2020 59th IEEE Conference on Decision and Control (CDC)
Jeju Island, Republic of Korea, December 14-18, 2020

Privacy-Preserving Resilience of Cyber-Physical Systems to Adversaries

Bhaskar Ramasubramanian[1], Luyao Niu[2], Andrew Clark[2], Linda Bushnell[1], and Radha Poovendran[1]

***Abstract*— A cyber-physical system (CPS) is expected to be resilient to more than one type of adversary. In this paper, we consider a CPS that has to satisfy a linear temporal logic (LTL) objective in the presence of two kinds of adversaries. The first adversary has the ability to tamper with inputs to the CPS to influence satisfaction of the LTL objective. The interaction of the CPS with this adversary is modeled as a stochastic game. We synthesize a controller for the CPS to maximize the probability of satisfying the LTL objective under any policy of this adversary. The second adversary is an eavesdropper who can observe labeled trajectories of the CPS generated from the previous step. It could then use this information to launch other kinds of attacks. A labeled trajectory is a sequence of labels, where a label is associated to a state and is linked to the satisfaction of the LTL objective at that state. We use differential privacy to quantify the indistinguishability between states that are related to each other when the eavesdropper sees a labeled trajectory. Two trajectories of equal length will be differentially private if they are differentially private at each state along the respective trajectories. We use a skewed Kantorovich metric to compute distances between probability distributions over states resulting from actions chosen according to policies from related states in order to quantify differential privacy. Moreover, we do this in a manner that does not affect the satisfaction probability of the LTL objective. We validate our approach on a simulation of a UAV that has to satisfy an LTL objective in an adversarial environment.**

I. Introduction

Cyber-physical systems (CPSs) consist of tightly coupled cyber and physical components that work in conjunction with algorithms and communication channels to satisfy complex objectives in dynamic environments [1]. The objectives that a CPS will need to meet often vary with time. Temporal logic frameworks like linear temporal logic [2] enable the specification of goals such as safety, stability, and reachability. In applications like power systems [3] and automobiles [4], the CPS must achieve its goals in scenarios that can be manipulated by an adversary that can disrupt nominal operation [5]. Disruptions to CPSs like power and water networks can inconvenience large sections of the population.

An adversary may not have the ability to directly launch attacks on the CPS. However, it could be capable of collecting information about the system, which can then be used to affect nominal operation [6]. Therefore, a well-designed system must protect information critical to maintaining normal operation. Differential privacy [7] is a security property that makes it difficult for an adversary to discern information about a system by providing probabilistic guarantees on the indistinguishabilty of observations of the system. This technique was originally used to protect sensitive data of individuals in databases in a manner that allowed for statistical analysis on aggregated information obtained from the data [8]. It has since been used to protect information represented as trajectories of a dynamical system [9]. Differential privacy of trajectories was accomplished by adding a carefully calibrated noise to sensitive trajectories so that an adversary could not glean information about the modified trajectory. Recent work has studied the enforcement of differential privacy on multi-agent systems to meet a temporal logic objective [10], or to minimize a quadratic cost [11].

We study differential privacy for systems represented as discrete state, discrete-action Markov decision processes (MDPs). Such systems encompass the operating environments of a broad range of practical systems, including robots and UAVs [12], and are especially suited for experimental analysis. In order to reason about trajectories that satisfy an LTL formula, we work with $PCTL^*$, a probabilistic computational tree logic which has LTL formulas as its path formulas [2]. Two states will be differentially private if the probabilities of satisfying a desired LTL objective starting from these states are sufficiently close. A trajectory will be differentially private if all states in the trajectory that are generated according to some policy are differentially private. In previous work that studied differential privacy for Markov chains [13], [14], the treatment was restricted to ensuring differential privacy of the initial state of the Markov chain. In comparison, we study differential privacy of states in an MDP along trajectories corresponding to satisfaction of an LTL objective in the presence of an adversary.

Different from the aforementioned works, we consider a setting with two kinds of adversaries that have different capabilities, and act independently of each other. We are interested in ensuring differential privacy of trajectories against one adversary (E). At the same time, we want to maximize the probability of satisfaction of the LTL objective under actions of the other adversary (A). Specifically, we assume that adversary A has the ability to inject signals to affect the control inputs to the CPS thereby influencing the transitions between CPS states. We would like to maximize the probability of satisfying the LTL goal under any sequence of actions played by this adversary. The interaction between the CPS and A is modeled as a zero-sum game. The synthesis of a CPS policy that maximizes the probability of satisfying the LTL goal under any adversary policy is

[1]Network Security Lab, Department of Electrical and Computer Engineering, University of Washington, Seattle, WA 98195, USA. {bhaskarr, lb2, rp3}@uw.edu

[2]Department of Electrical and Computer Engineering, Worcester Polytechnic Institute, Worcester, MA 01609 USA. {lniu,aclark}@wpi.edu

This work was supported by the U.S. Army Research Office, National Science Foundation, and the Office of Naval Research via Grants W911NF-16-1-0485, CNS-1941670, and N00014-17-S-B001 respectively.

related to reaching a Stackelberg equilibrium of this game. Once this objective is accomplished, we want that trajectories produced by the synthesized controller be indistinguishable to the adversary E. E can eavesdrop on these trajectories by observing a sequence of labels associated to states along the trajectories. That is, observing labeled trajectories should not allow E to gain information about the state of the system. To the best of our knowledge, this is the first work that studies resilient control in the presence of an adversary along with ensuring differential privacy of states along a trajectory satisfying the LTL objective. Prior work has studied these problems separately for a single kind of adversary (either an adversary of type A or of type E).

A. Contributions

We solve the problem of satisfying an LTL objective in the presence of an adversary that can tamper with actuator inputs, while also ensuring that trajectories of the system which satisfy the objective are differentially private to an eavesdropper who can observe labels of states on these trajectories. We make the following contributions:

- We show that the CPS policy that will maximize the probability of satisfaction of the LTL objective under any adversary policy is related to reaching certain subsets of a Markov chain associated to representations of the environment and the LTL goal under these policies.
- We use a fragment of the logic $PCTL^*$, whose semantics are defined over MDPs, to reason about LTL trajectories. Starting with a realization of the defender policy that satisfies the LTL objective, we define a symmetric relation on states and show that pairs of states in this relation will be differentially private.
- We present a value-iteration procedure and show that if a (skewed) distance between values of 'related' states is below a threshold, then these states will be differentially private. Two trajectories will be differentially private if the initial states of the trajectories are related, and if optimal defender policies from these states and all pairs of subsequently related states are 'close' to each other.
- A case-study on the satisfaction of a reach-avoid specification for a UAV influenced by adversarial inputs in an environment with an eavesdropper illustrates our results.

B. Outline of Paper

The rest of this paper is organized as follows: Section II summarizes related work, and we give a brief introduction to LTL, stochastic games, and differential privacy in Section III. We state our problem and detail the steps of our solution in Sections IV and V. Section VI shows an example illustrating our approach, and Section VII concludes the paper.

