



Linear Temporal Logic Satisfaction in Adversarial Environments Using Secure Control Barrier Certificates

Bhaskar Ramasubramanian^{1(✉)}, Luyao Niu², Andrew Clark², Linda Bushnell¹,
and Radha Poovendran¹

¹ 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

Abstract. This paper studies the satisfaction of a class of temporal properties for cyber-physical systems (CPSs) over a finite-time horizon in the presence of an adversary, in an environment described by discrete-time dynamics. The temporal logic specification is given in *safe-LTL_F*, a fragment of linear temporal logic over traces of finite length. The interaction of the CPS with the adversary is modeled as a two-player zero-sum discrete-time dynamic stochastic game with the CPS as defender. We formulate a dynamic programming based approach to determine a stationary defender policy that maximizes the probability of satisfaction of a *safe-LTL_F* formula over a finite time-horizon under any stationary adversary policy. We introduce *secure control barrier certificates* (S-CBCs), a generalization of barrier certificates and control barrier certificates that accounts for the presence of an adversary, and use S-CBCs to provide a lower bound on the above satisfaction probability. When the dynamics of the evolution of the system state has a specific underlying structure, we present a way to determine an S-CBC as a polynomial in the state variables using sum-of-squares optimization. An illustrative example demonstrates our approach.

Keywords: Linear temporal logic · *safe-LTL_F* · Dynamic programming · Secure control barrier certificate · Sum-of-squares optimization

1 Introduction

Cyber-physical systems (CPSs) use computing devices and algorithms to inform the working of a physical system [8]. These systems are ubiquitous, and vary in

This work was supported by the U.S. Army Research Office, the National Science Foundation, and the Office of Naval Research via Grants W911NF-16-1-0485, CNS-1656981, and N00014-17-S-B001 respectively.

© Springer Nature Switzerland AG 2019

T. Alpcan et al. (Eds.): GameSec 2019, LNCS 11836, pp. 385–403, 2019.

https://doi.org/10.1007/978-3-030-32430-8_23

size and scale from energy systems to medical devices. The wide-spread influence of CPSs such as power systems and automobiles makes their safe operation critical. Although distributed algorithms and systems allow for more efficient sharing of information among parts of the system and across geographies, they also make the CPS vulnerable to attacks by an adversary who might gain access to the distributed system via multiple entry points. Attacks on distributed CPSs have been reported across multiple application domains [20, 43, 44, 46]. In these cases, the damage to the CPS was caused by the actions of a stealthy, intelligent adversary. Thus, methods designed to only account for modeling and sensing errors may not meet performance requirements in adversarial scenarios. Therefore, it is important to develop ways to specify and verify properties that a CPS must satisfy that will allow us to provide guarantees on the operation of the system while accounting for the presence of an adversary.

In order to verify the behavior of a CPS against a rich set of temporal specifications, techniques from formal methods can be used [9]. Properties like safety, stability, and priority can be expressed as formulas in linear temporal logic (LTL) [19]. These properties can then be verified using off-the-shelf model solvers [15, 28] that take these formulas as inputs. If the state space and the actions available to the agents are both finite and discrete, then the environment can be represented as a Markov decision process (MDP) [38] or a stochastic game [11]. These representations have also been used as abstractions of continuous-state continuous action dynamical system models [10, 32]. However, a significant shortcoming is that the computational complexity of abstracting the underlying system grows exponentially with the resolution of discretization desired [14, 21].

The method of barrier certificates (or barrier functions), which are functions of the states of the system was introduced in [36]. Barrier functions provide a certificate that all trajectories of a system starting from a given initial set will not enter an unsafe region. The use of barrier functions does not require explicit computation of sets of reachable states, which is known to be undecidable for general dynamical systems [29], and moreover, it allows for the analysis of general nonlinear and stochastic dynamical systems. The authors of [36] further showed that if the states and inputs to the system have a particular structure, computationally efficient methods can be used to construct a barrier certificate.

Barrier certificates were used to determine probabilistic bounds on the satisfaction of an LTL formula by a discrete-time stochastic system in [22]. A more recent work by the same authors [23] used control barrier certificates to synthesize a policy in order to maximize the probability of satisfaction of an LTL formula.

Prior work that uses barrier certificates to study temporal logic satisfaction assumes a single agent, and does not study the case when the CPS is operating in an adversarial environment. To the best of our knowledge, this paper is the first to use barrier certificates to study temporal logic satisfaction for CPSs in adversarial environments. We introduce secure barrier certificates (S-CBCs), and use it to determine probabilistic bounds on the satisfaction of an LTL formula

under any adversary policy. Further, definitions of barrier certificates and control barrier certificates in prior work can be recovered as special cases of S-CBCs.

1.1 Contributions

In this paper, we consider the setting when there is an adversary whose aim is to ensure that the LTL formula is not satisfied by the CPS (defender). The temporal logic specification is given in *safe* – LTL_F , a fragment of LTL over traces of finite length. We make the following contributions:

- We model the interaction between the CPS and adversary as a two-player dynamic stochastic game with the CPS as defender. The two players take their actions simultaneously, and these jointly influence the system dynamics.
- We present a dynamic programming based approach to determine a stationary defender policy to maximize the probability of satisfaction of an LTL formula over a finite time-horizon under any stationary adversary policy.
- In order to determine a lower bound on the above satisfaction probability, we define a new entity called *secure control barrier certificates (S-CBCs)*. S-CBCs generalize barrier certificates and control barrier certificates to account for the presence of an adversary.
- When the evolution of the state of the dynamic game can be expressed as polynomial functions of the states and inputs, we use sum-of-squares optimization to compute an S-CBC as a polynomial function of the states.
- We present an illustrative example demonstrating our approach.

1.2 Outline of Paper

We summarize related work on control barrier certificates and temporal logic satisfaction in Sect. 2. Section 3 gives an overview of temporal logic and game-theoretic concepts that will be used to derive our results. The problem that is the focus of this paper is formulated in Sect. 4. Our solution approach is presented in Sect. 5, where we define a dynamic programming operator to synthesize a policy for the defender in order to maximize the probability of satisfaction of the LTL formula under any adversary policy. We define a notion of secure control barrier certificates to derive a lower bound on the satisfaction probability, and are able to explicitly compute an S-CBC under certain assumptions. Section 6 presents an illustrative example, and we conclude the paper in Sect. 7.

2 Related Work

The method of barrier functions was introduced in [36] to certify that all trajectories of a continuous-time system starting from a given initial set do not enter an unsafe region. Control barrier functions (CBFs) were used to provide guarantees on the safety of continuous-time nonlinear systems with affine inputs for an adaptive cruise control application in [6]. The notion of input-to-state CBFs that

ensured the safety of nonlinear systems under arbitrary input disturbances was introduced in [24], and safety was characterized in terms of the invariance of a set whose computation depended on the magnitude of the disturbance. The authors of [45] relaxed the supermartingale condition that a barrier certificate had to satisfy in [36] in order to provide finite-time guarantees on the safety of a system. The verification and control of a finite-time safety property for continuous-time stochastic systems using barrier functions was recently presented in [41]. Barrier certificates were used to verify LTL formulas for a deterministic, continuous-time nonlinear dynamical system in [49]. Time-varying CBFs were used to accomplish tasks specified in signal temporal logic in [30]. A survey of the use of CBFs to design safety-critical controllers is presented in [5]. The use of barrier certificates or CBFs in these works were all for continuous time dynamical systems and did not consider the effect of the actions of an adversarial player.

