

Verification of Hyperproperties for Dynamical Systems via Barrier Certificates

Mahathi Anand, *Member, IEEE*, Vishnu Murali, *Student Member, IEEE*, Ashutosh Trivedi, and Majid Zamani, *Senior Member, IEEE*

Abstract—Hyperproperties are system properties that require quantification over multiple execution traces of a system. Hyperproperties can express several specifications of interest for cyber-physical systems—such as opacity, robustness, and noninterference—which cannot be expressed using linear time properties. This paper presents for the first time a discretization-free approach for the formal verification of discrete-time dynamical systems against hyperproperties. The proposed approach involves the decomposition of complex hyperproperties into several verification conditions by exploiting the automata-based structures corresponding to the complements of the original specifications. These verification conditions are then discharged by synthesizing so-called *augmented barrier certificates*, which provide certain safety guarantees for the underlying system. For systems with polynomial-type dynamics, we present a sound procedure to synthesize polynomial-type augmented barrier certificates by reducing the problem to sum-of-squares optimizations. We demonstrate the effectiveness of our proposed approaches on two physical case studies against two important hyperproperties: *initial-state opacity* and *initial-state robustness*.

I. INTRODUCTION

Classical control theory provides theory, techniques, and tools to analyze *complex* dynamical systems against *simple* objectives such as stability. Traditional formal methods, on the other hand, tend to focus on developing approaches to verify software/hardware systems with *simpler* discrete dynamics against structurally *rich* logic-based specifications concerning safety and liveness. Cyber-physical systems (CPS)—systems characterized by structured software systems interacting with equally complex physical systems—blur this traditional demarcation between control systems and software systems. The critical role of CPS in modern society in safety- and security-critical applications has spurred interest in developing approaches to provide rigorous guarantees for such systems. As a result, in the last two decades, formal verification of CPS with complex continuous state-space dynamics against logic-based specifications has received considerable attention [1], [2].

This work was supported in part by the German Research Foundation (DFG) through the Research Training Group 2428 and the NSF under grant ECCS-2015403.

M. Anand is with the Institute of Systems Theory and Automatic Control, University of Stuttgart, Germany. V. Murali, A. Trivedi, and M. Zamani are with the Computer Science Department, University of Colorado Boulder, USA. M. Zamani is also with the Computer Science Department, LMU Munich, Germany. Emails: mahathi.anand@ist.uni-stuttgart.de, vishnu.murali@colorado.edu, ashutosh.trivedi@colorado.edu, majid.zamani@colorado.edu.

Verification of CPS has been typically performed against temporal logic specifications, notably linear temporal logic (LTL), signal temporal logic (STL) [3], and computation tree logic (CTL) [4], expressing properties of a set of desirable system executions. While these logics can describe a large number of specifications of interest that consider individual execution traces of CPS, many important information-flow properties and planning objectives involve relating multiple execution traces. These properties cannot be expressed by classical temporal logic specifications equipped to express properties of individual traces. To fill this gap, Clarkson and Schneider [5] introduced the notion of *hyperproperties* as properties of collective behavior relating multiple execution traces.

As an example of a hyperproperty, consider a security property in a CPS prone to *intrusion attacks*. Suppose that this system requires that secret information is never revealed, *i.e.*, observations from the outside remain indistinguishable from each other, despite the secret. This specification, known as opacity [6], requires us to relate and quantify two observation traces simultaneously. Similarly, an optimality objective [7] for a robotic system would require the existence of a trace that is more favorable than all the other traces of the system, again quantifying multiple execution traces at a time. Other examples for hyperproperties include noninterference [8] and observational determinism [9]. In order to formally specify hyperproperties, the hyper-temporal logic HyperLTL was introduced in [10]. HyperLTL, developed as an extension to LTL, uses trace variables to denote individual execution traces and utilizes universal (\forall) and existential (\exists) quantifiers before a quantifier-free formula over atomic propositions to specify on which traces the atomic propositions must hold.

Formal verification of hyperproperties has been studied in the context of finite-state systems. However, the techniques used for finite systems are not applicable to real-world CPS which evolve on continuous (or even hybrid) state spaces.

Contributions. We aim at bridging this gap by presenting a discretization-free, systematic, and sound verification procedure based on a notion of barrier certificates for discrete-time dynamical systems against hyperproperties. In particular, we consider those specifications that can be expressed by HyperLTL formulae [10]. The verification procedure is achieved by decomposing the given specification into simpler safety tasks, so-called *conditional invariance*, by constructing an implicitly quantified Büchi automaton corresponding to the complement

of the specification. We introduce *augmented barrier certificates* (ABCs), defined over an augmented system obtained by taking the product of the original system with itself (self-composition), which provide us with sufficient conditions ensuring the satisfaction of those conditional invariances.

Inspired by the results in [11], we propose an automata-theoretic approach to extend the applicability of ABCs beyond conditional invariance to HyperLTL specifications by finding barrier certificates ensuring the possibility of avoiding accepting traces of corresponding automata by disallowing certain transitions on different lassos (simple path followed by a simple accepting cycle). To do so, “existential” player is required to select a trace before knowing the choices of the “universal” player. This necessitates the need for a common ABC for some transitions of all lassos, which may be hard to ensure in practice. On the other hand, when the HyperLTL property belongs to $\forall^*\exists^*$ -fragment [10], we can exploit a similar approach as in [11] to use separate ABCs to provide the necessary guarantees by leveraging the structure of the automata corresponding to the negation of specifications. For systems with polynomial-type dynamics, we present a sum-of-squares (SOS) approach to compute polynomial-type ABCs for the individual conditional invariance. Finally, we demonstrate the effectiveness of our proposed approach by verifying two physical case studies with respect to initial-state opacity and initial-state robustness, respectively, which both can be described by HyperLTL formulae. Proofs of all statements are provided in the Appendix.

Related Literature. There have been several results in the literature for the verification and synthesis of CPS against temporal logic specifications. Many earlier results have utilized abstraction-based techniques based on state-space discretization. Examples include abstraction-based framework for linear systems [12], for nonlinear systems [13], synthesizing feedback strategies for piece-wise affine systems [14], and counterexample-guided abstraction refinement (CEGAR) for nonlinear systems [15] to name a few. More recently, automata-theoretic, discretization-free approaches via barrier certificates [16] have been utilized for the verification of LTL specifications in the context of nonlinear systems [11], hybrid systems [17], as well as stochastic systems [18], [19]. We should add that barrier certificates have also been used to verify STL specifications [20], [21].

Unfortunately, most of the existing results pertaining to hyperproperties are tailored to finite-state transition systems. For example, the results in [22] present a practical verification approach for finite-state systems with respect to alternation-free fragments of HyperLTL formulae. The proposed approaches in [10] present a model-checker for HyperLTL specifications with alternation depth of at most one. The results in [23] propose a new model checking algorithm based on model-counting for quantitative hyperproperties. A bounded model checking algorithm for hyperproperties is proposed in [24]. Verification of other types of hyperproperties such as k -safety hyperproperties and hyperliveness properties have also been studied in [25] and [26], respectively. Checking satisfiability of certain fragments of HyperLTL specifications, such as the

“ $\forall^*\exists^*$ ” fragment, are undecidable in general [27]. Formal verification of continuous state-space CPS against general hyperproperties remains largely unexplored. Hyperproperties have been studied for CPS in [28] as well as [29], but in the context of falsification and statistical model checking, respectively. These results are empirical and rely on experimental simulations, and therefore do not provide sound guarantees. Finally, we would like to mention that an extended abstract of this work was presented in [30].

II. PROBLEM DEFINITION

We write \mathbb{R} and \mathbb{N} to denote the set of real and non-negative integers, respectively. Appropriate subscripts are used to restrict the sets, e.g., $\mathbb{R}_{>0} = \{x \in \mathbb{R} : x > 0\}$ denotes the set of positive reals. We write \mathbb{R}^n to denote n -dimensional Euclidean space equipped with Euclidean norm $\|x\|$. For a finite set A , the cardinality of A is denoted by $|A|$.