II. Related Work

The satisfaction of an LTL objective for two-player stochastic games when the players had competing objectives was presented in [15], [16] for the case when states were fully observable and in [17], [18] for the partially observable case. The authors of [19], [20] extended this to the case when an adversary could tamper with both actuators and clocks of the CPS to affect the satisfaction of a time-sensitive temporal constraint. There, the same adversary was capable of effecting both kinds of attacks on the system, which makes it different from the assumptions made in this paper. Evaluating multiple traces corresponding to different executions of a system using hyperproperties [21] was proposed in [22]. The authors proposed a temporal logic $HyperPCTL$, in order to reason about probabilistic hyperproperties, and this approach was studied in the context of verification for CPSs in [23].

The authors of [24], [25] presented a logical characterization of differential privacy for labeled probabilistic transition systems using a trace metric that corresponded to *exact* differential privacy ($\delta = 0$ in Definition 7). The authors of [13], [14] computed distances that constituted sound upper bounds on a skewed total variation distance between initial states of a Markov chain. They used the Kantorovich metric [26] to lift a distance between states to a distance between probability distributions associated to these states. A large part of this analysis was motivated by the work on defining metrics for MDPs [27]–[29]. These metrics measured the distance between states in MDPs with large state spaces and then allowed aggregation of states that were 'close' to each other. They further showed that for such states, the values of the states were also 'close' to each other. Algorithms based on Monte-Carlo methods were used to give guarantees on the differential privacy of synthesized policies in [30].

Using differential privacy to protect information about trajectories of dynamical systems was first presented in [9]. Since then, it has been studied for privacy-preserving consensus in multi-agent dynamical systems [31], networked systems [32], and linear quadratic control for multi-agent systems [11]. We point the reader to the review in [33], and references therein for a detailed survey of recent developments using differential privacy in dynamical systems.

Differential privacy in the context of the satisfaction of temporal objectives is a relatively recent area of research. In [34], the authors introduced $dpCTL^*$, an extension of the logic $PCTL^*$ augmented with a differentially private operator. They presented a model-checking procedure to verify differential privacy on Markov chains. The authors assumed that actions are controlled by an adversary, which is different from the setup in this paper. The authors of [10] proposed a differentially private controller synthesis procedure for multi-agent systems to satisfy a metric temporal logic objective. In their setting, each agent added a noise term while communicating its location to a local hub, which in turn transmitted the information to a cloud controller that determined the optimal inputs for each agent.

III. Preliminaries

A. Linear Temporal Logic

Temporal logic frameworks enable the representation and reasoning about temporal information on propositional statements. *Linear temporal logic (LTL)* is one such framework, where the progress of time is 'linear'. An *LTL formula* [2] is defined over a set of atomic propositions $\mathcal{AP}$, and can be

written as $\varphi := \mathtt{T}|\sigma|\neg\varphi|\varphi\wedge\varphi|\mathbf{X}\varphi|\varphi\mathbf{U}\varphi$, where $\sigma \in \mathcal{AP}$, and $\mathbf{X}$ and $\mathbf{U}$ are temporal operators denoting the *next* and *until* operations respectively. The semantics of LTL are defined over (infinite) words in $2^{\mathcal{AP}}$. We write $\eta_0\eta_1\cdots := \eta \models \varphi$ when a trace $\eta \in (2^{\mathcal{AP}})^\omega$ satisfies an LTL formula φ.

Definition 1 (LTL Semantics)**.** *Let $\eta^i = \eta_i\eta_{i+1}\ldots$. Then, the semantics of LTL can be recursively defined as:*

1) $\eta \models \mathtt{T}$ *if and only if (iff)* η_0 *is true;*
2) $\eta \models \sigma$ *iff* $\sigma \in \eta_0$*;*
3) $\eta \models \neg\varphi$ *iff* $\eta \not\models \varphi$*;*
4) $\eta \models \varphi_1 \wedge \varphi_2$ *iff* $\eta \models \varphi_1$ *and* $\eta \models \varphi_2$*;*
5) $\eta \models \mathbf{X}\varphi$ *iff* $\eta^1 \models \varphi$*;*
6) $\eta \models \varphi_1\mathbf{U}\varphi_2$ *iff* $\exists j \geq 0$ *such that* $\eta^j \models \varphi_2$ *and for all* $k < j, \eta^k \models \varphi_1$.

Moreover, the logic admits derived formulas of the form: *i)* $\varphi_1 \vee \varphi_2 := \neg(\neg\varphi_1 \wedge \neg\varphi_2)$; *ii)* $\varphi_1 \Rightarrow \varphi_2 := \neg\varphi_1 \vee \varphi_2$; *iii)* $\mathbf{F}\varphi := \mathtt{T}\mathbf{U}\varphi$ (eventually); *iv)* $\mathbf{G}\varphi := \neg\mathbf{F}\neg\varphi$ (always).

Definition 2 (Deterministic Rabin Automaton)**.** *A* deterministic Rabin automaton (DRA) *is a quintuple $\mathcal{RA} = (Q, \Sigma, \kappa, q_0, F)$ where Q is a nonempty finite set of states, Σ is a finite alphabet, $\kappa : Q \times \Sigma \to Q$ is a transition function, $q_0 \in Q$ is the initial state, and $F := \{(L(i), K(i)\}_{i=1}^M$ where $L(i), K(i) \subseteq Q$ for all i, and M is a positive integer.*

A *run* of $\mathcal{RA}$ is a sequence of states $q_0q_1\ldots$ such that $q_i = \kappa(q_{i-1}, \alpha)$ for all i and for some $\alpha \in \Sigma$. The run is *accepting* if for some $(L, K) \in F$, the run intersects with L finitely many times, and with K infinitely often. An LTL formula φ over $\mathcal{AP}$ can be represented by a DRA with alphabet $2^{\mathcal{AP}}$ that accepts exactly those runs that satisfy φ.

B. Labeled Stochastic Games and Markov Chains

A stochastic game (SG) involves two players, and starts with the system in a particular state. Transitions to subsequent states are probabilistically determined by the current state and the actions chosen by each player.

Definition 3 (Stochastic Game)**.** *A* stochastic game *[15] is a tuple $\mathcal{G} := (S, U_{def}, U_{adv}, \mathbb{T}, \mathcal{AP}, \mathcal{L})$. S is a finite set of states, U_{def} and U_{adv} are finite sets of actions of the defender and adversary. The function $\mathbb{T} : S \times U_{def} \times U_{adv} \times S \to [0, 1]$ encodes $\mathbb{T}(s'|s, u_{def}, u_{adv})$, the probability of transition from state s to state s' when defender and adversary actions are u_{def} and u_{adv}. $\mathcal{AP}$ is a set of atomic propositions. $\mathcal{L} : S \to 2^{\mathcal{AP}}$ is a labeling function that maps a state to a subset of atomic propositions that are satisfied in that state.*

SGs can be viewed as an extension of Markov Decision Processes (MDPs) when there is more than one player taking an action. For a player in an SG, a *policy* is a mapping from sequences of states to actions, if it is deterministic, or from sequences of states to a probability distribution over actions, if it is randomized. A policy is *Markov* if it is dependent only on the most recent state. It is *stationary* if it does not depend on time. We denote the defender's policy by μ and the adversary's policy by τ.

