <http://www.ltl2dstar.de/>, 2007 [Online].
- [33] D. A. Huffman, "A method for the construction of minimum-redundancy codes," *Proceedings of the IRE*, vol. 40, no. 9, pp. 1098–1101, 1952.
- [34] D. P. Bertsekas and J. N. Tsitsiklis, "An analysis of stochastic shortest path problems," *Mathematics of Operations Research*, vol. 16, no. 3, pp. 580–595, 1991.
- [35] C. Baier, J. Klein, L. Leuschner, D. Parker, and S. Wunderlich, "Ensuring the reliability of your model checker: Interval iteration for Markov decision processes," in *International Conference on Computer Aided Verification*, 2017, pp. 160–180.
- [36] R. Serfozo, *Basics of Applied Stochastic Processes*. Springer, 2009.
- [37] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

## APPENDIX A

We first define the hitting probability of a set of states in an MC  $\mathcal{M}^\pi$  induced from an MDP  $\mathcal{M}$  by a policy  $\pi \in \Pi^S(\mathcal{M})$ . If a sequence  $s_k s_{k+1} \dots s_n$  in  $\mathcal{M}^\pi$  of states satisfies  $s_k \in B$ ,  $s_{k+1}, \dots, s_{n-1} \notin A$  and  $s_n \in A$ , we write  $s_k s_{k+1} \dots s_n \in B(S \setminus A)^* A$ .

**Definition 11:** For an induced MC  $\mathcal{M}^\pi$ , the *hitting probability* of a set  $A$  of states from a set  $B$

of states is  $\rho^\pi(B, A) := \sum_{s_k \dots s_n \in T} \prod_{k \leq i < n} \mathcal{P}_{s_i, s_{i+1}}^\pi$  where  $T = B(S \setminus A)^* A$ .

**Proof of Proposition 1** If  $\sup_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi) = \infty$ , equality follows from the fact that  $\Pi^S(\mathcal{M}) \subseteq \Pi(\mathcal{M})$ . If  $\sup_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi) < \infty$ , the result follows from Proposition 3 in this paper, Proposition 35 in [3], and Proposition 2 in [34]. Specifically,  $H(\mathcal{M}, \pi)$  can be written as an expected total cost with respect to a specific cost function on an MDP with a compact action set [3], Proposition 3, together with Theorem 1, implies that every stationary policy on this MDP is *proper*, i.e., all stochastic processes induced by stationary policies are guaranteed to reach an absorbing state within finite time step, and finally, the sufficiency of stationary policies to minimize the expected total cost on this MDP follows from [34].  $\square$

**Proof of Lemma 1:** We prove the sufficiency by contradiction and the necessity by construction.

( $\Rightarrow$ ) To obtain a contradiction, assume there exists a state  $s \in C$  that satisfies  $|\cup_{a \in D(s)} \text{Succ}(s, a)| = 1$  and is both stochastic and recurrent in an MC  $\mathcal{M}^\pi$  induced by a stationary policy  $\pi \in \Pi^S(\mathcal{M})$ . Since the state  $s$  is stochastic in  $\mathcal{M}^\pi$  and  $|\cup_{a \in D(s)} \text{Succ}(s, a)| = 1$ , the policy  $\pi$  satisfies  $\pi_s(a_2) > 0$  for some action  $a_2 \in D_0(s)$ . Therefore, there exists a state  $t$  in  $\mathcal{M}^\pi$  such that  $t \in \text{Succ}(s) \setminus C$  by the definition of  $D_0(s)$ .

Let  $u = \mathcal{P}_{s,t}^\pi > 0$ . Then, for all  $k \in \mathbb{N}$ ,  $(\mathcal{P}^\pi)_{s,s}^k \leq (1 - u')^k$  for some  $0 < u' \leq u$ . (Note that if  $u' = 0$  for some  $k \in \mathbb{N}$ , then there exists a path that starts from the state  $s$ , visits the state  $t$  and returns to the state  $s$  with probability 1. However, in this case we should have  $t \in C$ .) As a result,

$$\xi_s^\pi = \rho^\pi(s_0, s) \sum_{k=0}^{\infty} (\mathcal{P}^\pi)_{s,s}^k \leq \sum_{k=0}^{\infty} (1 - u')^k = \frac{1}{u'} < \infty, \quad (32)$$

where we use the fact that the hitting probability satisfies  $\rho^\pi(s_0, s) \leq 1$ . This raises a contradiction since the state  $s$  is recurrent, and it must satisfy  $\xi_s^\pi = \infty$ .

( $\Leftarrow$ ) Suppose there exists a state  $s \in C$  such that  $|\cup_{a \in D(s)} \text{Succ}(s, a)| > 1$ . Then, either (i) there exist actions  $a_i, a_j \in D(s)$  such that  $\text{Succ}(s, a_i) \setminus \text{Succ}(s, a_j) \neq \emptyset$ , or (ii) there exists an action  $a_i \in D(s)$  such that  $|\text{Succ}(s, a_i)| > 1$ . For case (i), we construct a policy  $\pi \in \Pi^S(\mathcal{M})$  such that  $\pi_s(a_i) > 0$  and  $\pi_s(a_j) > 0$  in the state  $s$ , and  $\pi_{s'}(a) = 1$  for some action  $a \in D(s')$  in states  $s' \in C \setminus \{s\}$ . Note that such actions exist by the definition of MEC. Finally, in states  $t \notin C$ , we choose actions such that the state  $s$  is reachable from the initial state. In the MC  $\mathcal{M}^\pi$  induced by  $\pi$ , the state  $s$  is an element of a bottom strongly connected component (BSCC) and  $|\text{Succ}(s)| > 1$ . Hence, it is both recurrent and stochastic. For the case (ii), we choose a policy  $\pi \in \Pi^S(\mathcal{M})$  such that  $\pi_s(a_i) = 1$  in the state  $s$  and  $\pi_{s'}(a) = 1$  for some  $a \in D(s')$  in states  $s' \in C \setminus \{s\}$ . In  $\mathcal{M}^\pi$ , the state  $s$  belongs to a BSCC and has multiple successor states. Hence, it is both recurrent and stochastic.  $\square$

**Proof of Theorem 1:** We first prove the necessary and sufficient conditions for an MDP to have infinite or unbounded maximum entropy. Then, we show that if the maximum entropy is not infinite and not unbounded, then it is finite and attainable by a stationary policy.