For a family $x_1 \in \mathbb{R}^{n_1}, x_2 \in \mathbb{R}^{n_2}, \dots, x_N \in \mathbb{R}^{n_N}$ of N vectors, we write (x_1, x_2, \dots, x_N) to denote the corresponding vector of dimension $\sum_i n_i$. For a set A , we write A^n for the n -ary Cartesian power of A , i.e. $A^n = \{(a_1, a_2, \dots, a_n) \mid a_i \in A \text{ for all } 1 \leq i \leq n\}$. For a tuple $t = (a_1, a_2, \dots, a_k)$ and $1 \leq i \leq k$ we write $t(i)$ for its i -th element and $t_{\leq i}$ for the tuple (a_1, a_2, \dots, a_i) . Similarly, for a mapping $f : A \rightarrow B$, we define its n -ary Cartesian power $f^n : A^n \rightarrow B^n$, as $(a_1, \dots, a_n) \mapsto (f(a_1), \dots, f(a_n))$.

An alphabet Σ is a finite set of letters. An ω -sequence $\sigma = w_0 w_1 \dots$ is an infinite concatenation of letters, i.e. for all $i \geq 0$ we have $w_i \in \Sigma$. A finite sequence is such a sequence but with a finite length. We write Σ^* and Σ^ω for the set of finite and ω -sequences over Σ , and we let $\Sigma^\infty = \Sigma^* \cup \Sigma^\omega$. For a sequence $\sigma = w_0 w_1 \dots \in \Sigma^\omega$, let $\sigma[i]$ be the i -th element w_i of σ and $\sigma[i, \infty]$ for the ω -sequence $w_i w_{i+1} \dots \in \Sigma^\omega$ of σ starting from i -th position.

We let $\text{zip} : (\Sigma^\omega)^p \rightarrow (\Sigma^p)^\omega$ denote a function that maps a p -tuple of sequences to a sequence of p -tuples, i.e.

$$(\sigma_1, \sigma_2, \dots, \sigma_p) \rightarrow$$

$$(\sigma_1[0], \sigma_2[0], \dots, \sigma_p[0])(\sigma_1[1], \sigma_2[1], \dots, \sigma_p[1]), \dots,$$

and $\text{unzip} : (\Sigma^p)^\omega \rightarrow (\Sigma^\omega)^p$ denotes the inverse of zip , i.e.

$$\sigma \rightarrow ((\sigma[0](1)\sigma1\dots), \dots, (\sigma[0](p)\sigma[1](p)\dots)),$$

where $\sigma[i](j)$ refers to the j th element of the i th p -tuple in the infinite sequence σ .

A. Discrete-Time Dynamical Systems

We consider a discrete-time dynamical system (or simply, system) defined by a tuple $\mathfrak{S} = (X, W, f)$, where $X \subseteq \mathbb{R}^n$ and $W \subseteq \mathbb{R}^m$ are the (potentially uncountable) state and exogenous input sets, and $f : X \times W \rightarrow X$ is the transition function that characterizes the state evolution. The evolution or *run* of the system \mathfrak{S} for a given initial state $x_0 \in X$ and exogenous input sequence $\nu : \mathbb{N} \rightarrow W$, denoted by $\mathbf{x}_{(x_0, \nu)}$, is given by sequence $\mathbf{x} : \mathbb{N} \rightarrow X$, where

$$\mathbf{x}(t) = \begin{cases} x_0 & \text{if } t = 0 \\ f(\mathbf{x}(t-1), \nu(t-1)) & \text{otherwise.} \end{cases} \quad (\text{II.1})$$

For a system \mathfrak{S} , we define its p -fold self-composition as p -fold augmented system $\mathfrak{S}^p = (X^p, W^p, f^p)$ where X^p and W^p are p -ary Cartesian powers of X , and W , respectively, and $f^p : X^p \times W^p \rightarrow X^p$ is equivalent to the p -ary Cartesian power of f , i.e. $f^p : (X \times W)^p \rightarrow X^p$ by using the zip function. We use these two types interchangeably. We use \tilde{x} and \tilde{w} to denote the state and exogenous input of the augmented system \mathfrak{S}^p , respectively. Similarly, we write $\tilde{x}_{(\tilde{x}_0, \tilde{\nu})}$ for the *run* of \mathfrak{S}^p starting from an initial state $\tilde{x}_0 \in X^p$ and under exogenous input sequence $\tilde{\nu} : \mathbb{N} \rightarrow W^p$.

B. Barrier Certificates

A function $\mathcal{B} : X \rightarrow \mathbb{R}$ is a *barrier certificate* for a system $\mathfrak{S} = (X, W, f)$ from the set $A \subseteq X$ to the set $B \subseteq X$, if:

$$\mathcal{B}(x) \leq 0, \quad \text{for all } x \in A, \quad (\text{II.2})$$

$$\mathcal{B}(x) > 0, \quad \text{for all } x \in B, \quad (\text{II.3})$$

and for all $x \in X$ and for all $w \in W$:

$$\mathcal{B}(f(x, w)) - \mathcal{B}(x) \leq 0. \quad (\text{II.4})$$

The existence of a barrier certificate is a sufficient condition to guarantee that if the system \mathfrak{S} ever visits a state from the set A , it will never visit a state from the set B in the future. If the set A characterizes initial states and the set B characterizes the set of bad or undesirable states, then the existence of barrier certificates from A to B can guarantee safety, see [16].

C. LTL Specifications

Consider a set of atomic propositions AP relevant to the underlying system and the alphabet $\Sigma = 2^{\text{AP}}$ characterized by the subsets of these propositions. We refer to an infinite sequence (ω -sequence) of letters from Σ as an infinite trace. We write Σ^ω for the set of all infinite traces over Σ .

Syntax. An LTL formula over AP can be built from the following production rules:

$$\psi ::= a \mid \neg\psi \mid \psi \vee \psi \mid X\psi \mid \psi \mathbf{U} \psi,$$

where $a \in \text{AP}$, X and \mathbf{U} are the **next** and **until** operators, respectively. Other popular temporal operators such as **globally** (G), **eventually** (F) and **release** (R) can be derived from these minimal ones in a standard manner.

Semantics. Given an infinite trace σ and an LTL formula ψ , the formula ψ is valid for σ , i.e. $\sigma \models \psi$, if:

- $\psi = a$ and $a \in \sigma(0)$,
- $\psi = \neg\psi$ and $\sigma \not\models \psi$,
- $\psi = \psi_1 \vee \psi_2$ and $\sigma \models \psi_1$ or $\sigma \models \psi_2$,
- $\psi = X\psi$ and $\sigma[1, \infty] \models \psi$,
- $\psi = \psi_1 \mathbf{U} \psi_2$ and $\sigma[i, \infty] \models \psi_2$ for some $i \geq 0$ and for all $0 \leq j < i$, we have that $\sigma[j, \infty] \models \psi_1$.

We refer the interested readers to [4] for more details on syntax and semantics of LTL properties. The LTL specifications can only express trace properties, i.e., properties of individual execution traces. However, they cannot specify properties over sets of execution traces, which is essential for many relevant security specifications.

D. HyperLTL Specifications

HyperLTL, unlike LTL which implicitly considers only a single trace at a time, can relate multiple traces simultaneously through the use of existential and universal quantifiers.