In this paper, we focus our attention on the Stackelberg setting [35], where the first player (leader) commits to a policy. The second player (follower) observes this and chooses its policy as the best response to the leader's policy, defined as the policy that maximizes the follower's utility. We assume that the players take their actions concurrently at each time step. We define the notion of a Stackelberg equilibrium (SE), which indicates that a solution to a Stackelberg game has been found. Let $Q_L(l, f)$ ($Q_F(l, f)$) be the utility gained by the leader (follower) by adopting a policy l (f).

Definition 4 (Stackelberg Equilibrium)**.** *A pair (l, f) is a* Stackelberg equilibrium *if $l = \arg\max_{l'} Q_L(l', BR(l'))$, where $BR(l') = \{f : f = \arg\max Q_F(l', f)\}$. That is, the leader's policy is optimal given that the follower observes the leader's policy and plays its best response.*

Given μ and τ, $\mathcal{G}$ is a *Markov chain (MC)* [36]. For $s, s' \in S$, s' is *accessible* from s, written $s \to s'$, if $\mathbb{P}(s_a|s)\mathbb{P}(s_b|s_a)\ldots\mathbb{P}(s_i|s_j)\mathbb{P}(s'|s_i) > 0$ for some (finite subset of) states $s_a, s_b, \ldots, s_i, s_j$. Two states *communicate* if $s \to s'$ and $s' \to s$. A state is *transient* if there is a nonzero probability of not returning to it when we start from that state, and is *positive recurrent* otherwise. In a finite state MC, every state is either transient or positive recurrent. We define labeled Markov chains, and a measure on this entity.

Definition 5 (Labeled Markov Chain)**.** *A labeled Markov chain is a tuple $\mathcal{M} = (S, \mathbb{T}, \mathcal{AP}, \mathcal{L})$, where $S, \mathcal{AP}, \mathcal{L}$ are as in Definition 3, and $\mathbb{T} : S \times S \to [0, 1]$ encodes $\mathbb{P}(s'|s)$*[1]*. A path on $\mathcal{M}$ is a sequence of states $S = s_0s_1, \ldots$ such that $s_i \in S$ and $\mathbb{T}(s_{i+1}, s_i) > 0$ for all i. We write $\mathtt{Paths}(s)$ to denote the set of paths in $\mathcal{M}$ that start from state s.*

A labeled MDP is defined by augmenting a finite set of actions Act to the Markov chain in Definition 5, and redefining the transition function as $\mathbb{T} : S \times Act \times S \to [0, 1]$ to encode $\mathbb{P}(s'|s, a)$. A policy on this labeled MDP is a map from sequences of states to actions (or, distributions over actions). In order to distinguish between policies synthesized to satisfy the LTL specification, and policies designed to ensure differential privacy, we will refer to the latter as a *scheduler*, denoted Γ, and the Markov chain so induced is $\mathcal{M}^\Gamma$. In this paper, we will assume that the output of the scheduler will depend only on the most recent state.

To quantitatively reason about $\mathcal{M}$, we need to define an appropriate probability space. We follow the treatment in [2], and let the sample space be $\mathtt{Paths}(s)$. The set of events $\mathcal{F}$ is the smallest $\sigma-$algebra on $\mathtt{Paths}(s)$ that contains all *cylinder sets*[2] spanned by finite length paths in $\mathcal{M}$. Then, there is a unique probability measure on $\mathcal{M}$ associated to this $\sigma-$algebra. However, this measure is a value on the cylinder sets. The next definition assigns a measure to an arbitrary state of $\mathcal{M}$ [13], [14].

[1]We further assume that exactly one atomic proposition will be true in a state of $\mathcal{M}$. Therefore, $\mathcal{L} : S \to \mathcal{AP}$. This is not restrictive since if labels l_1, l_2 are true, we can define $l_{12} = l_1 \wedge l_2$, and augment l_{12} to $\mathcal{L}$.

[2]The cylinder set of $\omega = s_0s_1\ldots, s_n$ is $Cyl(\omega) := \{\omega' \in \bigcup_s \mathtt{Paths}(s) | \omega \text{ is a prefix of } \omega'\}$.

Definition 6 (Measure on $s \in S$ in $\mathcal{M}$)**.** *Let* $\mathbb{T}^+(s_0, s_1, \ldots, s_k) = \prod_{i=0}^{k-1} \mathbb{P}(s_{i+1}|s_i)$ *and* $\mathcal{L}^+(s_0, \ldots, s_k) = \mathcal{L}(s_0)\ldots\mathcal{L}(s_k)$*. Then, for* $s \in S$*,* $\nu_s : \mathcal{F} \to [0,1]$ *is the unique measure on* $\mathcal{F}$ *such that for any cylinder set* $Cyl(\omega)$*,* $\nu_s(Cyl(\omega)) = \sum\{\mathbb{T}^+(p) : p \in \texttt{Paths}(s), \mathcal{L}^+(p) = \omega\}$

C. Differential Privacy

Differential privacy is a property that ensures that private data of an agent is protected, while allowing for statistical inferences from aggregates of the data [37]. This makes it unlikely that an adversary will learn anything meaningful about sensitive data. Attractive features of differential privacy include compositionality, resilience to post-processing, and robustness to side information. The notion of differential privacy that we use in this paper is defined over $\mathcal{M}$.

Definition 7 ($(\epsilon, \delta)-$Differential Privacy)**.** *Let* $R \subset S \times S$ *be a symmetric relation. Then, given* $\epsilon \geq 0$ *and* $\delta \in [0,1]$*, a labeled Markov chain* $\mathcal{M}$ *is* $(\epsilon, \delta)-$*differentially private with respect to* R *if for every* $s, s' \in S$ *such that* $(s, s') \in R$*,* $\nu_s(E) \leq e^{\epsilon}\nu_{s'}(E) + \delta$ *for every measurable subset* $E \in \mathcal{F}$*.*

A measure of how different states $s, s' \in S$ in $\mathcal{M}$ are can be quantified using the *total variation distance*, given by:

$$tv(s, s') := tv(\nu_s, \nu_{s'}) = \sup_{E \in \mathcal{F}} |\nu_s(E) - \nu_{s'}(E)|$$

IV. Problem Formulation

Problem 1. *Given a stochastic game* $\mathcal{G}$ *representing the environment, an LTL specification* φ*, and parameters* ϵ, δ*:* i)*: determine a defender policy* μ *to maximize the satisfaction probability of* φ *under any adversary policy* τ*;* ii)*: determine a symmetric relation* R *so that for any trajectory from* i) *that satisfies* φ*, there is some other trajectory such that states* s, s' *along the two trajectories are in* R*, and the two trajectories are* $(\epsilon, \delta)-$*differentially private to an eavesdropper.*

We consider two adversaries, each having different capabilities. The first adversary can inject inputs into the system in order to influence transitions between states in the CPS. This sequence of inputs of this adversary is determined by the policy τ in Problem 1. The second adversary is an eavesdropper, who can observe trajectories of the system as a result of the policy synthesized by the defender, and can potentially use this information to launch an attack on the system. Ensuring differential privacy of this aforementioned trajectory will ensure that states along the trajectory are relatively indistinguishable to this eavesdropper.