Barrier certificates in the discrete-time setting were used to analyze the reachable belief space of a partially observable Markov decision process (POMDP) with applications to verifying the safety of POMDPs in [2], and for privacy verification in POMDPs in [3]. The use of barrier certificates for the verification and synthesis of control policies for discrete-time stochastic systems to satisfy an LTL formula over a finite time horizon was presented in [22] and [23]. These papers also assumed a single agent, and did not account for the presence of an adversary.

The authors of [33] used barrier functions to solve a reference tracking problem for a continuous-time linear system subject to possible false data injection attacks by an adversary, with additional constraints on the safety and reachability of the system. Probabilistic reachability over a finite time horizon for discrete-time stochastic hybrid systems was presented in [1]. This was extended to a dynamic stochastic game setting when there were two competing agents in [18], and to the problem of ensuring the safety of a system that was robust to errors in the probability distribution of a disturbance input in [50]. These papers did not assume that a temporal specification had to be additionally satisfied.

Determining a policy for an agent in order to maximize the probability of satisfying an LTL formula in an environment specified by an MDP was presented in [19]. This setup was extended to the case when there were two agents—a defender and an adversary—who had competing objectives to ensure the satisfaction of the LTL formula in an environment specified as a stochastic game in [32]. These papers assume that the states of the system are completely observable, which might not be true in every situation. The satisfaction of an LTL formula in partially observable environments represented as POMDPs was studied in [42] and the extension to partially observable stochastic games with two competing agents, each with its own observation of the state of the system, was formulated in [39].

3 Preliminaries

In this section, we give a brief introduction to linear temporal logic and discrete-time dynamic stochastic games. Wherever appropriate, we consider a

probability space $(\Omega, \mathcal{F}, \mathbb{P})$. We write $(X, \mathcal{B}(X))$ to denote the measurable space X equipped with the Borel σ -algebra, and $\mathbb{R}_{\geq 0}$ to denote the set of non-negative real numbers.

3.1 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* [9] is defined over a set of atomic propositions \mathcal{AP} , and can be written as:

$$\phi := \mathbf{T} \mid \sigma \mid \neg\phi \mid \phi \wedge \phi \mid \mathbf{X}\phi \mid \phi \mathbf{U}\phi,$$

where $\sigma \in \mathcal{AP}$, and \mathbf{X} and \mathbf{U} are temporal operators denoting the *next* and *until* operations. The semantics of LTL are defined over (infinite) words in $2^{\mathcal{AP}}$.

The syntax of linear temporal logic over finite traces, denoted LTL_F [17], is the same as that of LTL. The semantics of LTL_F is expressed in terms of *finite-length words* in $2^{\mathcal{AP}}$. We denote a word in LTL_F by η , write $|\eta|$ to denote the length of η , and η_i , $0 < i < |\eta|$, to denote the proposition at the i^{th} position of η . We write $(\eta, i) \models \phi$ when the LTL_F formula ϕ is true at the i^{th} position of η .

Definition 1 (LTL_F Semantics). *The semantics of LTL_F can be recursively defined in the following way:*

1. $(\eta, i) \models \mathbf{T}$;
2. $(\eta, i) \models \sigma$ iff $\sigma \in \eta_i$;
3. $(\eta, i) \models \neg\phi$ iff $(\eta, i) \not\models \phi$;
4. $(\eta, i) \models \phi_1 \wedge \phi_2$ iff $(\eta, i) \models \phi_1$ and $(\eta, i) \models \phi_2$;
5. $(\eta, i) \models \mathbf{X}\phi$ iff $i < |\eta| - 1$ and $(\eta, i + 1) \models \phi$;
6. $(\eta, i) \models \phi_1 \mathbf{U}\phi_2$ iff $\exists j \in [i, |\eta|]$ such that $(\eta, j) \models \phi_2$ and for all $k \in [i, j)$, $(\eta, k) \models \phi_1$.

Finally, we write $\eta \models \phi$ if and only if $(\eta, 0) \models \phi$.

Moreover, the logic admits derived formulas of the form: (i) $\phi_1 \vee \phi_2 := \neg(\neg\phi_1 \wedge \neg\phi_2)$; (ii) $\phi_1 \Rightarrow \phi_2 := \neg\phi_1 \vee \phi_2$; (iii) $\mathbf{F}\phi := \mathbf{T}\mathbf{U}\phi$ (eventually); (iv) $\mathbf{G}\phi := \neg\mathbf{F}\neg\phi$ (always). The set $\mathcal{L}(\phi)$ comprises the language of finite-length words associated with the LTL_F formula ϕ . In this paper, we focus on a subset of LTL_F called *safe-LTL_F* [40], that explicitly considers only safety properties [26].

Definition 2 (safe-LTL_F Formula). *An LTL_F formula is a safe-LTL_F formula if it can be written in positive normal form (PNF)¹, using the temporal operators \mathbf{X} (next) and \mathbf{G} (always).*

¹ In PNF, negations occur only adjacent to atomic propositions.

Next, we define an entity that will serve as an equivalent representation of an LTL_F formula, and will allow us to check if the LTL_F formula is satisfied or not.

Definition 3 (Deterministic Finite Automaton). A deterministic finite automaton (DFA) is a quintuple $\mathcal{A} = (Q, \Sigma, \delta, q_0, F)$ where Q is a nonempty finite set of states, Σ is a finite alphabet, $\delta : Q \times \Sigma \rightarrow Q$ is a transition function, $q_0 \in Q$ is the initial state, and $F \subseteq Q$ is a set of accepting states.

Definition 4 (Accepting Runs). A run of \mathcal{A} of length n is a finite sequence of $(n + 1)$ states $q_0 \xrightarrow{\sigma_0} q_1 \xrightarrow{\sigma_1} \dots \xrightarrow{\sigma_{n-1}} q_n$ such that $q_i \in \delta(q_{i-1}, \sigma_{i-1})$ for all $i \in [1, n]$ and for some $\sigma_0, \dots, \sigma_{n-1} \in \Sigma$. The run is accepting if $q_n \in F$. We write $\mathcal{L}(\mathcal{A})$ to denote the set of all words accepted by \mathcal{A} .

Every LTL_F formula ϕ over \mathcal{AP} can be represented by a DFA \mathcal{A}_ϕ with $\Sigma = 2^{\mathcal{AP}}$ that accepts all and only those runs that satisfy ϕ , that is, $\mathcal{L}(\phi) = \mathcal{L}(\mathcal{A}_\phi)$ [16]. The DFA \mathcal{A}_ϕ can be constructed by using a tool like Rabinizer4 [25].

3.2 Discrete-Time Dynamic Stochastic Games

We model the interaction between the CPS (defender) and adversary as a two-player dynamic stochastic game that evolves according to some known (discrete-time) dynamics [7]. The evolution of the state of the game at each time step is affected by the actions of both players.