Syntax. We consider HyperLTL with syntax:

$$\phi ::= \exists \pi. \phi \mid \forall \pi. \phi \mid \psi$$

$$\psi ::= a_\pi \mid \neg\psi \mid \psi \vee \psi \mid X\psi \mid \psi \mathbf{U} \psi.$$

The key distinction over LTL formulae is the introduction of trace quantifiers \exists and \forall . The quantifier $\exists \pi$ stands for “for some trace π ” while the quantifier $\forall \pi$ stands for “for all traces π ”, respectively. The variable ψ generates standard LTL formulae with the exception that atomic propositions can refer to distinct trace variables. Hence, for every proposition $a \in \text{AP}$ and trace variable π , we use a_π to express that proposition a is referring to the trace π . A trace variable occurs free in a HyperLTL formula, if it is not bounded by any trace quantifier, i.e., if the trace variable is quantified by neither \exists nor \forall quantifiers. A HyperLTL formula with no free variable is called closed.

Semantics. Since HyperLTL formulae express the properties of multiple trace variables, one requires to assign these trace variables to specific traces for reasoning about the satisfaction of the formula. Let $\mathcal{V} = \{\pi_1, \pi_2, \dots\}$ be an infinite set of trace variables. The semantics of a HyperLTL formula ψ is defined over a set T of traces and a trace valuation function $\Pi : \mathcal{V} \rightarrow \Sigma^\omega$ that maps all the free trace variables occurring in the formula ψ to traces in the set Σ^ω . We use $\Pi[\pi \rightarrow \sigma]$ to express the trace valuation function Π' that agrees with Π for all trace variables except π and $\Pi'(\pi) = \sigma$. We define the trace valuation suffix $\Pi[i, \infty]$ as $\pi \mapsto \Pi(\pi)[i, \infty]$, i.e. $\Pi[i, \infty]$ maps π to the i -suffix of the trace mapped to π by Π .

We say that a HyperLTL formula ψ is satisfiable over a given set T of traces and trace valuation function $\Pi : \mathcal{V} \rightarrow \Sigma^\omega$, and we write $\Pi \models_T \psi$ if one of the following holds:

- $\phi = \exists \pi. \psi$ and there is $\sigma \in T$ such that $\Pi[\pi \rightarrow \sigma] \models_T \psi$,
- $\phi = \forall \pi. \psi$ and for all $\sigma \in T$, we have $\Pi[\pi \rightarrow \sigma] \models_T \psi$,
- $\phi = a_\pi$ and $a \in \Pi(\pi)(0)$,
- $\phi = \neg\psi$ and $\Pi \not\models_T \phi$,
- $\phi = \psi_1 \vee \psi_2$ and $\Pi \models_T \psi_1$ or $\Pi \models_T \psi_2$,
- $\phi = X\psi$ and $\Pi[1, \infty] \models_T \psi$,
- $\phi = \psi_1 \mathbf{U} \psi_2$ and there is $i \geq 0$ such that $\Pi[i, \infty] \models_T \psi_2$ and for all $0 \leq j < i$, we have that $\Pi[j, \infty] \models_T \psi_1$.

A closed HyperLTL formula ϕ is considered to be valid for a set of traces T , and we write $T \models \phi$ if the empty trace assignment satisfies the formula, i.e., $\emptyset \models_T \phi$. We refer interested readers to [10] for more details on syntax and semantics of HyperLTL properties.

Remark 2.1: It might be of interest to the readers to know the differences between computational tree logic (CTL) as well as HyperLTL specifications, due to the use of quantifiers also in the former. It must be noted that while CTL does indeed have explicit trace quantifiers, it cannot directly express hyperproperties since the atomic propositions do not refer to specific traces. More details on the comparison can be found in [5].

E. Büchi Automata

A Büchi automaton is a tuple $\mathcal{A} = (\Sigma, Q, q_0, \Delta, F)$, where Σ is a finite set of *alphabet*, Q is a finite set of *states*, $q_0 \in Q$ is the *initial state*, $\Delta \subseteq Q \times \Sigma \times Q$ is the transition relation and $F \subseteq Q$ is the set of *accepting states*. We say that a Büchi automaton is *deterministic* if Δ can be represented as a function that maps a state and alphabet pair to exactly one state, i.e., $\Delta : Q \times \Sigma \rightarrow Q$, and nondeterministic otherwise. We abbreviate deterministic and nondeterministic Büchi automata as DBA and NBA respectively. A *run* r of \mathcal{A} on $w \in \Sigma^\omega$ is an infinite sequence $r_0 w_0 r_1 w_1 \dots \in (Q \times \Sigma)^\omega$ such that $r_0 = q_0$ and, for $i > 0$, $(r_{i-1}, w_{i-1}, r_i) \in \Delta$. We write $\inf(r)$ for the set of states that appear infinitely often in the run r . We say that a trace w is accepted by \mathcal{A} , if, for the corresponding run r , we have $\inf(r) \cap F \neq \emptyset$. We call this run r the accepting run of \mathcal{A} . The *language*, $\mathcal{L}_{\mathcal{A}}$, of any automaton \mathcal{A} (or, *recognized* by \mathcal{A}) is the subset of traces in Σ^ω that have accepting runs in \mathcal{A} . A language is ω -*regular* if it is accepted by an NBA.

It is well known that for any LTL specification, one may compile a corresponding NBA [31] by using LTL to Büchi construction techniques [32]. However, in the case of HyperLTL, the presence of quantification of traces in the specification prevents an immediate compilation to an automaton. Nevertheless, given a HyperLTL specification $\phi = \mu_1 \pi_1 \dots \mu_p \pi_p \psi$ defined over a set of atomic propositions AP, where for all $1 \leq i \leq p$, $\mu_i \in \{\exists, \forall\}$, one may construct an NBA \mathcal{A}_ψ corresponding to the quantifier-free LTL formula ψ over the alphabet Σ^p .

In order to accommodate for the presence of quantification in the HyperLTL formula, a trace for this automaton is a p -tuple of individual traces, denoted by $\tilde{\sigma} = (\sigma_1, \dots, \sigma_p)$, i.e., $\text{unzip}(\tilde{\sigma}) \in \mathcal{L}_{\mathcal{A}_\psi}$ if $(\pi_1 \mapsto \sigma_1, \dots, \pi_p \mapsto \sigma_p)$ is accepted by \mathcal{A}_ψ .

F. HyperLTL Verification Problem

Let $\mathfrak{S} = (X, W, f)$ be a system and $L : X \rightarrow \Sigma$ be a measurable function labeling states of system \mathfrak{S} with letters of the alphabet. We can extend the labeling function from states to runs of \mathfrak{S} in a straightforward fashion: for every infinite state run $x = (x(0), x(1), \dots)$ of \mathfrak{S} , the function $L(x)$ defines the corresponding *trace* of the system $(\sigma_0, \sigma_1, \dots) \in \Sigma^\omega$, where $\sigma_i = L(x(i))$, for all $i \in \mathbb{N}$. Let $T(\mathfrak{S}, L)$ be the set of all traces of \mathfrak{S} . We now define the key problem of interest.

Problem 2.2 (HyperLTL Verification): Given a discrete-time dynamical system $\mathfrak{S} = (X, W, f)$, a labeling function $L : X \rightarrow \Sigma$, and a HyperLTL specification ϕ , the HyperLTL verification problem $\mathfrak{P} = (\mathfrak{S}, L, \phi)$ is to decide whether $T(\mathfrak{S}, L) \models \phi$.

While model checking of finite-state systems against HyperLTL specifications is decidable [22], the verification problem stated above is in general undecidable for continuous state-space systems considered in this paper. It follows readily from the fact that even simple reachability is undecidable for simple continuous state-space dynamical systems [33]. Our approach provides a sound procedure for Problem 2.2. Moreover, to better illustrate our approach, we use the following case study as a running example that we use throughout the article.