Assumption 1. *The two adversaries act independent of each other, and do not communicate with each other.*

V. Solution Approach

We adopt a two-step approach to solve Problem 1. In the first step, we will determine a defender policy to maximize the probability of satisfying the LTL objective under any adversary policy. In the second step, we will use optimal policies that satisfy φ from the previous step to define a symmetric relation R on the states to ensure $(\epsilon, \delta)-$differential privacy of a labeled trajectory observed by the eavesdropper. We will impose an additional constraint on the 'closeness' of policies from states that are 'related' to each other (this will be made precise later in this section) to ensure differential privacy of subsequent states along the trajectory while maintaining satisfaction of the LTL objective.

The representation of the environment when composed with that of φ yields a product game. For an instantiation of defender and adversary policies, this is a Markov chain. These policies will induce paths on the Markov chain, and we want to ensure differential privacy of states along these paths. We will use a fragment of the temporal logic PCTL* [2] to reason about probabilities over paths on this entity.

A. Synthesis of Policies to Satisfy LTL Objective

In order to find runs on $\mathcal{G}$ that would be accepted by the DRA $\mathcal{RA}$ corresponding to the LTL task φ, we construct an entity that composes representations of the environment ($\mathcal{G}$) and the goal ($\mathcal{RA}$). We call this a *product game*.

Definition 8 (Product Stochastic Game (PSG))**.** *Given SG* $\mathcal{G}$ *and DRA* $\mathcal{RA}$ *corresponding to LTL formula* φ*, a PSG is a tuple* $\mathcal{G}^{\varphi} := (S^{\varphi}, U_{def}, U_{adv}, \mathbb{T}^{\varphi}, F^{\varphi}, \mathcal{AP}, \mathcal{L}^{\varphi})$*, where* $S^{\varphi} = S \times Q$*,* $\mathbb{T}^{\varphi}((s', q')|(s, q), u_{def}, u_{adv}) = \mathbb{T}(s'|s, u_{def}, u_{adv})$ *iff* $\kappa(q, \mathcal{L}(s')) = q'$ *and* 0 *otherwise,* $F^{\varphi} := \{(L^{\varphi}(i), K^{\varphi}(i)\}_{i=1}^{M}$ *is such that* $L^{\varphi}(i), K^{\varphi}(i) \subseteq S^{\varphi}$*, and* $(s, q) \in L^{\varphi}(i)$ *iff* $q \in L(i)$ *and* $(s, q) \in K^{\varphi}(i)$ *iff* $q \in K(i)$*,* $\mathcal{L}^{\varphi}((s, q)) = \mathcal{L}(s)$*.*

As a first step, we are interested in the synthesis of defender policies that would satisfy the LTL objective under any adversary policy. This will be equivalent to reaching certain recurrent subsets of a Markov chain formed under instantiations of these policies. Maximizing the probability of satisfaction in this setting corresponds to reaching an equilibrium of a zero-sum Stackelberg game between the defender and adversary. A careful justification of this assertion with detailed proofs for fully and partially observable environments has been studied in our earlier works [15]–[18]. We only state relevant results that will be useful in our goal to further establish differential privacy of trajectories in $\mathcal{G}^{\varphi}$ that will satisfy φ under the respective agent policies.

Proposition 1. *Let* $v(s, q) = \max_{\mu} \min_{\tau} \mathbb{P}(\varphi|\mu, \tau, (s, q))$*. Then,*

$$v(s,q) = \max_{\mu} \min_{\tau} \sum_{u_d \in U_{def}} \sum_{u_a \in U_{adv}} \sum_{(s',q') \in S^{\varphi}} \mu(u_{def}|(s,q)) \times \tau(u_{adv}|(s,q))\mathbb{T}^{\varphi}((s',q')|(s,q), u_{def}, u_{adv})v((s',q'))$$

Proof. The proof can be found in Lemma 1 of [16]. □

Let $\mathcal{E}$ denote the set of *accepting states* of $\mathcal{G}^{\varphi}$. That is, a subset of recurrent states of the stochastic game which also satisfy φ. We note that in the terminology of [15], [16] this set is part of a *generalized maximal accepting end component*, while [17], [18] use the term $\varphi-$ *feasible recurrent set*. Let $\mathbb{P}(\texttt{reach}\ \mathcal{E}|s, \mu, \tau)$ denote the probability of reaching the set of states $\mathcal{E}$ in $\mathcal{G}^{\varphi}$. We have the following result for stationary Markov policies μ, τ.

Theorem 1. *For a stationary Markov defender policy μ, and initial state s, the following holds:*

$$\min_{\tau} \mathbb{P}(\mathcal{G}^{\varphi} \models \varphi | s, \mu, \tau) = \min_{\tau} \mathbb{P}(\texttt{reach}\ \mathcal{E} | s, \mu, \tau)$$

Proof. The proof can be found in Proposition 3 of [16]. An analogous result for the partially observable state setting can be found in Theorem 4.3 of [17]. □

Therefore, the problem of maximizing the probability of satisfaction of φ under any adversary policy is equivalent to reaching a subset of states under these policies of the product game $\mathcal{G}^{\varphi}$ that composes representations of the environment and the LTL objective. Moreover, results in [15]–[18] establish convergence of these policies to an equilibrium of the Stackelberg game between the defender and adversary.

B. The Temporal Logic PCTL*

The semantics of LTL, as seen in Definition 1 are defined over infinite words. In order to establish differential privacy, we require a means to reason over paths in a Markov chain or Markov decision process. The temporal logic PCTL* is appropriate in this setting. A PCTL* formula comprises state and path formulas and its semantics are expressed over Markov chains or MDPs. We will work with a fragment of PCTL* whose path formulas are LTL formulas.