**Infinite maximum entropy.** We prove that the maximum entropy of an MDP  $\mathcal{M}$  is infinite if and only if there exists a state  $s \in C$  such that  $|\cup_{a \in D(s)} \text{Succ}(s, a)| > 1$ , and conclude, by Lemma 1, that the claim holds.

( $\Rightarrow$ ) The proof is by contradiction. Assume that the maximum entropy of  $\mathcal{M}$  is infinite, i.e.  $\max_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi) = \infty$ , and  $|\cup_{a \in D(s)} \text{Succ}(s, a)| = 1$  for all states  $s \in C$ . We consider two cases: (i)  $D_0(s) = \emptyset$  for all  $s \in C$ , and (ii)  $D_0(s) \neq \emptyset$  for some  $s \in C$ .

**Case (i):** Suppose that  $D_0(s) = \emptyset$  for all  $s \in C$ . Then, for an arbitrarily chosen MC  $\mathcal{M}^\pi$  induced by a policy  $\pi \in \Pi^S(\mathcal{M})$ , we have  $|\text{Succ}(s)| = 1$  for all  $s \in C$ . Hence, for  $\mathcal{M}^\pi$ ,  $L^\pi(s) = 0$  for all  $s \in C$  due to (15). Recall that  $H(\mathcal{M}, \pi) < \infty$  if and only if for all  $s \in S$ ,  $\xi_s^\pi = \infty$  implies  $L^\pi(s) = 0$  [2], and note that, if  $\xi_s^\pi = \infty$ , then  $s \in C$ . Consequently, we have  $H(\mathcal{M}, \pi) < \infty$  for any MC  $\mathcal{M}^{\pi'}$  induced by a policy  $\pi' \in \Pi^S(\mathcal{M})$  since we choose  $\mathcal{M}^\pi$  arbitrarily. This implies that  $\max_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi) < \infty$  and raises a contradiction since for an MDP with infinite maximum entropy, we have  $\max_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi) = \infty$ .

**Case (ii):** Suppose that  $D_0(s) \neq \emptyset$  for some  $s \in C$ . For an arbitrarily chosen induced MC  $\mathcal{M}^\pi$ , the local entropy of any state  $s \in S$  is bounded by  $L^\pi(s) \leq \log|S|$  [2]. We assume, without loss of generality, that for all states  $s \in C$ ,  $|\text{Succ}(s)| > 1$  for  $\mathcal{M}^\pi$ . (If  $|\text{Succ}(s)| = 1$ , the state  $s$  has no contribution to the entropy of  $\mathcal{M}^\pi$  due to (8).) Recalling (32), for any  $s \in C$ , there exists a constant  $u_s > 0$  such that  $\xi_s^\pi \leq \frac{1}{u_s}$ . Then, for any  $\mathcal{M}^\pi$ ,  $\xi_s^\pi \leq \frac{1}{u'}$  for all  $s \in C$ , where  $u' = \min\{u_s : \xi_s^\pi \leq \frac{1}{u_s}, s \in C\}$ .

We now consider the states  $s' \notin C$ . For any state  $s' \notin C$ ,  $\mathbb{P}_{s' a s'} < 1$  for all  $a \in \mathcal{A}(s')$  since otherwise the state  $s'$  must belong to a MEC. Then, there exists a constant  $u'' > 0$  such that  $\xi_{s'}^\pi \leq \frac{1}{u''}$  for all  $s' \notin C$  for any induced MC  $\mathcal{M}^\pi$ . As a result,

$$\begin{aligned} H(\mathcal{M}, \pi) &= \sum_{s \in S} L^\pi(s) \xi_s^\pi = \sum_{s \in C} L^\pi(s) \xi_s^\pi + \sum_{s \notin C} L^\pi(s) \xi_s^\pi, \\ &\leq \frac{|S| \log|S|}{u'} + \frac{|S| \log|S|}{u''} < \infty, \end{aligned}$$

for any induced MC  $\mathcal{M}^\pi$  and for some  $u', u'' > 0$ . This implies  $\max_{\pi \in \Pi^S(\mathcal{M})} H(\mathcal{M}, \pi) < \infty$  and raises a contradiction.

( $\Leftarrow$ ) Using the proof of Lemma 1, we can construct a policy for which the induced MC contains a state that is both stochastic and recurrent. By Corollary 1 in [2], the entropy of the induced MC is infinite.