Example 2.3: Here, we consider the discrete-time evolution of the temperature $T(\cdot)$ in a room in the presence of a safety controller (as designed in [34]) given by

$$\begin{aligned} \mathfrak{S} : T(k+1) = & T(k) + \tau_s \alpha_e (T_e - T(k)) \\ & + \tau_s \alpha_h (T_h - T(k)) (c_1 T(k) + c_2), \end{aligned} \quad (\text{II.5})$$

where parameters $\alpha_e = 0.008$ and $\alpha_h = 0.0036$ are heat exchange coefficients, $T_e = 15^\circ\text{C}$ is the ambient temperature, $T_h = 55^\circ\text{C}$ is the heater temperature, $c_1 = -0.0024$ and $c_2 = 0.5357$ are controller parameters, and $\tau_s = 5$ minutes is the sampling time. We want to verify the initial-state robustness property, which requires that if a state run starting from a given initial condition remains safe, then all the state runs starting from δ -close initial conditions must also remain safe. Such a specification is especially useful when there are uncertainties arising from not knowing the exact initial state. Note that robustness is a commonly studied property in the classical control theory [35]. However, we provide here an alternative method to verify robustness by formulating it as a HyperLTL formula.

To describe our specification as a HyperLTL formula, we consider the state set $X = [20, 35]$. We further introduce the safe set as $X_1 = [20, 25]$ and the unsafe set as $X_2 = [25, 35]$. For the system \mathfrak{S} , the predefined initial state is given by $X_3 = \{21\}$. We also define the set $X_4 = [20.5, 21.5]$ to capture δ -close states with respect to the initial state, where $\delta = 0.5$. The set of atomic propositions is $AP = \{a_1, a_2, a_3, a_4\}$, where $L(x \in X_i) = a_i$, for all $i \in \{1, 2, 3, 4\}$. The HyperLTL formula for initial-state robustness specification is $\phi = \forall \pi_1 \forall \pi_2 (a_{3\pi_1} \wedge a_{4\pi_2}) \rightarrow G(a_{1\pi_1} \wedge a_{1\pi_2})$.

III. AUGMENTED BARRIER CERTIFICATES (ABCs)

Verification of HyperLTL formulae in the context of finite systems has been well studied [10], [22]. The verification procedure is based on automata-theoretic model-checking, where quantifier-free fragments of the desired HyperLTL formulae are compiled into ω -automata, and trace quantification is handled by appropriately composing these automata with the underlying Kripke Structure. The interleavings of automata and Kripke products lead to an automata whose language emptiness decides the satisfaction of the specification. Unfortunately, this approach cannot be extended to continuous state-space systems by simply using abstraction-based techniques as system relations (e.g., simulation relations) may not preserve hyperproperties [36].

To verify the specification ϕ against system \mathfrak{S} , we compile the negation of the HyperLTL formula into an implicitly quantified Büchi automata, and discharge the unsatisfiability of the specification to synthesizing appropriate barrier certificates. Our verification procedure is depicted in Figure 1. Here, given a HyperLTL specification ϕ and system \mathfrak{S} , we construct the p -fold augmented system \mathfrak{S}^p , and the NBA for $\neg\psi$. We then find an augmented barrier certificate (ABC) that acts as a proof certificate of conditional invariance (cf. Definition 3.2) for some transitions along every lasso. This acts as a “scissor” and allows us to conclude that the system \mathfrak{S} satisfies the specification ϕ . We note that our approach is not complete

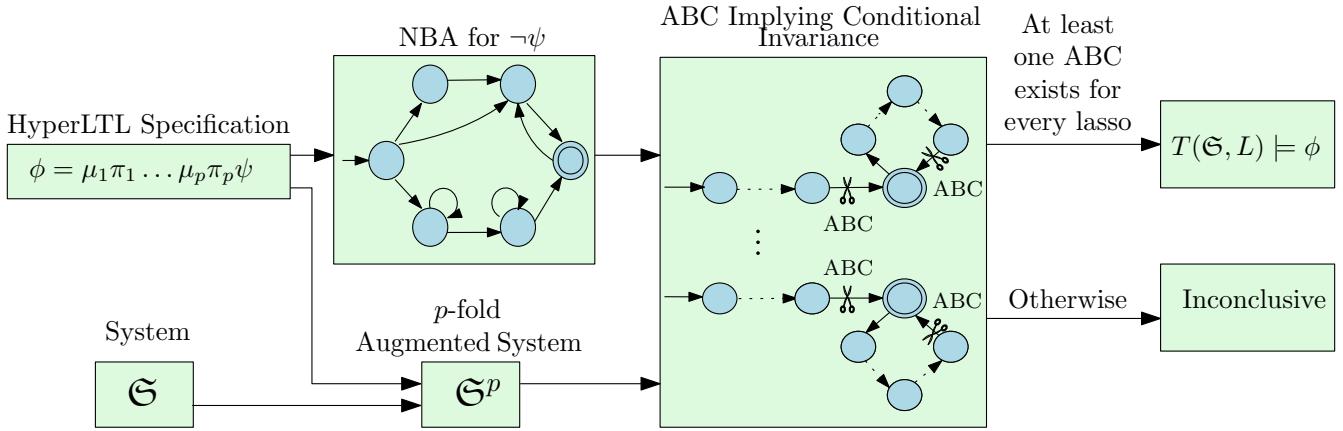


Fig. 1: A schematic block diagram illustrating the verification procedure.

in that if it cannot find an ABC for at least one transition along every lasso then we cannot conclude that the system does not satisfy ϕ . The rest of the section introduces the idea of augmented barrier certificates and provides an automata-theoretic sound verification approach for Problem 2.2.

A. HyperLTL Evaluation Game Semantics

We provide a game semantics to the HyperLTL verification problem $\mathfrak{P} = (\mathfrak{S}, L, \phi)$ as a two-player *stage-based evaluation game* played between two players, Eloise (\exists) and Abelard (\forall), where Eloise takes the role of a verifier and her goal is to prove that $T(\mathfrak{S}, L) \models \phi$, while the goal of Abelard (the spoiler) is the opposite. Given a HyperLTL formula $\mu_1\pi_1 \dots \mu_p\pi_p\psi$, we say that Eloise controls the quantifier μ_i if $\mu_i = \exists$, otherwise we say that Abelard controls the quantifier. The game continues in stages. In the first stage, the game begins with a token in the initial position $(\mu_1\pi_1\mu_2\pi_2 \dots \mu_p\pi_p\psi, \Pi = \emptyset)$ [37] and the player controlling the left-most quantifier μ_1 chooses a trace σ_1 from $T(\mathfrak{S}, L)$ and moves the token to the next position $(\mu_2\pi_2 \dots \mu_p\pi_p\psi, \Pi = \{\pi_1 \rightarrow \sigma_1\})$. The game from the next position continues in a similar fashion until we reach a position with a quantifier-free HyperLTL formula. We call such positions *terminal*. We say that a terminal position $(\psi, \Pi = \{\pi_1 \rightarrow \sigma_1, \dots, \pi_p \rightarrow \sigma_p\})$ is winning for Eloise if $\text{zip}(\sigma_1, \dots, \sigma_p) \models \psi$ with the standard LTL semantics for the formula ψ . We say that Eloise wins the multi-stage evaluation game if she has a way to choose her moves such that no matter how Abelard chooses his moves the game ends in a winning position for Eloise; otherwise Abelard wins the game. Notice that at every step, both players have access to complete infinite traces that have been chosen by players in earlier positions as a part of position description. It is trivial to see that Eloise wins the evaluation game if and only if $\Pi \models_T \phi$.