An *PCTL* state formula* [2] is defined over a set of atomic propositions $\mathcal{AP}$, and can be written as: $\Phi := \mathtt{T} | \sigma | \neg\Phi | \Phi \wedge \Phi | \mathbb{P}_J(\phi)$, where ϕ is a *path formula*, which is an LTL formula whose sub-state formulas are PCTL* state formulas. This is formed according to $\phi := \Phi | \neg\phi | \phi \wedge \phi | \mathbf{X}\phi | \phi \mathbf{U} \phi$, where Φ is a PCTL* state formula. $J \subseteq [0,1]$ is a non-empty interval with rational end points. The semantics of PCTL* are expressed over an MDP $\mathcal{M}^{\Gamma}$, where Γ is a scheduler. In the sequel, we omit the superscript Γ, and use it only when the context will not be clear otherwise.

Definition 9 (PCTL* Semantics)**.** *The semantics of PCTL* state formula are defined on states of a labeled MDP $\mathcal{M}$ as:*

1) $(\mathcal{M}, s) \models \mathtt{T}$ *iff* $\mathtt{Paths}(s)$ *is true for all* s
2) $(\mathcal{M}, s) \models \sigma$ *iff* $\sigma \in \mathcal{L}(s_0)$
3) $(\mathcal{M}, s) \models \neg\Phi$ *iff* $(\mathcal{M}, s) \not\models \Phi$
4) $(\mathcal{M}, s) \models \Phi_1 \wedge \Phi_2$ *iff* $(\mathcal{M}, s) \models \Phi_1$ *and* $(\mathcal{M}, s) \models \Phi_2$.
5) $(\mathcal{M}, s) \models \mathbb{P}^{\Gamma}_J(\phi)$ *iff* $\mathcal{M}$, $\mathbb{P}[(\mathcal{M}^{\Gamma}, s) \models \phi] \in J$ *for some scheduler* Γ.

The semantics of PCTL path formulas are defined as:*

1) $(\mathcal{M}, S) \models \Phi$ *iff* $(\mathcal{M}, s_0) \models \Phi$
2) $(\mathcal{M}, S) \models \mathbf{X}\phi$ *iff* $(\mathcal{M}, \mathtt{Paths}(s_1)) \models \phi$
3) $(\mathcal{M}, S) \models \phi_1 \mathbf{U} \phi_2$ *iff* $\exists j \geq 0$ *such that* $(\mathcal{M}, \mathtt{Paths}(s_j)) \models \phi_2$ *and for all* $k < j$, $(\mathcal{M}, \mathtt{Paths}(s_k)) \models \phi_1$.

Problem 1 can then be interpreted as finding a defender policy to satisfy φ under any adversary policy, and then synthesizing a scheduler so that $(\epsilon, \delta)-$differential privacy holds for states along trajectories generated according to the policies in the previous step. The scheduler will be related to defining a symmetric relation between states, and ensuring that policies computed at these states are 'close' to each other so that subsequent states along the trajectories remain related. In this manner, the original objective of maximizing the probability of satisfying φ will also not be affected.

C. Distances between States and Differential Privacy

The eavesdropper observes a sequence of labels, corresponding to labels on states along a trajectory. The eavesdropper is also assumed to have access to optimal defender policies at each state synthesized in the previous step. Informally, a trajectory will be differentially private if at each state along the trajectory, the adversary will be unable to associate a unique state $s = \mathcal{L}^{-1}(l)$, where $\mathcal{L}^{-1} : 2^{\mathcal{AP}} \to 2^{S}$.

To capture the relation between the notion of differential privacy in Definition 7 and paths in a Markov chain, we use the notion of a *skewed distance* from [13].

Definition 10 (Skewed Total Variation Distance)**.** *For $\alpha \geq 1$, the skewed distance is a quantity $\Delta_{\alpha} : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that $\Delta_{\alpha}(x, y) := \max\{x - \alpha y, y - \alpha x, 0\}$. Then, given $s, s' \in S$, the skewed total variation distance is*

$$tv_{\alpha}(s, s') := tv_{\alpha}(\nu_s, \nu_{s'}) = \sup_{E \in \mathcal{F}} \Delta_{\alpha}(\nu_s(E), \nu_{s'}(E))$$

This gives the following result [13]:

Proposition 2. *The labeled MC $\mathcal{M}$ is $(\epsilon, \delta)-$differentially private with respect to $R \subset S \times S$ if for every $s, s' \in S$ such that $(s, s') \in R$, we have $tv_{\alpha}(s, s') \leq \delta$, where $\alpha = e^{\epsilon}$.*

To establish differential privacy, we have to define a notion of equivalence on states of the MDP (or Markov chain) so that they are indistinguishable from each other. We use a skewed variant of the Kantorovich distance[3] [26] to translate the distance between MDP states to a distance over probability measures [28]. This is needed since the consequence of taking an action in a state often yields a distribution over the next state, which is why we need to '*lift*' tv_{α} from a relation over states to one over distributions.

Definition 11 (Skewed Kantorovich Metric)**.** *Let ω, ω' be probability distributions over states S of an MDP $\mathcal{M}$. Then, given a symmetric distance $d : S \times S \to [0, 1]$, the skewed Kantorovich distance between ω and ω', $K^{d}_{\alpha}(\omega, \omega')$, is the maximum value of a pair of linear programs- the first is given below; the second is got by reversing ω and ω':*

$$\max_{x_s} \sum_{s \in S} (\omega(s) - \alpha\omega'(s)) x_s \tag{1}$$

$$\text{subject to: } x_s - \alpha x_{s'} \leq d(s, s'), \quad \text{for all } s, s' \in S$$
$$0 \leq x_s \leq 1, \qquad \text{for all } s \in S$$

Lemma 1. *[13] Let $d(s, s') = \mathbf{I}_{\neq}(s, s')$, where $\mathbf{I}_{\neq}$ takes value 1 if its arguments are not equal, and 0 otherwise. Then $K^{d}_{\alpha}(\omega, \omega') = tv_{\alpha}(s, s')$.*

[3]This metric has its origins in optimal transportation theory, but has since been used in several applications in computer science [38].

D. Value Iteration for Differential Privacy

The approach we take is to ensure that differential privacy will hold at each state of the trajectory, starting from the initial state. To achieve this, we define a symmetric relation between states in a recursive manner. We further stipulate that at related states, optimal defender policies synthesized in the previous step are close to each other, and that subsequent states resulting from these policies be related. We formalize this approach and prove relevant results in this section.

We recall standard terminology from the MDP literature [39]. The *value* of a state s under scheduler Γ is $V^{\Gamma}(s) := \mathbb{E}[\sum_{t=0}^{\infty} \beta^t r_t | s_0 = s, \Gamma]$, where s_0 is the initial state, $\beta \in (0,1)$ is a discount factor, and r_t is a *reward* received at time t. The expectation is taken over the transitions of state induced by Γ. The goal is to determine Γ to maximize $V^{\Gamma}(s)$ for each s. The optimal value function V^* will then satisfy the *Bellman optimality equations*, given by:

$$V^*(s) = \max_{a \in A}(r_s^a + \beta \sum_{s_1 \in S} \mathbb{P}(s_1|s,a)V^*(s_1)), \forall s \in S \quad (2)$$

Lemma 2. *[39] Let $V_0(s) = 0$ and $V_{n+1}(s) = \max_{a \in A}(r_s^a + \beta \sum_{s_1 \in S} \mathbb{P}(s_1|s,a)V_n(s_1)$. Then, $\{V_n(s)\}_{n \geq 1}$ converges uniformly to $V^*(s)$ for each $s \in S$.*

Since our treatment is related to maximizing the satisfaction probability, we will assume that $r_s^a = 0$ at all states other than accepting states of the product game (where this reward will be set to 1). In this case, the value of a state will be related to the probability of satisfying the LTL goal, and therefore will be in $[0,1]$. Consequently, we can set $\beta = 1$.