**Unbounded maximum entropy.** ( $\Rightarrow$ ) The proof is by contradiction. Assume that the maximum entropy of  $\mathcal{M}$  is unbounded, and there exists  $s \in C$  such that  $|\cup_{a \in D(s)} \text{Succ}(s, a)| > 1$  or all MECs in  $\mathcal{M}$  are bottom strongly connected. First, suppose that  $H(\mathcal{M})$  is unbounded and there exists  $s \in C$  such that  $|\cup_{a \in D(s)} \text{Succ}(s, a)| > 1$ . Then, by case (i) of Theorem 1, the maximum entropy of  $\mathcal{M}$  is infinite, which is a contradiction. Second, suppose that  $H(\mathcal{M})$  is unbounded and all MECs in  $\mathcal{M}$  are bottom strongly connected. Then,  $H(\mathcal{M}, \pi) < \infty$  for all  $\pi \in \Pi^S(\mathcal{M})$  by the definition of unboundedness. Using case (i) of Theorem 1, we conclude that there is no state in MECs that is both stochastic and recurrent in an induced MC  $\mathcal{M}^\pi$ . Consequently, all states  $s \in C$  are deterministic for any induced MC since  $D_0(s) = \emptyset$

and  $|\cup_{a \in D(s)} \text{Succ}(s, a)| = 1$  for all  $s \in C$ . This implies that  $L^\pi(s) = 0$  for all  $s \in C$ . Since every state  $s' \notin C$  satisfies  $\mathbb{P}_{s' a, s'} < 1$  for all  $a \in \mathcal{A}(s')$ , there exists a constant  $u^* > 0$  such that for all  $s' \notin C$  and for all  $\pi \in \Pi^S(\mathcal{M})$ ,  $\xi_{s'}^\pi \leq \frac{1}{u^*}$ . As a result,

$$H(\mathcal{M}, \pi) = \sum_{s \in S} L^\pi(s) \xi_s^\pi = \sum_{s \notin C} L^\pi(s) \xi_s^\pi \leq \frac{|S| \log|S|}{u^*}$$

for any policy  $\pi \in \Pi^S(\mathcal{M})$ . Hence, the maximum entropy is bounded. Since we assumed at the beginning that the maximum entropy is unbounded, this raises a contradiction.

( $\Leftarrow$ ) The proof is by construction. Suppose that  $\mathcal{M}$  has a MEC  $C^*$  which is not bottom strongly connected. Then, there exists a state  $s \in C^*$  such that  $D_0(s) \neq \emptyset$ . Let  $R = \text{Succ}(s, a_i) \setminus C^*$  for  $a_i \in D_0(s)$ , i.e., the set of states that are reachable from the state  $s$  and do not belong to the MEC  $C^*$ . We construct a policy  $\pi \in \Pi^S(\mathcal{M})$  such that, for the state  $s$ ,  $\sum_{t \in R} \mathcal{P}_{s,t}^\pi = \epsilon$ , and for states  $s' \in C^* \setminus \{s\}$ ,  $\pi_{s'}(a) = 1$ , for some  $a \in D(s')$ . For states  $s'' \notin C^*$ , we choose actions so that state  $s$  is reachable from the initial state in the induced MC  $\mathcal{M}^\pi$ .

The induced MC  $\mathcal{M}^\pi$  has the property that  $\rho^\pi(t, s) = 1$  for all  $t \in C^* \setminus \{s\}$  and  $\rho^\pi(s, s) = 1 - \epsilon'$  for some  $0 < \epsilon' \leq \epsilon$ . Here, we note that state  $s$  is not recurrent in  $\mathcal{M}^\pi$ .

Since states  $s \in C^*$  are reachable from the initial state in  $\mathcal{M}^\pi$ , we have  $\rho^\pi(s_0, s) > 0$ . Additionally,  $\mathcal{K} := \{k | (\mathcal{P}^\pi)_{s_0, s}^k > 0\}$  is non-empty. Let  $k^* := \min(\mathcal{K})$  and  $\rho := (\mathcal{P}^\pi)_{s_0, s}^{k^*}$ . Then, we have  $\rho \leq \rho(s_0, s)$  because  $\rho$  only includes the first hitting probability. Moreover, the state  $s$  satisfies  $\rho^\pi(s, s) = 1 - \epsilon' < 1$ . Then,  $\xi_s^\pi = \frac{\rho^\pi(s_0, s)}{1 - \rho^\pi(s, s)} \geq \frac{\rho}{\epsilon'} \geq \frac{\rho}{\epsilon}$ , where the equality is a well-known result for finite-state MCs [35], [36].

The local entropy of the state  $s$  is the smallest when it has two outgoing transitions, one with probability  $\epsilon$  to a state  $t \in R$  and the other with probability  $1 - \epsilon$  to a state in  $C^*$  [2]. (It can be imagined as a Bernoulli random variable with parameter  $\epsilon$  where  $\epsilon$  can be arbitrarily small.) Hence,  $L^\pi(s) \geq -((\epsilon \log \epsilon) + ((1 - \epsilon) \log(1 - \epsilon)))$ . As a result,

$$L^\pi(s) \xi_s^\pi \geq -\frac{\rho((\epsilon \log \epsilon) + ((1 - \epsilon) \log(1 - \epsilon)))}{\epsilon}. \quad (33)$$

Note that  $\lim_{\epsilon \rightarrow 0^+} L^\pi(s) \xi_s^\pi = \infty$ . Therefore, for any policy  $\pi \in \Pi^S(\mathcal{M})$ , it is always possible to find another policy  $\pi' \in \Pi^S(\mathcal{M})$  that induces an MC with a greater entropy. Hence, the maximum entropy of the MDP is unbounded.