We consider another version of evaluation games that we dub *turn-based evaluation games*. These games are played on the p -fold self-composition $\mathfrak{S}^p = (X^p, W^p, f^p)$ of the system \mathfrak{S} . These games start with a token in some initial configuration $(x_1^0, x_2^0, \dots, x_p^0) \in X^p$ and at every round first the player controlling μ_1 chooses an action $w_1 \in W$, followed by the player controlling μ_2 , and so on. In this way, the players form

a set of joint actions $(w_1, w_2, \dots, w_p) \in W^p$ and the token is moved to a state $f^p((x_1^0, x_2^0, \dots, x_p^0), (w_1, w_2, \dots, w_p))$. The game continues in this fashion indefinitely and the players thus form an infinite run \tilde{x} of $\mathfrak{S}^p = (X^p, W^p, f^p)$. To relate the augmented system \mathfrak{S}^p with letters in Σ^p , we extend the definition of the labeling function to the augmented system domain by using $L^p : X^p \rightarrow \Sigma^p$ to map the states in the augmented state set to the alphabet Σ^p . We say that the trace \tilde{x} is winning for Eloise if $L^p(\tilde{x}) \models \psi$. We say that Eloise has a winning strategy in the turn-based evaluation game for $\mathfrak{P} = (\mathfrak{S}, L, \phi)$ if she can select her moves in such a way that no matter how Abelard chooses his moves (including using an arbitrary look-ahead), the resulting trace is winning for Eloise. Moreover, we say that Eloise has a positional winning strategy if to select actions in a given round her choice depends only on the current states and choices resolved before her turn for the other quantifiers in the current round.

Lemma 3.1: If Eloise has a positional winning strategy in a turn-based evaluation game for \mathfrak{P} then she has a winning strategy in the stage-based evaluation game.

It can readily be verified that the above statement is true, since turn-based evaluation game is more restrictive than a stage-based evaluation game. However, for the sake of completeness, we provide the proof of Lemma 3.1 in the Appendix.

B. Augmented Barrier Certificates

We reduce the search for a positional strategy for Eloise in a turn-based evaluation game to the search for barrier-functions like certificates that we call augmented barrier certificates (ABCs). Just like barrier certificates provide a proof that separates two sets over X for arbitrary traces of the system, ABCs provide a proof that separate two sets over X^p for appropriately chosen traces by players. To tie in the notion of ABCs with HyperLTL properties, we present a special class of properties that we call *conditional invariance* properties that generalize the notion of invariance.

Definition 3.2 (Conditional Invariance (CI)): We say that a HyperLTL formula $\chi = \mu_1\pi_1 \dots \mu_p\pi_p\xi$ is a conditional invariance (CI) if ξ is of the form $G(s_A \rightarrow G(\neg s_B))$ where s_A, s_B are Boolean combination of atomic propositions.

Definition 3.3 (Augmented Barrier Certificates (ABCs)): Consider a CI $\chi = \mu_1\pi_1 \dots \mu_p\pi_p G(s_A \rightarrow G(\neg s_B))$ and the

sets $L^{p-1}(s_A) = A \subseteq X^p$ and $L^{p-1}(s_B) = B \subseteq X^p$. We say that $\mathcal{B} : X^p \rightarrow \mathbb{R}$ is an *augmented barrier certificate* (ABC) for a system $\mathfrak{S} = (X, W, f)$ and property χ from the set $A \subseteq X^p$ to set $B \subseteq X^p$ if

$$\mathcal{B}(\tilde{x}) \leq 0, \quad \text{for all } \tilde{x} \in A, \quad (\text{III.1})$$

$$\mathcal{B}(\tilde{x}) > 0, \quad \text{for all } \tilde{x} \in B, \quad (\text{III.2})$$

and $\forall x_1 \in X, \mu_1 w_1 \in W, \forall x_2 \in X, \mu_2 w_2 \in W, \dots, \forall x_p \in X, \mu_p w_p \in W$ one has:

$$\mathcal{B}(f^p(\tilde{x}, \tilde{w})) - \mathcal{B}(\tilde{x}) \leq 0, \quad (\text{III.3})$$

where $\tilde{x} = (x_1, \dots, x_p)$ and $\tilde{w} = (w_1, \dots, w_p)$.

Remark 3.4: Note that all the states' components x_i of the augmented system \mathfrak{S}^p in condition (III.3) are quantified universally, while their corresponding exogenous inputs' components w_i are quantified according to μ_i . The components x_i cannot be quantified according to μ_i as it may result in the state runs of the augmented system \mathfrak{S}^p reaching the unsafe region B . To see this, consider a component $x_i(t)$ at some time step $t \in \mathbb{N}$ which is quantified by $\mu_i = \exists$. For that component, one may be able to pick a corresponding input component $w_i(t)$ such that the augmented barrier certificate is non-increasing according to condition (III.3). However, since x_i is quantified only existentially, one may fail to ensure the existence of input $w_i(t+1)$ for $x_i(t+1)$ at the next time step such that the augmented barrier certificate is still non-increasing. Due to this, one would fail to ensure that the augmented barrier certificate remains non-increasing at every time step, possibly resulting in safety violations. Note that such quantification of the states is without loss of generality and does not restrict the class of HyperLTL specifications considered in this paper.

Remark 3.5: For a CI $\chi = \mu_1 \pi_1 \dots \mu_p \pi_p \mathbb{G}(s_A \rightarrow \mathbb{G}(\neg s_B))$, if the set $L^{p-1}(s_A) \cap L^{p-1}(s_B)$ is non-empty, then there exists no ABC satisfying conditions (III.1)-(III.3). This is due to the conflict in the satisfaction of conditions (III.1) and (III.2).

Lemma 3.6: The existence of an ABC for a conditional invariance χ implies that $T(\mathfrak{S}, L) \models \chi$.

The proof of Lemma 3.6 is similar to proving safety using standard barrier certificates, and is provided in the Appendix.

C. ABCs as “scissors”

To extend CI guarantees obtained via ABCs to arbitrary HyperLTL specifications, we first construct an NBA $\mathcal{A}_{\neg\psi} = (\Sigma^p, Q, q_0, \Delta, F)$ corresponding to $\neg\psi$. Then, we employ ABCs as *scissors* disallowing transitions to the accepting states of $\mathcal{A}_{\neg\psi}$. To do so, we first present the following lemmas to guarantee disjunction or conjunction over a set of CIs.

Lemma 3.7 (ABCs for disjunction of CIs): Given a set of CIs $\{\chi_1, \dots, \chi_k\}$, the existence of an ABC \mathcal{B}_j for some CI χ_j , where $\chi_j = \mu_1 \pi_1 \dots \mu_p \pi_p \mathbb{G}(s_{A_j} \rightarrow \mathbb{G}(\neg s_{B_j}))$ for $1 \leq j \leq k$, implies that $T(\mathfrak{S}, L) \models \chi$, where

$$\chi = \mu_1 \pi_1 \dots \mu_p \pi_p \bigvee_{1 \leq j \leq k} \mathbb{G}(s_{A_j} \rightarrow \mathbb{G}(\neg s_{B_j})).$$

We say that a function \mathcal{B} is a common ABC for a set of CIs $\{\chi_1, \dots, \chi_k\}$ if \mathcal{B} is an ABC for all of the CIs.