Definition 5 (Discrete-time Dynamic Stochastic Game). A discrete-time dynamic stochastic game (DDSG) is a tuple $\mathcal{G} = (X, W, U_d, U_a, f, \mathcal{N}, \mathcal{AP}, L)$, where $X \subseteq \mathbb{R}^n$ and W are Borel-measurable spaces representing the state-space and uncertainty space of the system, $U_d \subseteq \mathbb{R}^d$ and $U_a \subseteq \mathbb{R}^a$ are compact Borel spaces that denote the action sets of the defender and adversary, $f : X \times U_d \times U_a \times W \rightarrow X$ is a Borel-measurable transition function characterizing the evolution of the system, $\mathcal{N} = \{0, 1, \dots, N - 1\}$ is an index-set denoting the stage of the game, \mathcal{AP} is a set of atomic propositions, and $L : X \rightarrow 2^{\mathcal{AP}}$ is a labeling function that maps states to a subset of atomic propositions that are satisfied in that state.

The evolution of the state of the system is given by:

$$x(k + 1) = f(x(k), u_d(k), u_a(k), w(k)); \quad x(0) = x_0 \in X; \quad k \in \mathcal{N}, \quad (1)$$

where $\{w(k)\}$ is a sequence of independent and identically distributed (i.i.d.) random variables with zero mean and bounded covariance.

In this paper, we focus on the *Stackelberg setting* with the defender as leader and adversary as follower. The leader selects its inputs anticipating the worst-case response by the adversary. We assume that the adversary can choose its action based on the action of the defender [18], and further, restrict our focus to stationary strategies for the two players. Due to the asymmetry in information available to the players, equilibrium strategies for the case when the game is zero-sum can be chosen to be deterministic strategies [13].

Definition 6 (Defender Strategy). A stationary strategy for the defender is a sequence $\mu^{(d)} := \{\mu_k^{(d)}\}_{k \in \mathcal{N}}$ of Borel-measurable maps $\mu_k^{(d)} : X \rightarrow U_d$.

Definition 7 (Adversary Strategy). A stationary strategy for the adversary is a sequence $\mu^{(a)} := \{\mu_k^{(a)}\}_{k \in \mathcal{N}}$ of Borel-measurable maps $\mu_k^{(a)} : X \times U_d \rightarrow U_a$.

4 Problem Formulation

For a DDSG \mathcal{G} , recall that the labeling function L indicates which atomic propositions are true in each state.

Assumption 1. We restrict our attention to labeling functions of the form $L : X \rightarrow \mathcal{AP}$. Then, if $\mathcal{AP} = (a_1, \dots, a_p)$, \mathcal{AP} and L will partition the state space as $X := \cup_{i=1}^p X_i$, where $X_i := L^{-1}(a_i)$. We further assume that $X_i \neq \emptyset$ for all i .

Remark 1. Through the remainder of the paper, we interchangeably use x_k or $x(k)$ to denote the state at time k .

Given a sequence of states $\mathbf{x}_{\mathcal{N}} := (x_0, x_1, \dots, x_{N-1})$, using Assumption 1, if $\eta_k = L(x_k)$ for all $k \in \mathcal{N}$, then we can write $L(\mathbf{x}_{\mathcal{N}}) = (\eta_0, \eta_1, \dots, \eta_{N-1})$.

Definition 8 (LTL Satisfaction by DDSG). For a DDSG \mathcal{G} and a safe-LTL_F formula ϕ , we write $\mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \phi\}$ to denote the probability that the evolution of the DDSG starting from $x(0) = x_0$ under player policies $\mu^{(d)}$ and $\mu^{(a)}$ satisfies ϕ over the time horizon $\mathcal{N} = \{0, 1, \dots, N-1\}$.

We are now ready to formally state the problem that this paper seeks to solve.

Problem 1. Given a discrete-time dynamic game $\mathcal{G} = (X, W, U_d, U_a, f, \mathcal{N}, \mathcal{AP}, L)$ that evolves according to the dynamics in Eq. (1) and a safe-LTL_F formula ϕ , determine a policy for the defender, $\mu^{(d)}$, that maximizes the probability of satisfying ϕ over the time horizon $\mathcal{N} = \{0, 1, \dots, N-1\}$ under any adversary policy $\mu^{(a)}$ for all $x_0 \in L^{-1}(a_j)$ for some $a_j \in \mathcal{AP}$. That is, compute:

$$\sup_{\mu^{(d)}} \inf_{\mu^{(a)}} \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \phi\} \quad (2)$$

5 Solution Approach

In this section, we present a dynamic programming approach to determine a solution to Problem 1. Our analysis is motivated by the treatment in [18] and [50].

We then introduce the notion of secure control barrier certificates (S-CBCs), and use these to provide a lower bound on the probability of satisfaction of the safe-LTL_F formula ϕ for a defender policy under any adversary policy

in terms of the accepting runs of length less than or equal to the length of the time-horizon of interest of a DFA associated with ϕ . For systems whose evolution of states can be written as a polynomial function of states and inputs, we present a sum-of-squares optimization approach in order to compute an S-CBC.

S-CBCs generalize barrier certificates [22] and control barrier certificates [23] to account for the presence of an adversary. A difference between the treatment in this paper and that of [22, 23] is that we define S-CBCs for stochastic dynamic games, while the latter papers focus on stochastic systems with a single agent.

5.1 Dynamic Programming for *safe* – LTL_F Satisfaction

We introduce a dynamic programming (DP) operator that will allow us to recursively solve a Bellman equation related to Eq. (2) backward in time. First, observe that we can write the satisfaction probability in Definition 8 as:

$$\mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \phi\} = \mathbb{E}_{\mu^{(d)}, \mu^{(a)}} \left\{ \prod_{k \in \mathcal{N}} \mathbf{1}(L(x_k) \models \phi) \mid x(0) = x_0 \right\}, \quad (3)$$

where $\mathbb{E}_{\mu^{(d)}, \mu^{(a)}}$ is the expectation operator under the probability measure $\mathbb{P}_{\mu^{(d)}, \mu^{(a)}}$ induced by agent policies $\mu^{(d)}$ and $\mu^{(a)}$. $\mathbf{1}(\cdot)$ is the indicator function, which takes value 1 if its argument is true, and 0 otherwise.

Assume that $V : X \rightarrow [0, 1]$ is a Borel-measurable function. A DP operator T can then be characterized in the following way:

$$V(x_{N-1}) = \mathbf{1}(L(x_{N-1}) \models \phi) \quad (4)$$

$$(TV)(x_k) := \sup_{u_d} \inf_{u_a} \mathbf{1}(L(x_k) \models \phi) \int_X V(f(x_k, u_d, u_a, w)) dx_{k+1}, \quad (5)$$

where $dx_{k+1} \equiv (dx_{k+1} \mid x_k, u_d, u_a)$ is a probability measure on the Borel space $(X, \mathcal{B}(X))$.

The following results adapts Theorem 1 of [18] to the case of temporal logic formula satisfaction over a finite time-horizon.