**Finite maximum entropy.** ( $\Leftarrow$ ) The result follows from the definition of the finite maximum entropy.

( $\Rightarrow$ ) Assume that the maximum entropy is not infinite and not unbounded. Hence,  $H(\mathcal{M}) = \sup_{\pi \in \Pi(\mathcal{M})} H(\mathcal{M}, \pi) < \infty$  which implies that, for any policy  $\pi \in \Pi^S(\mathcal{M})$ ,  $H(\mathcal{M}, \pi) < \infty$ . Then, for all states  $s \in C$ ,  $|\cup_{a \in D(s)} \text{Succ}(s, a)| = 1$  and all MECs are BSC by cases (i) and (ii) of Theorem 1, respectively. As a result, for all  $s \in C$ , we have  $L^\pi(s) = 0$  for any  $\pi \in \Pi^S(\mathcal{M})$ , and hence,  $H(\mathcal{M}, \pi) = \sum_{s \in S \setminus C} \xi_s^\pi L^\pi(s)$  for any  $\pi \in \Pi^S(\mathcal{M})$ .

Suppose that there exists a state  $s \in S \setminus C$  such that  $\xi_s^\pi = 0$  for some  $\pi \in \Pi^S(\mathcal{M})$ . Then, for the induced MC  $\mathcal{M}^\pi$ ,  $\xi_s^\pi L^\pi(s) = 0$  since  $L^\pi(s) \leq \log|S|$  is bounded. Therefore, without loss of generality, we can neglect unreachable states in any induced MC  $\mathcal{M}^\pi$  and assume  $\xi_s^\pi > 0$  for all states  $s \in S \setminus C$ . We pick an arbitrary state  $s \in S \setminus C$

and an arbitrary policy  $\pi \in \Pi^S(\mathcal{M})$ , and define a new function  $\lambda^\pi(s, a) = \sum_{k=0}^{\infty} (\mathcal{P}^\pi)_{s_0, s}^k \pi_s(a) = \xi_s^\pi \pi_s(a)$  which satisfies  $\lambda^\pi(s, a) \geq 0$  and  $\xi_s^\pi = \sum_{a \in \mathcal{A}(s)} \lambda^\pi(s, a)$ . Note that  $\xi_s^\pi < \infty$  since the state  $s \in S \setminus C$  is transient in  $\mathcal{M}^\pi$ . We also have  $\pi_s(a) = \frac{\lambda^\pi(s, a)}{\sum_{a \in \mathcal{A}(s)} \lambda^\pi(s, a)}$  since  $\xi_s^\pi = \sum_{a \in \mathcal{A}(s)} \lambda^\pi(s, a) > 0$  for reachable states. Plugging  $\xi_s^\pi$  and  $\pi_s(a)$  into (18), we obtain

$$H(\mathcal{M}) = \sup_{\substack{\sum_{a \in \mathcal{A}(s)} \lambda^\pi(s, a) > 0 \\ \lambda^\pi(s, a) \geq 0}} - \sum_{s \in S \setminus C} \sum_{t \in S} \left[ \sum_{a \in \mathcal{A}(s)} \lambda^\pi(s, a) \mathbb{P}_{s, a, t} \right] \log \frac{\left[ \sum_{a' \in \mathcal{A}(s)} \lambda^\pi(s, a') \mathbb{P}_{s, a', t} \right]}{\left[ \sum_{a'' \in \mathcal{A}(s)} \lambda^\pi(s, a'') \right]}. \quad (34)$$

Let  $M := \sup_{\pi \in \Pi(\mathcal{M})} \sum_{a \in \mathcal{A}(s)} \lambda^\pi(s, a) < \infty$ . Then, the function  $H(\mathcal{M})$  is continuous in  $\lambda^\pi(s, a)$  and bounded over the region  $\mathcal{R} := \{(\lambda^\pi(s, a))_{a \in \mathcal{A}(s)} \mid \lambda^\pi(s, a) \geq 0, \sum_{a \in \mathcal{A}(s)} \lambda^\pi(s, a) \leq M\}$  where, if  $\sum_{a \in \mathcal{A}(s)} \lambda^\pi(s, a) = 0$ , we use the convention  $0 \log \frac{0}{0} = 0$  which preserves continuity. Note that the set  $\mathcal{R}$  is closed. It is also compact since we have  $\max_{\pi \in \Pi(\mathcal{M})} \xi_s^\pi = \sup_{\pi \in \Pi(\mathcal{M})} \xi_s^\pi$  for all  $s \in S \setminus C$ , which can be shown by formulating a reward maximization problem and noting that the maximum expected reward is attainable by deterministic stationary policies. We omit the explicit construction of the reward maximization problem here for brevity and refer the reader to Chapter 2 in [26] for details. Finally, since we have a continuous function maximized over a compact set in the right hand side of (34), the supremum is achievable.  $\square$

**Proof of Proposition 3:** Since  $\mathcal{M}$  has a finite maximum entropy, all states  $s \in C$  have a single successor state, i.e.,  $|Succ(s, a)| = 1$ , due to Theorem 1. Additionally, all MECs are BSC due to Theorem 1. Hence, all states  $s \in C$  are either unreachable or recurrent, and have zero local entropy  $L^\pi(s) = 0$  in any MC  $\mathcal{M}^\pi$  induced by a policy  $\pi \in \Pi^S(\mathcal{M})$ . Recall that for MCs with finite total entropy, the sum in (16) is taken only over the transient states. Therefore, changing the successors of the states  $s \in C$  does not affect the maximum entropy of  $\mathcal{M}$  as long as the conditions  $|Succ(s, a)| = 1$  and  $Succ(s, a) \subseteq C$  are not violated. Making states in MECs absorbing does not violate these conditions, and hence, the result follows.  $\square$

**Proof of Proposition 4:** All constraints are affine in variables  $\lambda(s, a)$  and  $\lambda(s)$ . Hence, we need only to show that the objective function is concave over the domain  $\lambda(s, a) \geq 0$ . For a given state  $s \in S \setminus C$ , define the vectors  $\mathbf{\Gamma}_s = (\eta(s, t))_{t \in S}$  and  $\mathbf{N}_s = \nu(s) \mathbf{1}$  where  $\mathbf{1} \in \mathbb{R}^n$ . Recalling that  $\eta(s, t)$  and  $\nu(s)$  are functions of  $\lambda(s, a)$ , we define the function  $f(\mathbf{\Gamma}_s, \mathbf{N}_s) := \sum_{t \in S} \eta(s, t) \log \left( \frac{\eta(s, t)}{\nu(s)} \right)$  over the domain  $\nu(s) = \sum_{a \in \mathcal{A}(s)} \lambda(s, a) \geq 0$  and use the convention (based on continuity arguments) that  $f(\mathbf{\Gamma}_s, \mathbf{0}) = 0$ .