Lemma 3.8 (ABCs for conjunction): The existence of a common ABC for a set of CIs $\{\chi_1, \chi_2, \dots, \chi_k\}$, where $\chi_i =$

Algorithm 1 Algorithm for verification of HyperLTL formulae

Input: $\mathfrak{S}, \phi = \mu_1 \pi_1 \dots \mu_p \pi_p \psi, L$

Construct NBA $\mathcal{A}_{\neg\psi}$ for $\neg\psi$

Identify lassos $\mathcal{R} = \{r_1, r_2, \dots, r_k\}$ of $\mathcal{A}_{\neg\psi}$

for $i \leftarrow 1$ to k **do**

Identify consecutive transition pairs

$$\mathcal{S}_{\mathcal{A}_{\neg\psi}} = \{(s_{A_{i,1}}, s_{B_{i,1}}), \dots, (s_{A_{i,v_i}}, s_{B_{i,v_i}})\}$$

for $j \leftarrow 1$ to v_i **do**

$$\chi_{i,j} \leftarrow \mu_1 \pi_1 \dots \mu_p \pi_p \mathbb{G}(s_{A_{i,j}} \rightarrow \mathbb{G}(\neg s_{B_{i,j}}))$$

Construct augmented system \mathfrak{S}^p

Find common ABC \mathcal{B} for $\chi_{i,j}$ for all $i \in \{1, \dots, k\}$ and some $j \in \{1, \dots, v_i\}$

if \mathcal{B} exists **then**

return $T(\mathfrak{S}, L) \models \phi$

else

return Inconclusive

$\mu_1 \pi_1 \dots \mu_p \pi_p \mathbb{G}(s_{A_i} \rightarrow \mathbb{G}(\neg s_{B_i}))$, for $1 \leq i \leq k$, implies that $T(\mathfrak{S}, L) \models \chi$, where

$$\chi = \mu_1 \pi_1 \dots \mu_p \pi_p \bigwedge_{1 \leq i \leq k} \mathbb{G}(s_{A_i} \rightarrow \mathbb{G}(\neg s_{B_i})).$$

Lemma 3.9 (ABCs for conjunctive normal forms): Given a set of sets of conditional invariances

$$\{\{\chi_{1,1}, \dots, \chi_{1,v_1}\}, \{\chi_{2,1}, \dots, \chi_{2,v_2}\}, \dots, \{\chi_{k,1}, \dots, \chi_{k,v_k}\}\},$$

where $\chi_{i,j} = \mu_1 \pi_1 \dots \mu_p \pi_p \mathbb{G}(s_{A_{i,j}} \rightarrow \mathbb{G}(\neg s_{B_{i,j}}))$, the existence of a common ABC for $\chi_{i,j}$ for every $1 \leq i \leq k$ and some $1 \leq j \leq v_i$ implies that $T(\mathfrak{S}, L) \models \chi$, where

$$\chi = \mu_1 \pi_1 \dots \mu_p \pi_p \bigwedge_{(1 \leq i \leq k)} \bigvee_{(1 \leq j \leq v_i)} \mathbb{G}(s_{A_{i,j}} \rightarrow \mathbb{G}(\neg s_{B_{i,j}})).$$

Proofs of Lemmas 3.7-3.9 can be found in the Appendix. To find the solution to Problem 2.2, we consider the NBA $\mathcal{A}_{\neg\psi}$ corresponding to $\neg\psi$, obtained from the HyperLTL specification $\phi = \mu_1 \pi_1 \dots \mu_p \pi_p \psi$. Then, we decompose $\mathcal{A}_{\neg\psi}$ to consecutive transition pairs along its accepting lassos, or simply lassos (see Appendix for more details.) Such a lasso consists of a simple path from the initial state to some accepting state, followed by a simple cycle on the accepting state. Note that the number of lassos in $\mathcal{A}_{\neg\psi}$ is finite since the NBA has finitely many edges that lead to finitely many simple paths to an accepting state and simple cycles over the accepting state. Any two consecutive edges along these lassos constitute a transition pair that corresponds to a CI specification. Utilizing Lemma 3.9, the following theorem characterizes a condition to solve Problem 2.2.

Theorem 3.10: Given a HyperLTL specification $\phi = \mu_1 \pi_1 \dots \mu_p \pi_p \psi$, the existence of a common ABC \mathcal{B} for some consecutive transition pair along every lasso of $\mathcal{A}_{\neg\psi}$ guarantees that $T(\mathfrak{S}, L) \models \phi$.

The proof of Theorem 3.10 can be obtained by considering the decomposition of $\mathcal{A}_{\neg\psi}$ into several conditional invariance specifications by considering the lassos in $\mathcal{A}_{\neg\psi}$, and consequently applying Lemma 3.9. This is described in detail in the Appendix. To decompose the NBA $\mathcal{A}_{\neg\psi}$ into consecutive transition pairs (s_A, s_B) , one can view $\mathcal{A}_{\neg\psi}$ as a graph and utilize variants of depth-first search algorithms [38]. This

decomposition of the NBA into a collection of transition pairs was presented in [11], [18] but in the context of verification for LTL specifications. Algorithm 1 demonstrates a sound verification procedure used in this paper. We refer the interested readers to [18, Section 4] for more details on the decomposition procedure.

Example 2.3 (continued): For the room temperature regulation system \mathfrak{S} described in (II.5), we first construct the augmented system $\mathfrak{S} = \mathfrak{S} \times \mathfrak{S}$. Then, for the HyperLTL specification $\phi = \forall \pi_1 \forall \pi_2 (a_{3\pi_1} \wedge a_{4\pi_2}) \rightarrow G(a_{1\pi_1} \wedge a_{1\pi_2})$, we construct the NBA $\mathcal{A}_{\neg\psi}$ corresponding to $\neg\psi$. This is obtained as shown in Figure 2. Then, we decompose $\mathcal{A}_{\neg\psi}$ into consecutive transition pairs as explained in this section. We notice that there is one lasso for \mathcal{A}_{ψ} , resulting in one transition pair $((a_2, a_3), \neg(a_1, a_1))$. Computing an ABC for this transition pair allows determining whether $T(\mathfrak{S}, L) \models \phi$.

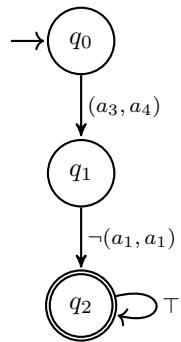


Fig. 2: NBA $\mathcal{A}_{\neg\psi}$ corresponding to $\neg\psi$

Remark 3.11: We note that the problem of finding the collection of consecutive transition pairs (thus conditional invariances), one from each lasso, which admit a common ABC is intractable. To show this, we first assume that we are given an oracle that determines whether a collection of consecutive transition pairs admits a common ABC. We then consider a relaxed version of this problem as follows. We assume that for any state r in $\mathcal{A}_{\neg\psi}$ with some incoming edge labeled s_A and outgoing edges s_{B_1}, \dots, s_{B_r} , if there exists an ABC \mathcal{B} for the pair (s_A, s_{B_j}) for some $1 \leq j \leq r$, then the function \mathcal{B} acts as an ABC for every pair (s_A, s_{B_j}) for all $1 \leq j \leq r$. Then, the problem of finding a suitable collection of transition pairs is reduced to finding a collection of edges such that their removal causes the accepting states to not be reachable from the initial state. This corresponds to a cut [39] that partitions the accepting states from the initial state. To determine whether a cut allows for a common ABC, we must make use of the oracle, and in the worst case, we need to enumerate all possible cuts in $\mathcal{A}_{\neg\psi}$. Since, the number of possible cuts is exponential in the number of edges of $\mathcal{A}_{\neg\psi}$ [39], the problem is clearly intractable.

The requirement of a common ABC for a collection of consecutive transition pairs is necessary to provide guarantees via Theorem 3.10. This is due to the fact that, in condition (III.3) of Definition 3.3, existential quantifiers may precede the universal quantifiers depending on the HyperLTL specification. In such cases, Eloise does not have access to the full-state information of the augmented system and the

choices made by Abelard in the turn-based game. However, Abelard has access to the states as well as Eloise's choices. Then, different ABCs for different transition pairs would imply that for each transition pair, Eloise picks a different strategy. This may lead to conflicts. For example, let us assume that the HyperLTL formula is of the form $\phi = \exists \pi_1 \forall \pi_2 \psi$, and consider two conditional invariances corresponding to pairs (s_{A_1}, s_{B_1}) and (s_{A_2}, s_{B_2}) . The first component of the state of the augmented system is controlled by Eloise, and the second one by Abelard. Due to a lack of full-state information of the augmented system for Eloise, she is only able to observe the label of the first component, and therefore may be unable to differentiate between s_{A_1} and s_{A_2} . Thus, having two different strategies corresponding to each of these pairs may result in ambiguity for Eloise. Moreover, picking the first strategy corresponding to (s_{A_1}, s_{B_1}) at state s_{A_2} could lead to Abelard choosing an input that violates the second conditional invariance corresponding to (s_{A_2}, s_{B_2}) , and vice-versa as Abelard selects a trace after Eloise selects her trace. This results in violation of the original specification. Unfortunately, even though a common ABC is necessary to provide verification guarantees, its existence may be difficult to find.