Define an operator F_{tv_α} as follows:

$$F_{tv_\alpha}(s,s') := \begin{cases} \max_{a \in A}(tv_\alpha(s,s')) & \mathcal{L}(s) = \mathcal{L}(s') \\ 1 & \mathcal{L}(s) \neq \mathcal{L}(s') \end{cases}$$

Let $F_{tv_\alpha}^n$ denote the composition of F_{tv_α} n times.

Theorem 2. $\Delta_\alpha(V_n(s), V_n(s')) \leq F_{tv_\alpha}^n(s,s')\ \forall s, s' \in S$.

Proof. From the definition of $F_{tv_\alpha}(s,s')$, the claim holds when $\mathcal{L}(s) \neq \mathcal{L}(s')$. Now, let $\mathcal{L}(s) = \mathcal{L}(s')$. Since $V_0(s) = 0$ for all s, the base case of the induction holds. Now, suppose $\Delta_\alpha(V_k(s) - V_k(s')) \leq F_{tv_\alpha}^k(s,s')$ is true for some k. Then, we have:

$$\begin{aligned} &V_{k+1}(s) - \alpha V_{k+1}(s') \\ &= \max_a[\sum_{u \in S} \mathbb{P}(u|s,a)V_n(u)] - \alpha \max_a[\sum_{u \in S} \mathbb{P}(u|s',a)V_k(u)] \\ &\leq \max_a[\sum_{u \in S}(\mathbb{P}(u|s,a) - \alpha \mathbb{P}(u|s',a))V_k(u)] \end{aligned}$$

Writing down the above set of inequalities for $V_{k+1}(s') - \alpha V_{k+1}(s)$, we observe that $V_k(u)$ is a feasible solution to the skewed Kantorovich metric $tv_\alpha^k(s,s')$ (Definition 11). Therefore, we get: $\Delta_\alpha(V_{k+1}(s) - V_{k+1}(s')) \leq \max_a[tv_\alpha^k(s,s')] = F_{tv_\alpha}(F_{tv_\alpha}^k(s,s')) = F_{tv_\alpha}^{k+1}(s,s')$, completing the proof. □

Proposition 3. *Let $F_{tv_\alpha}^*(s,s')$ be the least fix point of $F_{tv_\alpha}(s,s')$. As $n \to \infty$, $\Delta_\alpha(V^*(s), V^*(s')) \leq F_{tv_\alpha}^*(s,s')$.*

Proof. This follows from the fact that in $K_\alpha^d(\omega, \omega')$, if we equip $d : S \times S \to [0,1]$ with a point-wise ordering, we get a complete lattice. When $K_\alpha^d(\omega, \omega') = tv_\alpha(s,s')$, then $F_{tv_\alpha}(s,s')$ is monotone with respect to this order. Consequently, it admits a least fixed point [40]. □

These results allow us to aggregate related states. Further, we restrict the actions (the set A in Equation (2)) to those determined by the optimal defender policy μ^*, where $\mu^* := \arg\max_{\mu} \min_{\tau} \mathbb{P}(\varphi)$. If $\mu^*(s)$ denotes the actions available at state s per the policy μ^*, we write $A(s) := \{a : a \in \mu^*(s)\}$.

Assumption 2. *Assume that at each non-terminal state s along the trajectory that satisfies φ, there is some action a such that the LTL objective will be satisfied.*

Assumption 2 stipulates that trajectories that will be generated according to the policies μ^* in Section V-A satisfy the LTL objective with non-zero probability. A more thorough analysis of ensuring differential privacy of trajectories that may not satisfy the LTL goal at all is left as future work.

In the sequel, we abuse notation to say that a state of the labeled MDP is a *terminal state* if it corresponds to an accepting state of the product game formed by composing representations of the environment and the LTL objective.

Theorem 3. *Consider a symmetric relation R defined recursively as: $(s,s') \in R$ if and only if:*

1) *$\mathcal{L}(s) = \mathcal{L}(s')$,*
2) *$tv(\mu^*(s), \mu^*(s')) \leq M$ where $M \ll 1/\alpha$,*
3) *$\exists u, u' \in S, a_s \in \mu^*(s), a_{s'} \in \mu^*(s')$ such that $\mathbb{P}(u|s,a_s) > 0, \mathbb{P}(u'|s',a_{s'}) > 0$, and $(u,u') \in R$ for all non-terminal states u, u',*
4) *$tv_\alpha(s,s') = 0$ for terminal states s, s'.*

Consider two trajectories, one starting from $s_0 \in S$, and the other from $s_0' \in S$. Then, if $(s_0, s_0') \in R$, each state along the two trajectories, written (s,s'), will be $(\epsilon, \delta)-$differential private for $\alpha = e^\epsilon$ if $\delta \geq \alpha M + \max_{(s,s') \in R} F_{tv_\alpha}^*(s,s')$.

Proof. $F_{tv_\alpha}^*(s,s') \geq tv_\alpha(s,s')$ since $F_{tv_\alpha}^*$ is a least fixed-point. From Thm. 2 and Proposition 3, $\Delta_\alpha(V^*(s), V^*(s')) \leq F_{tv_\alpha}^*(s,s')$. The second condition for R allows us to relax the requirement that the optimal action from s and that from s' is the same. This will allow two states to be related if the optimal (stochastic) defender policies from these states are 'close' to each other. The third condition ensures that these policies lead to transitions to states that will also be related to each other. Now, since $V^*(s) = \max_\Gamma \mathbb{P}(\varphi|s)$, where Γ comprises actions from μ^*, $\Delta_\alpha(V^*(s), V^*(s'))$ is $|\mathbb{P}(\varphi|\texttt{Paths}(s)) - \alpha \mathbb{P}(\varphi|\texttt{Paths}(s'))|$. This is the same as $tv_\alpha(s,s')$, but along measurable sets determined by Γ. From Proposition 2, we want this quantity, plus an additional term that accounts for the closeness of optimal policies to be less than δ. Therefore, $\delta \geq F_{tv_\alpha}^*(s,s') + \alpha M$. The recursive definition of R requires that this property be maintained at every subsequent state until a terminal state is reached. The lower bound on δ is got by taking the maximum over least fixed points corresponding to each pair of related states. □

In Theorem 3, we compare trajectories of equal lengths. The length of a trajectory is the number of actions in the trajectory, plus one (for the initial state). Future work will study differential privacy for trajectories of different lengths.