The function  $f(\mathbf{\Gamma}_s, \mathbf{N}_s)$  is the relative entropy between the vectors  $\mathbf{\Gamma}_s$  and  $\mathbf{N}_s$ , and thus, it is convex over the domain  $\sum_{a \in \mathcal{A}(s)} \lambda(s, a) > 0$  [37]. Moreover, since  $\nu(s) \geq \eta(s, t) \geq 0$  for all  $s \in S \setminus C$  and  $t \in S$ ,  $f(\mathbf{\Gamma}_s, \mathbf{N}_s) \leq 0$  for all  $s \in S \setminus C$ . Therefore, for states  $s \in S \setminus C$ , we can include the point  $\sum_{a \in \mathcal{A}(s)} \lambda(s, a) = 0$  to the domain over which the function  $f$  is convex. Now, note that the objective function in (21a)

is equal to  $\sum_{s \in S \setminus C} -f(\mathbf{\Gamma}_s, \mathbf{N}_s)$ . Since the sum of convex functions is convex and the negation of a convex function is concave [37], the objective function (21a) is concave over the domain  $\lambda(s, a) \geq 0$ .  $\square$

**Proof of Theorem 2:** Assuming  $H(\mathcal{M}) < \infty$ , we have  $H(\mathcal{M}') = H(\mathcal{M})$  due to Proposition 3, and hence, an optimal policy for  $\mathcal{M}'$  is also optimal for  $\mathcal{M}$ . We first prove that for a given modified MDP  $\mathcal{M}'$ , the objective function (21a) of the convex program in (21a)-(21g) is the maximum entropy  $H(\mathcal{M}')$  of  $\mathcal{M}'$ . Then, we construct an optimal policy for  $\mathcal{M}$  using the optimal variables  $\lambda^*(s, a)$  that solve the program in (21a)-(21g) for  $(\mathcal{M}', C)$ , where  $C$  is the set of all states in MECs in  $\mathcal{M}'$ .

We utilize the results of [22] to relate the variables  $\lambda(s, a)$  with the expected residence time in states. In [22], it is shown that variables  $\lambda(s, a) = \sum_{k=0}^{\infty} (\mathcal{P}^\pi)_{s_0, s}^k \pi_s(a) = \xi_s^\pi \pi_s(a)$  satisfy the constraint in (21b) and corresponds to the expected residence time in a state-action pair  $(s, a)$  in an induced MC  $\mathcal{M}^\pi$ . Additionally,  $\lambda(s)$  corresponds to the reachability probability of states  $s \in C$ . Then, it is clear that for states  $s \in S \setminus C$ ,

$$\xi_s^\pi = \sum_{a \in \mathcal{A}(s)} \lambda(s, a). \quad (35)$$

Additionally, if  $\sum_{a \in \mathcal{A}(s)} \lambda(s, a) > 0$ , we have

$$\pi_s(a) = \frac{\lambda(s, a)}{\sum_{a \in \mathcal{A}(s)} \lambda(s, a)}. \quad (36)$$

Recall that for all  $s \in C$  and  $\pi \in \Pi^S(\mathcal{M}')$ , we have  $L^\pi(s) = 0$  since  $H(\mathcal{M}') < \infty$ . Therefore,

$$H(\mathcal{M}') = \sup_{\pi \in \Pi^S(\mathcal{M}')} \left[ \sum_{s \in S} \xi_s^\pi L^\pi(s) \right] \quad (37)$$

$$= \max_{\pi \in \Pi^S(\mathcal{M}')} \left[ \sum_{s \in S \setminus C} \xi_s^\pi L^\pi(s) \right], \quad (38)$$

due to Proposition 2 and Theorem 1. Our aim is to show that the expression in (38) is equal to the objective in (21a).

For an arbitrary  $\pi \in \Pi^S(\mathcal{M}')$ , define the set  $G^\pi := \{s \in S \setminus C \mid \xi_s^\pi = 0\}$ . Note that if  $G^\pi \neq \emptyset$  for some  $\pi \in \Pi^S(\mathcal{M}')$ , the states  $s \in G^\pi$  do not affect the value of (38) by the definition of  $G^\pi$ . We consider two cases: (1)  $G^\pi \neq \emptyset$  and (2)  $G^\pi = \emptyset$ . For case 1, we will show that states  $s \in G^\pi$  do not affect the value of (21a), and for case 2, we will show that the expression in (38) is equal to the objective in (21a).