However, in specifications where Eloise has access to full state information and all of Abelard's choices, the requirement of a common ABC may be relaxed. This is especially true for specifications in the $\forall^* \exists^*$ fragment, where all the universal quantifiers precede the existential ones. In fact, the $\forall^* \exists^*$ fragment holds great importance as it comprises of many relevant security properties. For example, a variant of the noninterference property [40] requires that, for all traces, the low-security variables should not see any difference in observation when high-security variables are changed and replaced by dummy variables. This can be expressed by the HyperLTL specification

$$\forall \pi_1 \exists \pi_2 (G h_{\pi_2}) \wedge \bigwedge_{l \in LS} l_{\pi_1} \leftrightarrow l_{\pi_2},$$

where h_{π_2} implies that the high security variables in π_2 are all set to a dummy variable h that is always true, and $LS \in AP$ denotes the set of low security variables. Similarly, initial-state opacity specification [36] is also in the $\forall^* \exists^*$ fragment (cf. case study). Considering the importance of this fragment, we now provide a separate algorithm to allow for multiple ABCs for different lassos under some conditions.

D. Algorithm for $\forall^* \exists^*$ -fragment of HyperLTL

From the above discussion, it can be understood that specifications in the $\forall^* \exists^*$ fragment enable the relaxation of common ABC requirement and allow for different ABCs in different lassos. In particular, Eloise can take advantage of the full state information of the augmented system available to her as well as the knowledge of Abelard's choices to use different ABCs for different consecutive transition pairs in every lasso. However, to do so, one must take the structure of the automata $\mathcal{A}_{\neg\psi}$ into consideration, as in the presence of states with two or more outgoing edges, there may be an ambiguity for Eloise in selecting strategies. Moreover, in the presence of

nondeterminism in the automaton, Eloise may fail to select strategies due to lack of information on the history of visited states. These challenges are demonstrated in the following examples.

Example 3.12 (States with a fork): In this example, we show the issue of utilizing multiple ABCs in the presence of a state with multiple outgoing edges. Consider the NBA $\mathcal{A}_{\neg\psi}$ shown in Figure 3 constructed from a set of atomic propositions $AP = \{a, b, c, d\}$ corresponding to some HyperLTL specification $\phi = \forall\pi_1 \dots \forall\pi_l \exists\pi_{l+1} \dots \exists\pi_p \psi$.

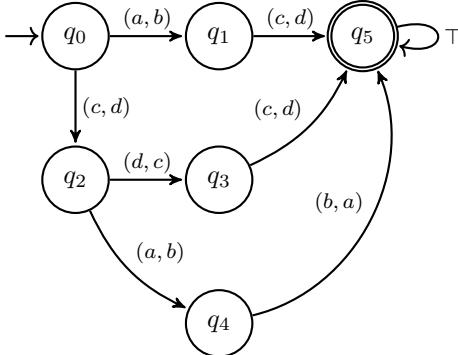


Fig. 3: NBA $\mathcal{A}_{\neg\psi}$ for Example 3.12

From $\mathcal{A}_{\neg\psi}$, we can identify $k = 3$ lassos as

$$\begin{aligned}\mathcal{R} = \{\mathbf{r}_1 = (q_0, q_1, q_5, q_5), \mathbf{r}_2 = (q_0, q_2, q_3, q_5, q_5), \\ \mathbf{r}_3 = (q_0, q_2, q_4, q_5, q_5)\}.\end{aligned}$$

For every $\mathbf{r} \in \mathcal{R}$, we obtain the consecutive transition pairs as

$$\begin{aligned}\mathcal{S}(\mathbf{r}_1) &= \{((a, b), (c, d)), ((c, d), \top)\}, \\ \mathcal{S}(\mathbf{r}_2) &= \{((c, d), (d, c)), ((d, c), (c, d)), ((c, d), \top)\}, \\ \mathcal{S}(\mathbf{r}_3) &= \{((c, d), (a, b)), ((a, b), (b, a)), ((b, a), \top)\}.\end{aligned}$$

Naturally, it is preferable to obtain different ABCs for at least one transition pair in every lasso to guarantee the satisfaction of the specification. However, this might cause problems for lassos \mathbf{r}_2 and \mathbf{r}_3 , where there are two outgoing edges from a single state q_2 . This leads to two different transition pairs $((c, d), (d, c))$ and $((c, d), (a, b))$. Having different ABCs for these pairs would result in different winning strategies for Eloise to avoid the sets corresponding to (d, c) and (a, b) , from the set corresponding to (c, d) , respectively. Choosing the first ABC and its corresponding strategy could lead to the violation of condition (III.3) for the second ABC and vice versa. However, the existence of a common ABC for both the pairs guarantees that Eloise has a winning strategy to avoid both (d, c) and (a, b) if she encounters a state corresponding to (c, d) . Therefore, for this specification, one would require to obtain a common ABC for the pairs $((c, d), (d, c))$ and $((c, d), (a, b))$ from lassos \mathbf{r}_2 and \mathbf{r}_3 , respectively, and a different ABC may be obtained for the pair $((a, b), (c, d))$ from the lasso \mathbf{r}_1 . However, if such a common ABC cannot be found, one can consider other transition pairs in \mathbf{r}_2 and \mathbf{r}_3 , and in that case, different ABCs may be used.

Example 3.13 (Nondeterminism): In this example, we show the issue of using multiple ABCs in the presence of nondeterminism in the automaton. Consider the NBA $\mathcal{A}_{\neg\psi}$ shown in

Figure 4 constructed from a set of atomic propositions $AP = \{a, b, c, d\}$ corresponding to some HyperLTL specification $\phi = \forall\pi_1 \dots \forall\pi_l \exists\pi_{l+1} \dots \exists\pi_p \psi$.

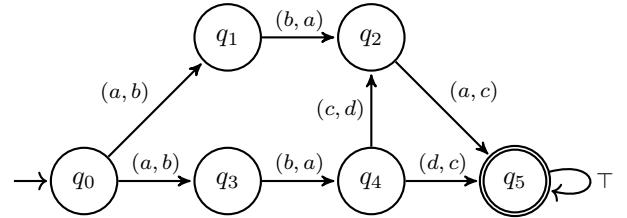


Fig. 4: NBA $\mathcal{A}_{\neg\psi}$ for Example 3.13

From $\mathcal{A}_{\neg\psi}$, we can identify $k = 3$ lassos as

$$\begin{aligned}\mathcal{R} = \{\mathbf{r}_1 = (q_0, q_1, q_2, q_5, q_5), \mathbf{r}_2 = (q_0, q_3, q_4, q_5, q_5), \\ \mathbf{r}_3 = (q_0, q_3, q_4, q_2, q_5, q_5)\}.\end{aligned}$$

For every $\mathbf{r} \in \mathcal{R}$, we obtain the consecutive transition pairs as

$$\begin{aligned}\mathcal{S}(\mathbf{r}_1) &= \{((a, b), (b, a)), ((b, a), (a, c)), ((a, c), \top)\}, \\ \mathcal{S}(\mathbf{r}_2) &= \{((a, b), (b, a)), ((b, a), (d, c)), ((d, c), \top)\}, \\ \mathcal{S}(\mathbf{r}_3) &= \{((a, b), (b, a)), ((b, a), (c, d)), ((c, d), (a, c)), \\ &\quad ((a, c), \top)\}.\end{aligned}$$

Consider lassos \mathbf{r}_1 , \mathbf{r}_2 and \mathbf{r}_3 , where there is a nondeterministic transition from the initial state q_0 to the states q_1 and q_3 under the label (a, b) . Ideally, a single ABC for the pair $((a, b), (b, a))$ would effectively disallow the transitions in all the lassos \mathbf{r}_1 , \mathbf{r}_2 and \mathbf{r}_3 . However, the problem arises when such an ABC cannot be found. In order to guarantee the satisfaction of the specification, other transition pairs in the lassos must be disallowed. Now, in \mathbf{r}_1 , there is a transition from (b, a) to (a, c) , while in \mathbf{r}_2 and \mathbf{r}_3 , there are two transitions from (b, a) to (d, c) and (c, d) , respectively. At any point in time, Eloise cannot uniquely determine the history of the states visited in $\mathcal{A}_{\neg\psi}$. As a result, after a nondeterministic transition from (a, b) to (b, a) , Eloise has no way of knowing whether to block further transitions from (b, a) to (a, c) , or from (b, a) to (d, c) and (c, d) . Therefore, the approach of using different ABCs fails in the presence of nondeterminism.