Remark 1. *Computing $tv_\alpha(s, s')$ is not known to be decidable [41]. However, in our setting, since $V(s) = \mathbb{P}(\varphi|s)$, Theorems 2, 3, and Proposition 3 will allow us to circumvent the need to explicitly compute $tv_\alpha(\cdot, \cdot)$.*

VI. Example

We consider a CPS in the form of a drone/ UAV that has to carry out persistent surveillance of a target region. At the same time, it must avoid certain regions of the environment. This could be either due to the presence of a physical obstacle, or other factors that might compromise it (e.g. entering a region might allow the drone to be detected by a radar). This goal can be represented by the LTL formula $\varphi = \mathbf{GF}\texttt{tar} \wedge \mathbf{G}\neg\texttt{obs}$, where `tar` and `obs` are labels denoting the target and obstacle respectively. The DRA corresponding to φ will have two states q_0, q_1, with $F = (\{\emptyset\}, \{q_1\})$.

The dynamics model of the UAV motion is inspired from [12]. The satisfaction of φ can be affected by an adversary that can influence trajectories of the UAV. A stochastic-game abstraction of the UAV dynamics was presented in [16], and we will assume that we have this abstraction for the remainder of this section. The environment of the drone is then an $M \times N$ grid, $S := \{s_i : i = x + My, x \in \{0, \ldots, M-1\}, y \in \{0, \ldots, N-1\}\}$. We assume that the drone's actions are $U_{def} = \{R, L, U, D\}$ denoting right, left, up, and down, and the actions of the adversary are $U_{adv} = \{A, NA\}$, denoting attack, and not attack respectively. Transition probabilities for $(u_{def}, u_{adv}) = (R, NA)$ and (R, A) are defined below. Probabilities for other action pairs can be defined similarly. Let N_{s_i} denote the neighbors of s_i.

$$\mathbb{T}(s_j|s_i, R, NA) = \begin{cases} 0.8 & j = i+1,\ (i+1) \not\equiv 0 \bmod M \\ \frac{0.2}{|N_{s_i}|} & (s_j \in \{s_i\} \cup N_{s_i} \setminus \{s_{i+1}\}), \\ & (i+1) \not\equiv 0 \bmod M \\ 1 & j = i \text{ and } (i+1) \equiv 0 \bmod M \end{cases}$$

$$\mathbb{T}(s_j|s_i, R, A) = \begin{cases} 0.6 & j = i+1,\ (i+1) \not\equiv 0 \bmod M \\ \frac{0.4}{|N_{s_i}|} & (s_j \in \{s_i\} \cup N_{s_i} \setminus \{s_{i+1}\}), \\ & (i+1) \not\equiv 0 \bmod M \\ 1 & j = i \text{ and } i+1 \equiv 0 \bmod M \end{cases}$$

We use the method from Section V-A to synthesize a drone policy that will maximize its probability of satisfying this objective under any adversary policy. Once this has been done, we want to ensure that a second, independent adversary eavesdrops on a labeled trajectory that satisfies the LTL objective, will not be able to discern the position of the drone. We use the skewed total variation distance in order to establish differential privacy of the trajectory at each state along the trajectory. Specifically, two states will be related if they satisfy the conditions in Theorem 3.

We report results of our experiments for $M = N = 10$, as shown in Figure 1. In the first step, using the results in Section V-A, we synthesize a policy for the defender that will maximize the probability of satisfying φ under any adversary policy. We consider a trajectory that satisfies φ, and use results in Sections V-C and V-D to observe the effects of values of ϵ and δ on the number of differentially private trajectories that can be realized. Specifically, different values of ϵ and δ will yield varying numbers of differentially private trajectories, under an additional assumption that the optimal defender policies at states which are related to each other are sufficiently close. We observe that larger values of ϵ and δ allow for more differentially private trajectories, as seen in Figure 2. Figure 1 also shows a fragment of length 25

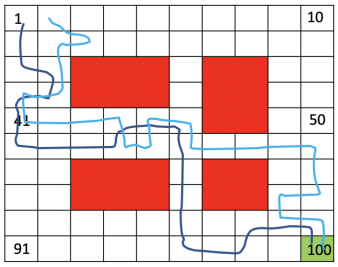

Fig. 1: The environment as a 10×10 grid. The objective for the CPS, given by the LTL formula $\varphi = \mathbf{GF}\texttt{tar} \wedge \mathbf{G}\neg\texttt{obs}$, is to visit the target (green state) infinitely often while always avoiding obstacles (red states). The CPS has to satisfy φ in the presence of an adversary who can insert inputs that affect transitions between successive states. After a trajectory that satisfies φ has been realized, a second eavesdropper adversary should not be to distinguish this trajectory from trajectories 'sufficiently close' to it. We use differential privacy as a metric to decide if the latter objective will be satisfied. The figure shows a fragment of length 25 of two trajectories that satisfy φ and are differentially private for $\epsilon = 1, \delta = 0.001$. The trajectories start from states that are related to each other, and at each state, the distance between the optimal defender policies at the states is below a threshold, and subsequent states are related to each other. An eavesdropper will not be able to distinguish between these trajectories only by observing labels on the states.

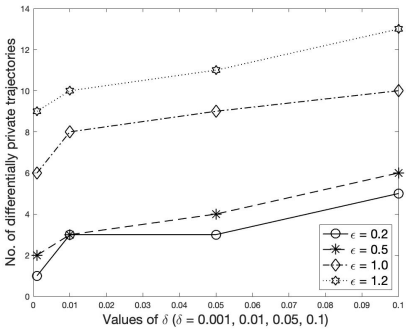

Fig. 2: Number of trajectories differentially private to one satisfying trajectory for different values of ϵ and δ. Larger ϵ and δ allow for more differentially private trajectories.

of two trajectories that are $(1, 0.001)-$differentially private. That is, by simply observing labels on the states, and with knowledge of the optimal defender policies at each state, with high probability, the eavesdropper adversary will not be able to distinguish between the two trajectories.

VII. Conclusion

This paper presented a solution to the problem of ensuring satisfaction of an LTL objective φ while maintaining privacy of trajectories in the presence of two kinds of adversaries. We modeled the interaction between the CPS and an adversary A who could tamper with actuator inputs as a stochastic game. Maximizing the probability of satisfying φ under actions of A was equivalent to reaching a Stackelberg equilibrium of this game. We then characterized the indistinguishability of trajectories to an eavesdropper E by ensuring differential privacy of states along trajectories that satisfied φ. We showed that if a distance between probabilities of satisfying of φ starting from states that were related to each other was below a threshold, then this ensured that these states were differentially private. An additional requirement on the distance between optimal policies at related states ensured differential privacy at every state along the trajectories. We validated our approach on a simulation of a UAV that had to satisfy an LTL objective in the presence of adversarial inputs in an environment with an eavesdropper.