**Case 1:** We assume that  $G^\pi \neq \emptyset$  and show that for any  $s \in G^\pi$ ,

$$\sum_{t \in S} \eta(s, t) \log \left( \frac{\eta(s, t)}{\nu(s)} \right) = 0. \quad (39)$$

Considering (21d)-(21f), and noting that  $0 \leq \mathbb{P}_{s, a, t} \leq 1$  for all  $t \in S$ , we have  $\nu(s) \geq \eta(s, t) \geq 0$  for all  $t \in S$ . Hence, for any  $s \in G^\pi$ , we have  $\nu(s) = \eta(s, t) = 0$  for all  $t \in S$  due to the definition of the set  $G^\pi$  and (35). We use the convention  $0 \log \frac{0}{0} = 0$  which is based on continuity arguments and the fact that whenever  $\nu(s) = 0$ , we have  $\eta(s, t) = 0$  for all  $t \in S$ . As a result, we conclude that the states  $s \in G^\pi$  do not affect the value of the objective in (21a).

**Case 2:** We assume that  $G^\pi = \emptyset$ . In this case, for any  $\pi \in \Pi^S(\mathcal{M}')$ , we have  $\xi_s^\pi = \sum_{a \in \mathcal{A}(s)} \lambda(s, a) > 0$  and, (36) holds for all  $s \in S \setminus C$  and  $a \in \mathcal{A}(s)$ . Plugging (35) and (36) into (38), we obtain the objective function in (21a). (Note that  $\eta(s, t)$  and  $\nu(s)$  variables can be written in terms of  $\lambda(s, a)$  using (21d)-(21e).) We conclude that the problem in (21a)-(21g) computes the maximum entropy of  $\mathcal{M}'$ .

Now, we construct an optimal policy for  $\mathcal{M}$ . We show in (38) that states  $s \in C$  does not affect the value of  $H(\mathcal{M}')$ . Therefore, an arbitrary assignment of actions in states  $s \in C$  does not affect the optimality of a policy. Similarly, for a given optimal policy  $\pi^*$ , an arbitrary assignment of actions in states  $s \in G^{\pi^*}$  does not affect the optimality due to (39). Additionally, by the construction given in (36), an optimal policy for states  $s \in S \setminus (C \cup G^{\pi^*})$  satisfies  $\pi_s^*(a) = \frac{\lambda^*(s, a)}{\sum_{a \in \mathcal{A}(s)} \lambda^*(s, a)}$ , where  $\lambda^*(s, a)$  are optimal variables for the problem in (21a)-(21g). Since an optimal policy for  $\mathcal{M}'$  is also optimal for  $\mathcal{M}$  due to Proposition 3, we conclude that Algorithm 2 returns an optimal policy for  $\mathcal{M}$ .  $\square$

**Proof of Lemma 2:** For an MC  $\mathcal{M}^\pi$  induced by a policy  $\pi \in \Pi^S(\mathcal{M})$ , let  $S_B$  and  $S_{B_0}$  be the union of its BSCCs and the set of its transient states, respectively. Moreover, let  $T := \text{Paths}_{\text{fin}}^\pi(\mathcal{M}) \cap (S \setminus S_B)^* S_B$ . For states  $s \in S_{B_0}$  and  $t \in S$ , define sets

$$A_{s,k} := \{s_0 \dots s_n \in T : n \in \mathbb{N}, \sum_{i=0}^n \mathbb{1}_{\{s_i=s\}} = k\},$$

$$B_{(s,t),k} := \{s_0 \dots s_n \in T : n \in \mathbb{N}, \sum_{i=1}^n \mathbb{1}_{\{(s_{i-1}, s_i)=(s,t)\}} = k\}.$$

Note that  $A_{s,k}$  is the collection of all paths along which the state  $s$  is  $k$  times visited and a BSCC in  $\mathcal{M}^\pi$  is reached. Similarly, the set  $B_{(s,t),k}$  is the collection of all paths along which the edge between state  $s$  and state  $t$  is  $k$  times traversed, and a BSCC in  $\mathcal{M}^\pi$  is reached.

It is known that any finite MC almost surely reaches a BSCC [23]. Thus, for any  $s \in S_{B_0}$  and  $t \in S$ , we have

$$\sum_{k=0}^{\infty} \Pr_{\mathcal{M}}^\pi(A_{s,k}) = \sum_{k=0}^{\infty} \Pr_{\mathcal{M}}^\pi(B_{(s,t),k}) = 1. \quad (40)$$

One can show using (40) that

$$\xi_s^\pi = \sum_{k=0}^{\infty} k \Pr_{\mathcal{M}}^\pi(A_{s,k}) \quad (41)$$

for transient states  $s \in S_{B_0}$  in  $\mathcal{M}^\pi$ . (We omit the derivation here. The result can be obtained by using the countability of  $A_{s,k}$  and performing a series of algebraic manipulations to obtain (1). A similar derivation can also be found in [14].)

Let  $\xi_{s,t}^\pi$  denote the expected number of transitions from a state  $s \in S_{B_0}$  to state  $t \in S$ . Then, we have

$$\xi_{s,t}^\pi = \sum_{k=0}^{\infty} k \Pr_{\mathcal{M}}^\pi(B_{(s,t),k}), \quad (42)$$

analogously to (41). Additionally, the relation between (41) and (42) is given by  $\mathcal{P}_{s,t}^\pi \xi_s^\pi = \xi_{s,t}^\pi$ , which can be obtained by using a method similar to the one used in [36] to derive (1).

Let  $N_{s,t}^{s_0 \dots s_n}$  be the number of transitions made from state  $s \in S_{B_0}$  to state  $t \in S$  along a finite path fragment  $s_0 \dots s_n \in T$ . Then, we have

$$\mathcal{P}_{s,t}^\pi \xi_s^\pi = \xi_{s,t}^\pi = \sum_{k=0}^{\infty} \sum_{s_0 \dots s_n \in B_{(s,t),k}} k \mathcal{P}^\pi(s_0 \dots s_n) \quad (43)$$

$$= \sum_{s_0 \dots s_n \in T} N_{s,t}^{s_0 \dots s_n} \mathcal{P}^\pi(s_0 \dots s_n), \quad (44)$$

where the equality in (43) follows from the fact that set  $B_{(s,t),k}$  is countable and each element  $s_0 \dots s_n \in B_{(s,t),k}$  is measurable. The equality in (44) is due to the fact that any finite path fragment  $s_0 \dots s_n \in T$  is an element of one and only one set  $B_{(s,t),k}$ , and that for a given path fragment  $s_0 \dots s_n$ , we have  $k = N_{s,t}^{s_0 \dots s_n}$  by definition.