Theorem 1. *Assume that the DDSG \mathcal{G} has to satisfy a *safe* – LTL_F formula ϕ over horizon \mathcal{N} . Let the DP operator T be defined as in Eq. (5). Additionally, if $dx_k \equiv (dx_{k+1} \mid x_k, u_d, u_a)$ is continuous, then,*

$$\sup_{\mu^{(d)}} \inf_{\mu^{(a)}} \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \phi\} = (T^N V)(x_0), \quad (6)$$

where $T^N := T \circ T \circ \dots \circ T$ (N times) is the repeated composition of the operator T .

Proof. Consider a particular pair of stationary agent policies $\mu^{(d)}$ and $\mu^{(a)}$. For these policies, define measurable functions $V_k^{\mu^{(d)}, \mu^{(a)}} : X \rightarrow [0, 1]$, $k = 0, 1, \dots, N-1$:

$$V_{N-1}^{\mu^{(d)}, \mu^{(a)}}(x_{N-1}) := \mathbf{1}(L(x_{N-1}) \models \phi) \quad (7)$$

$$V_k^{\mu^{(d)}, \mu^{(a)}}(x_k) := \mathbb{E}_{\mu^{(d)}, \mu^{(a)}} \left\{ \prod_{i=k}^{N-1} \mathbf{1}(L(x_i) \models \phi) \mid x(k) = x_k \right\}, k = 0, 1, \dots, N-2 \quad (8)$$

Therefore, we have $\mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \phi\} = V_0^{\mu^{(d)}, \mu^{(a)}}(x_0)$.

Now, consider strategies of the agents at a stage k . Define the operator $T_{\mu_k^{(d)}, \mu_k^{(a)}}^{(d), (a)}$:

$$(T_{\mu_k^{(d)}, \mu_k^{(a)}}^{(d), (a)} V)(x_k) := \mathbf{1}(L(x_k) \models \phi) \int_X V(f(x_k, u_d, u_a, w)) dx_{k+1} \quad (9)$$

Expanding Eq. (8) using the definition of the expectation operator will allow us to write $V_k^{\mu^{(d)}, \mu^{(a)}}(x) = (T_{\mu_{k+1}^{(d)}, \mu_{k+1}^{(a)}}^{(d), (a)} V)(x)$.

The result follows by an induction argument which uses the fact that $T_{\mu_k^{(d)}, \mu_k^{(a)}}^{(d), (a)}$ is a monotonic operator. We refer to [18] for details. Further, this procedure also guarantees the existence of a defender policy that will maximize the probability of satisfaction of ϕ under any adversary policy. \square

5.2 Secure Control Barrier Certificates

Definition 9. A continuous function $B : X \rightarrow \mathbb{R}_{\geq 0}$ is a secure control barrier certificate (S-CBC) for the DDSG \mathcal{G} if for any state $x \in X$ and some constant $c \geq 0$,

$$\inf_{u_d} \sup_{u_a} \mathbb{E}_w[B(f(x, u_d, u_a, w)) \mid x] \leq B(x) + c. \quad (10)$$

Intuitively, for some defender action u_d , the increase in the value of an S-CBC is bounded from above along trajectories of \mathcal{G} under any adversary action u_a .

Remark 2. S-CBCs generalize control barrier certificates and barrier certificates seen in prior work. If $f(x, u_d, u_{a_1}, w) \sim f(x, u_d, u_{a_2}, w)$ for every $u_{a_1}, u_{a_2} \in U_a$, then we recover the definition of a control barrier certificate [23]. The definition of a barrier certificate [22, 36] is got by additionally requiring that $f(x, u_{d_1}, u_{a_1}, w) \sim f(x, u_{d_2}, u_{a_2}, w)$ for every $u_{d_1}, u_{d_2} \in U_d$ and $u_{a_1}, u_{a_2} \in U_a$. Here \sim denotes stochastic equivalence of the respective stochastic processes [35]. In the latter case, when $c = 0$, the function B is a *super-martingale*. For this case, along with some additional assumptions on the system dynamics, asymptotic guarantees on the satisfaction of properties over the infinite time-horizon can be established [36].

Remark 3. Although our definition of S-CBCs in Definition 9 bears resemblance to the notion of a *worst-case barrier certificate* introduced in [36], there are some distinctions. While the entity in [36] considers a dynamical system with a single

disturbance input, our setting considers three terms that influence the evolution of the state of the system: we want to find a defender input that will allow the barrier function to satisfy a certain property under any adversary input and disturbance. A second point of difference is that while [36] focuses on asymptotic analysis, we consider properties over a finite time horizon.

We limit our attention to stationary strategies for both players. Studying the effects of other strategies is left as future work. The following preliminary result will be used subsequently to determine a bound on the probability of reaching a subset of states under particular agent policies over a finite time-horizon.

Lemma 1. *Consider a DDSG \mathcal{G} and let $B : X \rightarrow \mathbb{R}_{\geq 0}$ be an S -CBC as in Definition 9. Then, for some constants $\lambda > 0$ and $c \geq 0$, initial state $x_0 \in X$, and a stationary defender policy, $\mu^{(d)} : X \rightarrow U_d$, the following holds under any stationary adversary policy $\mu^{(a)} : X \times U_d \rightarrow U_a$:*

$$\inf_{\mu^{(d)}} \sup_{\mu^{(a)}} \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \left[\sup_{0 \leq k < N} B(x(k)) \geq \lambda \right] \leq \frac{B(x_0) + cN}{\lambda} \quad (11)$$

Proof. The proof follows from the result of Chapter III, Theorem 3 and Corollary 2-1 in [27], Definition 9, and the fact that the agents adopt stationary policies. \square

Definition 10 (s -Reachability). *For the DDSG \mathcal{G} with dynamics in Eq. (1), let $s \in [0, 1]$ and $X_0 \subset X$ be the set of possible initial states and $X_1 \subset X$ be disjoint from X_0 . Then, \mathcal{G} is s -reachable with respect to X_1 , if $\sup_{k \in \mathcal{N}} \mathbb{P}[x_k \in X_1] \leq s$. That is, the probability of reaching a state in X_1 starting from X_0 in the time horizon $[0, N]$ is upper bounded by s .*

Theorem 2. *With X_0 and X_1 known, and $X_0 \cap X_1 = \emptyset$, assume there exists an S -CBC $B : X \rightarrow \mathbb{R}_{\geq 0}$, stationary policies, $\mu^{(d)} : X \rightarrow U_d$ and $\mu^{(a)} : X \times U_d \rightarrow U_a$, and constant $c \geq 0$. Additionally, if there is a constant $\delta \in [0, 1]$ such that:*

1. $B(x) \leq \delta$ for all $x \in X_0$,
2. $B(x) > 1$ for all $x \in X_1$,

then the DDSG \mathcal{G} starting from $x_0 \in X_0$ is $(\delta + cN)$ -reachable with respect to X_1 .