Future work will seek to generalize our setup to ensuring differential privacy of trajectories that may not satisfy φ. We will also extend our analysis to the case where trajectories may have different lengths.

References

[1] R. Baheti and H. Gill, "Cyber-physical systems," *The Impact of Control Technology*, vol. 12, no. 1, pp. 161–166, 2011.

[2] C. Baier and J.-P. Katoen, *Principles of Model Checking*. MIT Press, 2008.

[3] J. E. Sullivan and D. Kamensky, "How cyber-attacks in Ukraine show the vulnerability of the US power grid," *The Electricity Journal*, vol. 30, no. 3, pp. 30–35, 2017.

[4] Y. Shoukry, P. Martin, P. Tabuada, and M. Srivastava, "Non-invasive spoofing attacks for anti-lock braking systems," in *Cryptographic Hardware and Embedded Systems*. Springer, 2013, pp. 55–72.

[5] A. Banerjee, K. K. Venkatasubramanian, T. Mukherjee, and S. K. Gupta, "Ensuring safety, security, and sustainability of mission-critical CPSs," *Proceedings of the IEEE*, vol. 100, no. 1, pp. 283–299, 2012.

[6] R. Focardi and R. Gorrieri, "A taxonomy of trace-based security properties for CCS," in *The Computer Security Foundations Workshop VII*. IEEE, 1994, pp. 126–136.

[7] C. Dwork, "Differential privacy: A survey of results," in *International Conference on Theory and Applications of Models of Computation*. Springer, 2008, pp. 1–19.

[8] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography Conference*. Springer, 2006, pp. 265–284.

[9] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, 2013.

[10] Z. Xu, K. Yazdani, M. T. Hale, and U. Topcu, "Differentially private controller synthesis with metric temporal logic specifications," in *American Control Conference (ACC)*, 2020, pp. 4745–4750.

[11] M. Hale, A. Jones, and K. Leahy, "Privacy in feedback: The differentially private LQG," in *American Control Conference (ACC)*. IEEE, 2018, pp. 3386–3391.

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

[13] D. Chistikov, A. S. Murawski, and D. Purser, "Bisimilarity distances for approximate differential privacy," in *Symposium on Automated Technology for Verification and Analysis*, 2018, pp. 194–210.

[14] ——, "Asymmetric distances for approximate differential privacy," in *International Conference on Concurrency Theory*, 2019.

[15] L. Niu and A. Clark, "Secure control under LTL constraints," in *Proc. American Control Conference (ACC)*, 2018, pp. 3544–3551.

[16] ——, "Optimal secure control with linear temporal logic constraints," *IEEE Transactions on Automatic Control*, vol. 65, no. 6, pp. 2434–2449, 2019.

[17] B. Ramasubramanian, A. Clark, L. Bushnell, and R. Poovendran, "Secure control under partial observability with temporal logic constraints," in *Proc. American Control Conference*, 2019, pp. 1181–1188.

[18] B. Ramasubramanian, L. Niu, A. Clark, L. Bushnell, and R. Poovendran, "Secure control in partially observable environments to satisfy LTL specifications," *arXiv preprint arXiv: 2007.12501*, 2020.

[19] L. Niu, B. Ramasubramanian, A. Clark, L. Bushnell, and R. Poovendran, "Control synthesis for cyber-physical systems to satisfy metric interval temporal logic objectives under timing and actuator attacks," in *ACM/ IEEE International Conference on Cyber-physical Systems*. IEEE, 2020, pp. 162–173.

[20] ——, "Robust satisfaction of metric interval temporal logic objectives in adversarial environments," *Submitted*, 2020.

[21] M. R. Clarkson and F. B. Schneider, "Hyperproperties," *Journal of Computer Security*, vol. 18, no. 6, pp. 1157–1210, 2010.

[22] E. Ábrahám and B. Bonakdarpour, "HyperPCTL: A temporal logic for probabilistic hyperproperties," in *International Conference on Quantitative Evaluation of Systems*. Springer, 2018, pp. 20–35.

[23] Y. Wang, M. Zarei, B. Bonakdarpour, and M. Pajic, "Statistical verification of hyperproperties for cyber-physical systems," *Transactions on Embedded Computing Systems*, vol. 18, no. 5s, pp. 1–23, 2019.

[24] J. Yang, Y. Cao, and H. Wang, "Differential privacy in probabilistic systems," *Information and Computation*, vol. 254, pp. 84–104, 2017.

[25] V. Castiglioni, K. Chatzikokolakis, and C. Palamidessi, "A logical characterization of differential privacy," *Science of Computer Programming*, vol. 188, p. 102388, 2020.

[26] L. Kantorovich, "On the transfer of masses," in *Doklady Akademii Nauk USSR*, vol. 37, no. 7–8, 1942, pp. 227–229.

[27] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden, "Metrics for labelled Markov processes," *Theoretical computer science*, vol. 318, no. 3, pp. 323–354, 2004.

[28] N. Ferns, P. Panangaden, and D. Precup, "Metrics for finite Markov decision processes." in *Conference on Uncertainty in Artificial Intelligence*, vol. 4, 2004, pp. 162–169.

[29] ——, "Bisimulation metrics for continuous MDPs," *SIAM Journal on Computing*, vol. 40, no. 6, pp. 1662–1714, 2011.

[30] B. Balle, M. Gomrokchi, and D. Precup, "Differentially private policy evaluation," in *International Conference on Machine Learning*, 2016.

[31] Y. Mo and R. M. Murray, "Privacy preserving average consensus," *IEEE Transactions on Automatic Control, 62(2)*, pp. 753–765, 2016.

[32] J. Cortés, G. E. Dullerud, S. Han, J. Le Ny, S. Mitra, and G. J. Pappas, "Differential privacy in control and network systems," in *Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 4252–4272.

[33] S. Han and G. J. Pappas, "Privacy in control and dynamical systems," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 309–332, 2018.

[34] D. Liu, B.-Y. Wang, and L. Zhang, "Model checking differentially private properties," in *Asian Symposium on Programming Languages and Systems*. Springer, 2018, pp. 394–414.

[35] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, 1991.

[36] S. P. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability*. Springer Science & Business Media, 2012.

[37] C. Dwork, A. Roth *et al.*, "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.

[38] Y. Deng and W. Du, "The Kantorovich metric in computer science: A brief survey," *Electronic Notes in Theoretical Computer Science*, vol. 253, no. 3, pp. 73–82, 2009.

[39] M. L. Puterman, *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

[40] A. Tarski, "A lattice-theoretical fixpoint theorem and its applications," *Pacific Journal of Mathematics*, vol. 5, no. 2, pp. 285–309, 1955.

[41] S. Kiefer, "On computing the total variation distance of hidden Markov models," in *International Colloquium on Automata, Languages, and Programming (ICALP)*, 2018.