We next express the probability of a finite path fragment in terms of the number of transition appearances. Then, we have

$$\mathcal{P}^\pi(s_0 \dots s_m) = \prod_{(s,t) \in S_{B_0} \times S : N_{s,t}^{s_0 \dots s_m} > 0} (\mathcal{P}_{s,t}^\pi)^{N_{s,t}^{s_0 \dots s_m}}. \quad (45)$$

By assumption, we have  $H(\mathcal{M}, \pi) < \infty$ . If  $s_0 \in S_B$ , both the entropy and the entropy of paths for  $\mathcal{M}^\pi$  are zero; hence, we only analyze the case  $s_0 \in S_{B_0}$ . In this case, the summation in (16) is taken over transient states  $s \in S_{B_0}$  since  $H(\mathcal{M}, \pi) < \infty$ . As a result,

$$H(\mathcal{M}, \pi) = - \sum_{s \in S_{B_0}} \xi_s^\pi \sum_{t \in S} \mathcal{P}_{s,t}^\pi \log \mathcal{P}_{s,t}^\pi \quad (46)$$

$$= - \sum_{(s,t) \in S_{B_0} \times S : \mathcal{P}_{s,t}^\pi > 0} \xi_s^\pi \mathcal{P}_{s,t}^\pi \log \mathcal{P}_{s,t}^\pi \quad (47)$$

$$= - \sum_{(s,t) \in S_{B_0} \times S : \mathcal{P}_{s,t}^\pi > 0} \sum_{s_0 \dots s_n \in T} N_{s,t}^{s_0 \dots s_n} \mathcal{P}^\pi(s_0 \dots s_n) \log \mathcal{P}_{s,t}^\pi \quad (48)$$

$$= - \sum_{(s,t) \in S_{B_0} \times S : N_{s,t}^{s_0 \dots s_n} > 0} \sum_{s_0 \dots s_n \in T} N_{s,t}^{s_0 \dots s_n} \mathcal{P}^\pi(s_0 \dots s_n) \log \mathcal{P}_{s,t}^\pi, \quad (49)$$

where (47) follows by removing transitions  $\mathcal{P}_{s,t}^\pi = 0$  and using the convention  $0 \log 0 = 0$ , (48) follows from (44), and (49) is obtained by removing state pairs  $(s, t) \in S_{B_0} \times S$  for which  $N_{s,t}^{s_0 \dots s_n} = 0$ .

Now, we analyze the entropy of paths. The entropy of paths  $H(\text{Paths}^\pi(\mathcal{M}))$  for the induced MC  $\mathcal{M}^\pi$  can be written as

$$H(\text{Paths}^\pi(\mathcal{M})) = - \sum_{s_0 \dots s_n \in T} \mathcal{P}^\pi(s_0 \dots s_n) \log \mathcal{P}^\pi(s_0 \dots s_n), \quad (50)$$

$$= - \sum_{s_0 \dots s_m \in T} \sum_{(s,t) \in S_{B_0} \times S : N_{s,t}^{s_0 \dots s_n} > 0} N_{s,t}^{s_0 \dots s_n} \mathcal{P}^\pi(s_0 \dots s_n) \log \mathcal{P}_{s,t}^\pi, \quad (51)$$

where (51) is obtained by plugging (45) into (50).

Since (49) and (51) are equal, we conclude that

$H(\mathcal{M}, \pi) = H(\text{Paths}^\pi(\mathcal{M}))$  for any  $\pi \in \Pi^S(\mathcal{M})$ , under the assumption that  $H(\mathcal{M}, \pi) < \infty$  for all  $\pi \in \Pi^S(\mathcal{M})$ .  $\square$

## APPENDIX B

In this appendix, we provide procedures to solve entropy maximization and constrained entropy maximization problems for MDPs with infinite maximum entropy.

1) *Entropy maximization*: In this case, for a given MDP  $\mathcal{M}$  with the union  $(C, D)$  of its MECs, there exists at least one state  $s^* \in C$  such that  $|\cup_{a \in D(s^*)} \text{Succ}(s^*, a)| > 1$  due to Theorem 1. We aim to synthesize a policy that induces an MC where the state  $s^*$  is both stochastic and recurrent. For simplicity, we assume that there exists only one state  $s^*$  such that  $|\cup_{a \in D(s^*)} \text{Succ}(s^*, a)| > 1$ . We form the modified MDP  $\mathcal{M}'$  by replacing each BSC MEC in  $\mathcal{M}$  with an absorbing state. Let  $S_B$  and  $S_{NB}$  be the set of all states in BSC MECs and non-BSC MECs in  $\mathcal{M}$ , respectively. We consider two cases, namely  $s^* \in S_B$  and  $s^* \in S_{NB}$ . If  $s^* \in S_B$ , let  $C'$  be the union of all absorbing states in  $\mathcal{M}'$  that are replaced with BSC MECs in  $\mathcal{M}$ , and  $C_*$  be the absorbing state that is replaced with the MEC that  $s^*$  is contained in. We solve the problem in (21a)-(21g) for  $(\mathcal{M}', C')$  together with the constraint  $\lambda(C_*) > 0$ . (Note that if there is a non-BSC MEC in  $\mathcal{M}'$ , the constraint (23) should also be included to this optimization problem.) We then use step 3 of Algorithm 2 to obtain a policy for states  $s \notin C'$ , and choose actions in state  $s^*$  such that  $|\text{Succ}(s^*)| > 1$  in the induced MC. By construction, the state  $s^*$  is both stochastic and recurrent in the induced MC, and due to Proposition 2, the entropy of the induced MC is infinite. If  $s^* \in S_{NB}$ , we replace the MEC that state  $s^*$  is contained in with an absorbing state and follow steps similar to the ones in the case  $s^* \in S_B$  to obtain an optimal policy.