Proof. Observe that $X_1 \subseteq \{x \in X : B(x) \geq 1\}$. Therefore, $\mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} [\exists k \in \mathcal{N} : x(k) \in X_1] \leq \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} [B(x(k)) \geq 1]$. Since this should be true for arbitrary k , we have:

$$\begin{aligned} \sup_{k \in \mathcal{N}} \mathbb{P}[x_k \in X_1] &\leq \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \left\{ \sup_{k \in \mathcal{N}} B(x(k)) \geq 1 \right\} \leq \inf_{\mu^{(d)}} \sup_{\mu^{(a)}} \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \left\{ \sup_{k \in \mathcal{N}} B(x(k)) \geq 1 \right\} \\ &\leq B(x_0) + cN \leq \delta + cN \end{aligned}$$

The second line of the above system of inequalities follows by setting $\lambda = 1$ in Lemma 1, and the fact that $B(x) \leq \delta$ for all $x \in X_0$. \square

5.3 Automaton-Based Verification

In order to verify that $\{L(\mathbf{x}_{\mathcal{N}}) \models \phi\}$ under agent policies $\mu^{(d)}$ and $\mu^{(a)}$, we need to establish that $(\eta_0, \eta_1, \dots, \eta_{N-1}) \subseteq \mathcal{L}(\mathcal{A}_{\phi})$. To do this, we first construct a DFA $\mathcal{A}_{\neg\phi}$, that accepts all and only those words over \mathcal{AP} that do not satisfy the *safe* – LTL_F formula ϕ . We have the following result:

Lemma 2. [9] *For $L(\mathbf{x}_{\mathcal{N}}) = (\eta_0, \eta_1, \dots, \eta_{N-1})$ and a DFA \mathcal{A}_{ϕ} , the following is true:*

$$(\eta_0, \eta_1, \dots, \eta_{N-1}) \subseteq \mathcal{L}(\mathcal{A}_{\phi}) \Leftrightarrow (\eta_0, \eta_1, \dots, \eta_{N-1}) \cap \mathcal{L}(\mathcal{A}_{\neg\phi}) = \emptyset$$

The construction of $\mathcal{A}_{\neg\phi}$ can also be carried out in Rabinizer4 [25]. The accepting runs of $\mathcal{A}_{\neg\phi}$ of length less than or equal to N can be computed using a depth-first search algorithm [47]. For the purposes of this section, it is important to understand that the accepting runs of $\mathcal{A}_{\neg\phi}$ of length less than or equal to N will give a bound on the probability that a particular pair of agent policies $(\mu^{(d)}, \mu^{(a)})$ will not satisfy ϕ over the time horizon \mathcal{N} . Using Definition 4 and following the treatment of [22] and [23] define the following terms (the reader is also referred to these works for an example that offers a detailed treatment of the procedure):

$$\mathcal{R}_N(\mathcal{A}_{\neg\phi}) := \{\mathbf{q} = (q_0, \dots, q_n) \in \mathcal{L}(\mathcal{A}_{\neg\phi}) : n \leq N, q_i \neq q_{i+1} \forall i < n\} \quad (12)$$

$$\mathcal{R}_N^a(\mathcal{A}_{\neg\phi}) := \{\mathbf{q} = (q_0, \dots, q_n) \in \mathcal{R}_N(\mathcal{A}_{\neg\phi}) : a \in \mathcal{AP} \text{ and } q_0 \xrightarrow{a} q_1\} \quad (13)$$

$$\mathcal{P}^a(\mathbf{q}) := \begin{cases} \{(q_i, q_{i+1}, q_{i+2}, T(\mathbf{q}, q_{i+1})) : 0 \leq i \leq n-2\} & \mathbf{q} \in \mathcal{R}_N^a(\mathcal{A}_{\neg\phi}), |\mathbf{q}| > 2 \\ \emptyset & \text{otherwise} \end{cases} \quad (14)$$

$$T(\mathbf{q}, q_{i+1}) := \begin{cases} N+2-|\mathbf{q}| & \exists a \in \mathcal{AP} : q_{i+1} \xrightarrow{a} q_{i+1} \\ 1 & \text{otherwise} \end{cases} \quad (15)$$

Intuitively, $\mathcal{R}_N(\mathcal{A}_{\neg\phi})$ is the set of accepting runs in $\mathcal{A}_{\neg\phi}$ of length not greater than N , and without counting any self-loops in the states of the DFA. The set $\mathcal{R}_N^a(\mathcal{A}_{\neg\phi})$ is the set of runs in $\mathcal{R}_N(\mathcal{A}_{\neg\phi})$ with the first state transition labeled by $a \in \mathcal{AP}$. For an element of $\mathcal{R}_N^a(\mathcal{A}_{\neg\phi})$, $\mathcal{P}^a(\mathbf{q})$ defines the set of paths of length 3 augmented with a ‘loop-bound’. The ‘loop-bound’ $T(\mathbf{q}, q_{i+1})$ is an indicator of the number of ‘self-loops’ the run in the DFA can make at state q_{i+1} while still keeping its length less than or equal to N . We assume that $T(\mathbf{q}, q_{i+1}) = 1$ when the run cannot make a self-loop at q_{i+1} .

5.4 Satisfaction Probability Using S-CBCs and $\mathcal{A}_{\neg\phi}$

In this section, we show that an accepting run of $\mathcal{A}_{\neg\phi}$ of length less than or equal to N gives a lower bound on the probability that a particular pair of agent policies will not satisfy the *safe* – LTL_F formula ϕ . We use this in conjunction with the S-CBC to derive an upper bound on the probability that ϕ will be

satisfied for a particular choice of defender policy under any adversary policy. Specifically, we use Theorem 2 over each accepting run of $\mathcal{A}_{\neg\phi}$ of length less than or equal to N to give a bound on the overall satisfaction probability.

Theorem 3. Assume that the DDSG \mathcal{G} has to satisfy a safe- LTL_F formula ϕ over horizon \mathcal{N} . Let $\mathcal{A}_{\neg\phi}$ be the DFA corresponding to the negation of ϕ , and for this DFA, assume that the quantities in Eqs. (12)–(14) have been computed. Then, for some $a_j \in \mathcal{AP}$ and all $x_0 \in L^{-1}(a_j)$ the maximum value of the probability of satisfaction of ϕ for a defender policy $\mu^{(d)}$ under any adversary policy $\mu^{(a)}$ satisfies the following inequality:

$$\sup_{\mu^{(d)}} \inf_{\mu^{(a)}} \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \phi\} \geq 1 - \sum_{\mathbf{q} \in \mathcal{R}_N^{a_j}(\mathcal{A}_{\neg\phi})} \prod_{\rho \in \mathcal{P}^{a_j}(\mathbf{q})} (\delta_\rho + c_\rho T),$$

where $\rho = (q, q', q'', T) \in \mathcal{P}^{a_j}(\mathbf{q})$ is the set of paths of length 3 with loop bound T for $a_j \in \mathcal{AP}$ in an accepting run of length N in $\mathcal{A}_{\neg\phi}$.