2) *Constrained entropy maximization*: We suppose that the feasible policy space for the problem in (30a)-(30b) is not empty. The product MDP  $\mathcal{M}_p$  contains a MEC  $(C, D)$  such that  $|\cup_{a \in D(s^*)} \text{Succ}(s^*, a)| > 1$  for some  $s^* \in C$  due to Theorem 1. We assume that there exists at least one non-BSC MEC in  $\mathcal{M}_p$  and there is only one state  $s^*$  in  $\mathcal{M}_p$  such that  $|\cup_{a \in D(s^*)} \text{Succ}(s^*, a)| > 1$ . These assumptions are introduced just to simplify the case analysis. We first partition the states into three disjoint sets  $B, S_0$  and  $S_r$  as explained in Section VI-C. Then, we form the modified MDP  $\mathcal{M}'_p$  by replacing each BSC MEC in  $\mathcal{M}_p$  with an absorbing state. Let  $S_B$  and  $S_{NB}$  be the set of all states in BSC MECs and non-BSC MECs in  $\mathcal{M}_p$ , respectively. We consider two cases: 1)  $s^* \in S_B$  and 2)  $s^* \in S_{NB}$ .

*Case 1*: If  $s^* \in S_B$ , let  $C'$  be the union of all absorbing states in  $\mathcal{M}'_p$  that are replaced with BSC MECs in  $\mathcal{M}_p$ , and  $C_*$  be the absorbing state that is replaced with the MEC that  $s^*$  is contained in. We obtain a policy for states  $s \in S_r$  by solving the problem in (21a)-(21g) for  $(\mathcal{M}'_p, C', \beta, \Gamma)$  together with the constraints (23), (31) and  $\lambda(C_*) > 0$ . Then, we choose actions in state  $s^*$  such that  $|\text{Succ}(s^*)| > 1$  in the induced MC. Note that if this problem is infeasible, then there exists no policy that induces an MC with infinite entropy whose paths satisfies the LTL specification with probability  $\beta$ . In this case, the maximum constrained entropy is unbounded, and we follow

the steps that are explained in Section VI-C2 to synthesize a policy that induces an MC with arbitrarily large entropy.

*Case 2*: If  $s^* \in S_{NB}$ , we consider two cases, namely  $s^* \in B \cup S_0$  and  $s^* \in S_r$ . If  $s^* \in B \cup S_0$ , we replace the MEC that  $s^*$  belongs to in  $\mathcal{M}'_p$  with an absorbing state. Then, we synthesize a policy that induces an MC with infinite entropy whose paths satisfy the LTL specification with probability  $\beta$  as explained in Case 1. Additionally, to ensure that the state  $s^*$  is recurrent in the induced MC, we choose actions in states that belong to the same MEC with  $s^*$  such that the MEC forms a BSCC in the induced MC. If  $s^* \in S_r$ , the maximum constrained entropy is not infinite because no state  $s \in S_r$  can be recurrent in an induced MC that satisfies the LTL specification with probability  $\beta > 0$ . In this case, the maximum constrained entropy is unbounded, and we use the procedure explained in Section VI-C2 to synthesize a policy that induces an MC with arbitrarily large entropy.



**Yagiz Savas** joined the Department of Aerospace Engineering at the University of Texas at Austin as a Ph.D. student in Fall 2017. He received his B.S. degree in Mechanical Engineering from Bogazici University in 2017. His research focuses on developing theory and algorithms that guarantee desirable behavior of autonomous systems operating in adversarial environments.



**Melkior Ornik** is an assistant professor in the Department of Aerospace Engineering and the Co-ordinated Science Laboratory at the University of Illinois at Urbana-Champaign. He received his Ph.D. degree from the University of Toronto in 2017. His research focuses on developing theory and algorithms for learning and planning of autonomous systems operating in uncertain, complex and changing environments, as well as in scenarios where only limited knowledge of the system is available.



**Murat Cubuktepe** joined the Department of Aerospace Engineering at the University of Texas at Austin as a Ph.D. student in Fall 2015. He received his B.S degree in Mechanical Engineering from Bogazici University in 2015. His main current research interests are verification and synthesis of uncertain, parametric and partially observable probabilistic systems. He also focuses on applications of convex optimization in formal methods and controls.



**Mustafa O. Karabag** joined the Department of Electrical and Computer Engineering at the University of Texas at Austin as a Ph.D. student in Fall 2017. He received his B.S. degree in Electrical and Electronics Engineering from Bogazici University in 2017. His research focuses on developing theory and algorithms for non-inferable planning in adversarial environments.



**Ufuk Topcu** joined the Department of Aerospace Engineering at the University of Texas at Austin as an assistant professor in Fall 2015. He received his Ph.D. degree from the University of California at Berkeley in 2008. He held research positions at the University of Pennsylvania and California Institute of Technology. His research focuses on the theoretical, algorithmic and computational aspects of design and verification of autonomous systems through novel connections between formal methods, learning theory and controls.