Proof. For $a_j \in \mathcal{AP}$, consider $\mathbf{q} \in \mathcal{R}_N^{a_j}(\mathcal{A}_{\neg\phi})$ (Eq. (13)) and the set $\mathcal{P}^{a_j}(\mathbf{q})$ (Eqs. (14) and (15)). Consider an element $\rho = (q, q', q'', T) \in \mathcal{P}^{a_j}(\mathbf{q})$. From Theorem 2, for some stationary defender policy $\mu^{(d)}$, the probability that a trajectory of \mathcal{G} starting from $x_0 \in L^{-1}(\sigma : q \xrightarrow{\sigma} q')$ and reaching $x_1 \in L^{-1}(\sigma : q' \xrightarrow{\sigma} q'')$ under stationary adversary policy $\mu^{(a)}$ over the time horizon T is at most $\delta_\rho + c_\rho T$. Therefore, the probability of an accepting run in $\mathcal{A}_{\neg\phi}$ of length at most N starting from $x_0 \in L^{-1}(a_j)$ is upper bounded by:

$$\inf_{\mu^{(d)}} \sup_{\mu^{(a)}} \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \neg\phi\} \leq \sum_{\mathbf{q} \in \mathcal{R}_N^{a_j}(\mathcal{A}_{\neg\phi})} \prod_{\rho \in \mathcal{P}^{a_j}(\mathbf{q})} (\delta_\rho + c_\rho T)$$

Now consider Eq. (2) of Problem 1. We have the following set of equivalences and inequalities:

$$\begin{aligned} \sup_{\mu^{(d)}} \inf_{\mu^{(a)}} \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \phi\} &= \sup_{\mu^{(d)}} (-\sup_{\mu^{(a)}} (-\mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \phi\})) \\ &= -\inf_{\mu^{(d)}} \sup_{\mu^{(a)}} (-\mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \phi\}) \\ &= -\inf_{\mu^{(d)}} \sup_{\mu^{(a)}} (-1 + \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \neg\phi\}) \\ &\geq 1 - \inf_{\mu^{(d)}} \sup_{\mu^{(a)}} \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}_{\mathcal{N}}) \models \neg\phi\} \\ &\geq 1 - \sum_{\mathbf{q} \in \mathcal{R}_N^{a_j}(\mathcal{A}_{\neg\phi})} \prod_{\rho \in \mathcal{P}^{a_j}(\mathbf{q})} (\delta_\rho + c_\rho T) \end{aligned}$$

□

Theorem 3 generalizes Theorem 5.2 of [23] to provide a lower bound for a stationary defender policy that maximizes the probability that the safe- LTL_F formula is satisfied by the DDSG \mathcal{G} over the time horizon \mathcal{N} , starting from $x_0 \in L^{-1}(a_j)$ for some $a_j \in \mathcal{AP}$ for any stationary adversary policy.

5.5 Computing an S-CBC

The use of barrier functions will circumvent the need to explicitly compute sets of reachable states, which is known to be undecidable for general dynamical systems [29]. However, computationally efficient methods can be used to construct a barrier certificate if the system dynamics can be expressed as a polynomial [36]. This will allow for determining bounds on the probability of satisfaction of the LTL formula without discretizing the state space. In contrast, if the underlying state space is continuous, computing the satisfaction probability and the corresponding agent policy using dynamic programming will necessitate a discretization of the state space in order to approximate the integral in Eq. (5).

We propose a sum-of-squares (SOS) optimization [34] based approach that will allow us to compute an S-CBC if the evolution of the state of the DDSG has a specific structure. The key insight is that if a function can be written as a sum of squares of different polynomials, then it is non-negative.

Assumption 2. *The sets X, U_d, U_a in the DDSG \mathcal{G} are continuous, and $f(x, u_d, u_a, w)$ in Eq. (1) can be written as a polynomial in x, u_d, u_a for any w . Further, the sets $X_i = L^{-1}(a_i)$ in Assumption 1 can be represented by polynomial inequalities.*

Proposition 1. *Under the conditions of Assumption 2, suppose that sets $X_0 := \{x \in X : g_0(x) \geq 0\}$, $X_1 := \{x \in X : g_1(x) \geq 0\}$, and $X := \{x \in X : g(x) \geq 0\}$, where the inequalities are element-wise. Assume that there is an SOS polynomial $B(x)$, constants $\delta \in [0, 1]$ and c , SOS (vector) polynomials $s_0(x), s_1(x)$, and $s(x)$, and polynomials $s_{u_i}^d(x)$ corresponding to the i^{th} entry in u_d , such that:*

$$-B(x) - s_0^\top(x)g_0(x) + \delta \quad (16)$$

$$B(x) - s_1^\top(x)g_1(x) - 1 \quad (17)$$

$$\forall u_a \in U_a : -\mathbb{E}_w[B(f(x, u_d, u_a, w))|x] + B(x) - \sum_i (u_{d_i} - s_{u_i}^d(x)) - s^\top(x)g(x) + c \quad (18)$$

are all SOS polynomials. Then, $B(x)$ satisfies the conditions of Theorem 2, and $u_{d_i} = s_{u_i}^d(x)$ is the corresponding defender policy.

Proof. The proof of this result follows in a manner similar to Lemma 7 in [49] and Lemma 5.6 in [23], and we do not present it here. \square

The authors of [23] discuss an alternative approach in the case when the input set has finite cardinality. A similar treatment is beyond the scope of the present paper, and will be an interesting future direction of research.

6 Example

We present an example demonstrating our solution approach to Problem 1.

Example 1. Let the dynamics of the DDSG \mathcal{G} with $X = W = \mathbb{R}^2$, U_d is a compact subset of \mathbb{R} , $U_a = [-1, 1]$, and $w_1(k), w_2(k) \sim \text{Unif}[-1, 1]$ (and i.i.d.) be given by:

$$x_1(k+1) = -0.5x_1(k)x_2(k) + w_1(k) \quad (19)$$

$$x_2(k+1) = x_1(k)x_2(k) + 0.1x_2^2(k) + u_d(k) + 0.6u_a(k) + w_2(k) \quad (20)$$

Let $\mathcal{AP} = \{a_0, a_1, a_2, a_3, a_4\}$, and sets X_0, X_1, X_2, X_3, X_4 such that for $x \in X_i$, $L(x) = a_i$. The sets X_i are defined by:

$$X_0 := \{(x_1, x_2) : x_1^2 + x_2^2 \leq 0.9\},$$

$$X_1 := \{(x_1, x_2) : (2 \leq x_1 \leq 6) \wedge (-2 \leq x_2 \leq 2)\},$$

$$X_2 := \{(x_1, x_2) : x_1^2 + (x_2 - 10)^2 \leq 4\},$$

$$X_3 := \{(x_1, x_2) : (-10 \leq x_1 \leq -3) \wedge (-4 \leq x_2 \leq -2)\},$$

$$X_4 := X \setminus \bigcup_i X_i.$$

The aim for an agent is to determine a sequence of inputs $\{u_d\}$ such that starting from X_0 , for any sequence of adversary inputs $\{u_a\}$, it avoids obstacles in its environment, defined by the sets X_1, X_2 , and X_3 for 10 units of time. The corresponding *safe* - LTL_F formula is $\phi = [a_0 \wedge \mathbf{G}\neg(a_1 \vee a_2 \vee a_3)]$. The DFA that accepts $\neg\phi$ is shown in Fig. 1. Suppose we are interested in determining a bound on the probability of ϕ being satisfied for a time-horizon of length 10. Using Eqs. (12)–(15), we have $\mathcal{P}^{a_0}(q_0, q_1, q_2) = \{(q_0, q_1, q_2, 9)\}$, and $\mathcal{P}^{a_j} = \emptyset$ for $j = 1, 2, 3, 4$.




Fig. 1. The DFA that accepts $\neg\phi$ for the *safe* - LTL_F formula $\phi = [a_0 \wedge \mathbf{G}\neg(a_1 \vee a_2 \vee a_3)]$ and $\mathcal{AP} = \{a_0, a_1, a_2, a_3, a_4\}$.

We use a sum-of-squares optimization toolbox, SOSTOOLS [37] along with SDPT3 [48], a semidefinite program solver. The barrier function $B(x) = B(x_1, x_2)$ was assumed to be a polynomial of degree-two. For the case $c = 0$, we determine the smallest value of δ that will satisfy the conditions in Proposition 1 to compute an S-CBC. The output of the program was an S-CBC given by

$$B(x) = 0.1915x_1^2 + 0.1868x_1x_2 - 0.144x_1 + 0.1201x_2^2 + 0.1239x_2$$

The environment and the obstacles denoted by the sets X_1, X_2, X_3 and the contours of the S-CBC is shown in Fig. 2. We observe that $B(x)$ is less than 1 in some part of X_1 . A possible reason is that when solving for the second condition in Proposition 1, we work with the union of the sets X_1, X_2 , and X_3 , which may lead to a conservative estimate of the S-CBC.




Fig. 2. The regions X_0, X_1, X_2, X_3, X_4 along with the computed secure control barrier certificate (S-CBC): $B(x) = 0.1915x_1^2 + 0.1868x_1x_2 - 0.144x_1 + 0.1201x_2^2 + 0.1239x_2$. The regions with red boundaries (X_1, X_2, X_3) denote obstacles in the environment. X_0 is the set from which the agent starts at time 0. The contours show the values of the S-CBC of degree 2 ranging from 1 to 100. (Color figure online)

From Theorem 2 and the computed value of δ , we have that

$$\sup_{\mu^{(d)}} \inf_{\mu^{(a)}} \mathbb{P}_{\mu^{(d)}, \mu^{(a)}}^{x_0} \{L(\mathbf{x}, \mathcal{N}) \models \phi\} \geq 0.9922.$$

This bound is conservative in the sense that we consider defender inputs u_d for only the extreme values of $u_a = -1$ and $u_a = 1$. However, for the dynamics in Eq. (20), if the last inequality in Proposition 1 is non-negative for both $u_a = -1$ and $u_a = 1$, then for any $u_a \in [-1, 1]$, this quantity will be non-negative.

Determining methods to explicitly compute a defender policy and considering S-CBCs of higher degree is an area of future research.

7 Conclusion

This paper introduced a new class of barrier certificates to provide probabilistic guarantees on the satisfaction of temporal logic specifications for CPSs that may be affected by the actions of an intelligent adversary. We presented a solution to the problem of maximizing the probability of satisfying a temporal logic specification in the presence of an adversary. The interaction between the CPS and adversary was modeled as a discrete-time dynamic stochastic game with the CPS as defender. The evolution of the state of the game was influenced jointly by the actions of both players. A dynamic programming based approach was used to synthesize a policy for the defender in order to maximize this satisfaction probability under any adversary policy. We introduced secure control barrier certificates, an entity that allowed us to determine a lower bound on the satisfaction probability. The S-CBC was explicitly computed for a certain class of dynamics using sum-of-squares optimization. An example illustrated our approach.

Our example may have resulted in conservative bounds for the satisfaction probabilities since we restrict our focus to barrier certificates that are second degree polynomials and to stationary policies for the two agents. Future work will seek to study conditions under which possibly more effective non-stationary agent policies and higher degree S-CBCs can be deployed to solve the problem. A second interesting problem over a finite time-horizon is to investigate if explicit time bounds can be enforced on the temporal logic formula. An example of such a property is that the agent is required to reach a subset of states of the system between 3 and 5 min. This formula cannot be encoded in LTL, but there are other temporal logic frameworks like metric interval temporal logic [4] or signal temporal logic [31] that will allow us to express it. We propose to study the case when the system will have to satisfy other kinds of timed temporal specifications [12] in the presence of an adversary in dynamic environments.

References

1. Abate, A., Prandini, M., Lygeros, J., Sastry, S.: Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica* **44**(11), 2724–2734 (2008)
2. Ahmadi, M., Jansen, N., Wu, B., Topcu, U.: Control theory meets POMDPs: A hybrid systems approach. arXiv preprint [arXiv:1905.08095](https://arxiv.org/abs/1905.08095) (2019)
3. Ahmadi, M., Wu, B., Lin, H., Topcu, U.: Privacy verification in POMDPs via barrier certificates. In: IEEE Conference on Decision and Control, pp. 5610–5615 (2018)
4. Alur, R., Feder, T., Henzinger, T.A.: The benefits of relaxing punctuality. *J. ACM* **43**(1), 116–146 (1996)
5. Ames, A.D., Coogan, S., Egerstedt, M., Notomista, G., Sreenath, K., Tabuada, P.: Control barrier functions: Theory and applications. In: European Control Conference, pp. 3420–3431 (2019)

6. Ames, A.D., Xu, X., Grizzle, J.W., Tabuada, P.: Control barrier function based quadratic programs for safety critical systems. *IEEE Trans. Autom. Control* **62**(8), 3861–3876 (2016)
7. Başar, T., Olsder, G.J.: *Dynamic Noncooperative Game Theory*, vol. 23. SIAM, Philadelphia (1999)
8. Baheti, R., Gill, H.: Cyber-physical systems. *Impact Control Technol.* **12**(1), 161–166 (2011)
9. Baier, C., Katoen, J.P.: *Principles of Model Checking*. MIT Press, Cambridge (2008)
10. Belta, C., Yordanov, B., Aydin Gol, E.: *Formal Methods for Discrete-Time Dynamical Systems. SSDC*, vol. 89. Springer, Cham (2017). <https://doi.org/10.1007/978-3-319-50763-7>
11. Bertsekas, D.P.: *Dynamic Programming and Optimal Control*, Volumes I and II, 4th edn. Athena Scientific, Nashua (2015)
12. Bouyer, P., Laroussinie, F., Markey, N., Ouaknine, J., Worrell, J.: Timed temporal logics. In: Aceto, L., Bacci, G., Bacci, G., Ingólfssdóttir, A., Legay, A., Mardare, R. (eds.) *Models, Algorithms, Logics and Tools. LNCS*, vol. 10460, pp. 211–230. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63121-9_11
13. Breton, M., Alj, A., Haurie, A.: Sequential Stackelberg equilibria in two-person games. *J. Optim. Theory Appl.* **59**(1), 71–97 (1988)
14. Chow, C.S., Tsitsiklis, J.N.: An optimal one-way multigrid algorithm for discrete-time stochastic control. *IEEE Trans. Autom. Control* **36**(8), 898–914 (1991)
15. Cimatti, A., Clarke, E., Giunchiglia, F., Roveri, M.: NuSMV: A new symbolic model verifier. In: Halbwachs, N., Peled, D. (eds.) *CAV 1999. LNCS*, vol. 1633, pp. 495–499. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48683-6_44
16. De Giacomo, G., Vardi, M.: Synthesis for LTL and LDL on finite traces. *Int. Joint Conf. Artif. Intell.* **15**, 1558–1564 (2015)
17. De Giacomo, G., Vardi, M.Y.: Linear temporal logic and linear dynamic logic on finite traces. In: *International Joint Conference on Artificial Intelligence*, pp. 854–860 (2013)
18. Ding, J., Kamgarpour, M., Summers, S., Abate, A., Lygeros, J., Tomlin, C.: A stochastic games framework for verification and control of discrete time stochastic hybrid systems. *Automatica* **49**(9), 2665–2674 (2013)
19. Ding, X., Smith, S.L., Belta, C., Rus, D.: Optimal control of MDPs with linear temporal logic constraints. *IEEE Trans. Autom. Control* **59**(5), 1244–1257 (2014)
20. Farwell, J.P., Rohozinski, R.: Stuxnet and the future of cyber war. *Survival* **53**(1), 23–40 (2011)
21. Gordon, G.J.: *Approximate solutions to Markov decision processes*. School of Computer Science, Carnegie-Mellon University, Pittsburgh, PA, Technical report (1999)
22. Jagtap, P., Soudjani, S., Zamani, M.: Temporal logic verification of stochastic systems using barrier certificates. In: Lahiri, S.K., Wang, C. (eds.) *ATVA 2018. LNCS*, vol. 11138, pp. 177–193. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01090-4_11
23. Jagtap, P., Soudjani, S., Zamani, M.: Formal synthesis of stochastic systems via control barrier certificates. *arXiv preprint arXiv:1905.04585* (2019)
24. Kolathaya, S., Ames, A.D.: Input-to-state safety with control barrier functions. *Control Syst. Lett.* **3**(1), 108–113 (2018)
25. Křetínský, J., Meggendorfer, T., Sickert, S., Ziegler, C.: Rabinizer 4: From LTL to your favourite deterministic automaton. In: Chockler, H., Weissenbacher, G. (eds.) *CAV 2018. LNCS*, vol. 10981, pp. 567–577. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96145-3_30

26. Kupferman, O., Vardi, M.Y.: Model checking of safety properties. In: Halbwachs, N., Peled, D. (eds.) CAV 1999. LNCS, vol. 1633, pp. 172–183. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48683-6_17
27. Kushner, H.J.: Stochastic Stability and Control. Academic Press, Cambridge (1967)
28. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 585–591. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_47
29. Lafferriere, G., Pappas, G.J., Yovine, S.: Symbolic reachability computation for families of linear vector fields. *J. Symb. Comput.* **32**(3), 231–253 (2001)
30. Lindemann, L., Dimarogonas, D.V.: Control barrier functions for signal temporal logic tasks. *IEEE Control Syst. Lett.* **3**(1), 96–101 (2019)
31. Maler, O., Nickovic, D.: Monitoring temporal properties of continuous signals. In: Lakhnech, Y., Yovine, S. (eds.) FORMATS/FTRTFT -2004. LNCS, vol. 3253, pp. 152–166. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30206-3_12
32. Niu, L., Clark, A.: Secure control under LTL constraints. In: IEEE American Control Conference, pp. 3544–3551 (2018)
33. Niu, L., Li, Z., Clark, A.: LQG reference tracking with safety and reachability guarantees under false data injection attacks. In: IEEE American Control Conference, pp. 2950–2957 (2019)
34. Parrilo, P.A.: Semidefinite programming relaxations for semialgebraic problems. *Math. Programm.* **96**(2), 293–320 (2003)
35. Pola, G., Manes, C., van der Schaft, A.J., Di Benedetto, M.D.: Bisimulation equivalence of discrete-time stochastic linear control systems. *IEEE Trans. Autom. Control* **63**(7), 1897–1912 (2017)
36. Prajna, S., Jadbabaie, A., Pappas, G.J.: A framework for worst-case and stochastic safety verification using barrier certificates. *Trans. Autom. Control* **52**(8), 1415–1428 (2007)
37. Prajna, S., Papachristodoulou, A., Parrilo, P.A.: Introducing SOSTOOLS: A general purpose sum of squares programming solver. In: IEEE Conference on Decision and Control, vol. 1, pp. 741–746 (2002)
38. Puterman, M.L.: Markov Decision Processes: Discrete Stochastic Dynamic Programming. Wiley, Hoboken (2014)
39. Ramasubramanian, B., Clark, A., Bushnell, L., Poovendran, R.: Secure control under partial observability under temporal logic constraints. In: IEEE American Control Conference, pp. 1181–1188 (2019)
40. Saha, I., Ramaithitima, R., Kumar, V., Pappas, G.J., Seshia, S.A.: Automated composition of motion primitives for multi-robot systems from safe LTL specifications. In: Proceedings of International Conference on Intelligent Robots and Systems, pp. 1525–1532 (2014)
41. Santoyo, C., Dutreix, M., Coogan, S.: Verification and control for finite-time safety of stochastic systems via barrier functions. In: IEEE Conference on Control Technology and Applications (2019)
42. Sharan, R., Burdick, J.: Finite state control of POMDPs with LTL specifications. In: IEEE American Control Conference, pp. 501–508 (2014)
43. Shoukry, Y., Martin, P., Tabuada, P., Srivastava, M.: Non-invasive spoofing attacks for anti-lock braking systems. In: Bertoni, G., Coron, J.-S. (eds.) CHES 2013. LNCS, vol. 8086, pp. 55–72. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40349-1_4

44. Slay, J., Miller, M.: Lessons learned from the maroochy water breach. In: Goetz, E., Sheno, S. (eds.) ICCIP 2007. IIFIP, vol. 253, pp. 73–82. Springer, Boston, MA (2008). https://doi.org/10.1007/978-0-387-75462-8_6
45. Steinhardt, J., Tedrake, R.: Finite-time regional verification of stochastic non-linear systems. *Int. J. Robot. Res.* **31**(7), 901–923 (2012)
46. Sullivan, J.E., Kamensky, D.: How cyber-attacks in Ukraine show the vulnerability of the US power grid. *Electr. J.* **30**(3), 30–35 (2017)
47. Tarjan, R.: Depth-first search and linear graph algorithms. *SIAM J. Comput.* **1**(2), 146–160 (1972)
48. Toh, K.C., Todd, M.J., Tütüncü, R.H.: SDPT3: A MATLAB software package for semidefinite programming. *Optim. Methods Softw.* **11**, 545–581 (1999)
49. Wongpiromsarn, T., Topcu, U., Lamperski, A.: Automata theory meets barrier certificates: Temporal logic verification of nonlinear systems. *IEEE Trans. Autom. Control* **61**(11), 3344–3355 (2015)
50. Yang, I.: A dynamic game approach to distributionally robust safety specifications for stochastic systems. *Automatica* **94**, 94–101 (2018)