Unfortunately, problems arising due to nondeterminism cannot be directly resolved. Therefore, to circumvent this issue, we instead consider the automaton $\mathcal{A}_{\neg\psi}$ to be deterministic. In particular, for our exposition we focus on the case where $\mathcal{A}_{\neg\psi}$ is a deterministic Büchi automaton (DBA).

The general verification procedure for determining whether Eloise has a strategy to ensure that the acceptance condition of $\mathcal{A}_{\neg\psi}$ is violated for specifications in the $\forall^* \exists^*$ fragment is provided in Algorithms 2 and 3. Having a DBA $\mathcal{A}_{\neg\psi} = (\Sigma^p, Q, q_0, \Delta, F)$, we first identify all the lassos $\mathcal{R} = \{\mathbf{r}_1, \dots, \mathbf{r}_k\}$ that reach and cycle on some state in F . We then prune $\mathcal{A}_{\neg\psi}$ and remove any states that are not in the lassos and all transitions to and from such states. Then, beginning from the initial state, for any label s_A , we identify the state that is reachable via s_A . Let the outgoing transitions from this state be s_{B_1}, \dots, s_{B_n} . Let $S = \{(s_A, s_{B_1}), \dots, (s_A, s_{B_n})\}$. A suitable common ABC is searched for all the transition pairs in

Algorithm 2 Algorithm for verification of $\forall^* \exists^*$ fragment of HyperLTL

Input: $\mathfrak{S}, \phi = \forall \pi_1 \dots \forall \pi_l \exists \pi_{l+1} \dots \exists \pi_p \psi, L$

Construct DBA $\mathcal{A}_{\neg\psi}$ for $\neg\psi$

Identify lassos $\mathcal{R} := \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_k\}$ of $\mathcal{A}_{\neg\psi}$

$R_d \leftarrow \emptyset$

$\mathcal{A}'_{\neg\psi} := (\Sigma^p, Q', q_0, \Delta', F') \leftarrow \text{Prune}(\mathcal{A}_{\neg\psi})$

$\text{visited} \leftarrow [0, \dots, 0]$ \triangleright Array of size $|Q'|$

$V \leftarrow \{q_0\}$

$\text{visited}[q_0] \leftarrow 1$

while $V \neq \emptyset$ **do**

for each $r \in V$ **do**

$R_d \leftarrow R_d \cup \text{ABC_FIND}(\mathcal{A}'_{\neg\psi}, r)$

$\mathcal{R}_m \leftarrow \mathcal{R} \setminus R_d$

if $\mathcal{R}_m = \emptyset$ **then**

return $T(\mathfrak{S}, L) \models \phi$

for each $s_A \in \Sigma^p$ **do**

$r' \leftarrow \Delta'(r, s_A)$

$G \leftarrow \{\mathbf{r} \in \mathcal{R}_m \mid (r, r') \in \mathbf{r}\}$

if $G \neq \emptyset$ **and** $\text{visited}[r'] < 1$ **then**

$V \leftarrow V \cup \{r'\}$

$\text{visited}[r'] \leftarrow \text{visited}[r'] + 1$

$V \leftarrow V \setminus \{r\}$

return Inconclusive

a set $S_a \subseteq S$. For every pair $(s_A, s_{B_j}) \in S_a$, for all $0 \leq j \leq n$, we also have that $(s_A, s_{B_j}) \in \mathcal{S}(\mathbf{r}_i)$ for some $0 \leq i \leq k$. If such a common ABC exists, then such lassos \mathbf{r}_i can be discarded from further consideration as existence of ABCs disallows the transitions in those lassos and are collected in the set \mathcal{R}_d . Note that, unfortunately, there is no systematic way to obtain the set $S_a \subseteq S$ consisting of all transition pairs that admit a common ABC. S_a is first picked in a trial-and-error fashion and then an *oracle* (c.f Remark 3.11) is used to determine if the transition pairs in S_a admit a common ABC.

This procedure is then repeated for every transition label at the initial state, and the discarded lassos are iteratively added to \mathcal{R}_d . Once all the outgoing transition labels are covered, we move on to the next state reachable from the initial state and repeat the procedure to find ABCs for only those pairs that belong to $\mathcal{S}(\mathbf{r}_i)$ such that $\mathbf{r}_i \in \mathcal{R} \setminus \mathcal{R}_d$. This continues in a breadth-first search fashion until all the lassos are discarded, *i.e.*, $\mathcal{R}_d = \mathcal{R}$, in which case we can conclude that $T(\mathfrak{S}, L) \models \phi$, or all the states of $\mathcal{A}_{\neg\psi}$ have been considered. If $\mathcal{R}_d \subset \mathcal{R}$, it means that there are lassos for which no ABC could be found, rendering the verification procedure inconclusive.

IV. ABCs VIA SUM-OF-SQUARES PROGRAMMING

In the previous sections, we showed that existence of an ABC satisfying conditions (III.1)-(III.3) for a transition pair (s_A, s_B) is vital to verify that a system \mathfrak{S} satisfies a desired HyperLTL specification ϕ . In general, synthesizing such ABCs is a difficult problem. However, under some assumptions on the type of ABCs considered, the dynamics of the systems, and the geometry of the state sets, one can efficiently compute ABCs that can sufficiently prove the satisfaction of conditional invariance guarantees. This can in turn be utilized to verify the satisfaction of HyperLTL specifications. Specifically, we

Algorithm 3 Function *ABC_FIND*

Input: $\mathcal{A}'_{\neg\psi} := (\Sigma^p, Q', q_0, \Delta', F'), r$

$R_d \leftarrow \emptyset$

for each $s_A \in \Sigma^p$ **do**

$S \leftarrow \emptyset$

$r' \leftarrow \Delta'(r, s_A)$

for each $s_B \in \Sigma^p$ **do**

if $\Delta'(r', s_B) \neq \emptyset$ **then** $S \leftarrow S \cup \{(s_A, s_B)\}$

 Find a common ABC for a set $S_a \subseteq S$.

for each $(s_A, s_B) \in S_a$ **do**

$r' \leftarrow \Delta'(r, s_A)$

$r'' \leftarrow \Delta'(r', s_B)$

$R_d \leftarrow R_d \cup \{\mathbf{r} \in \mathcal{R} \mid (r, r', r'') \in \mathbf{r}\}$

return R_d

see that when the dynamics of the systems are restricted to polynomial functions and the state set X , exogenous input set W as well as the safe and unsafe sets obtained from (s_A, s_B) are semi-algebraic sets [41], one can utilize sum-of-squares (SOS) programming techniques [42] to compute polynomial ABCs of predefined degrees. We now formally state the following assumption.

Assumption 4.1: The system \mathfrak{S} has a continuous state set $X \subseteq \mathbb{R}^n$ and continuous exogenous input set $W \in \mathbb{R}^m$, and its transition function $f : X \times W \rightarrow X$ is a polynomial function of the state x and input w .

Under Assumption 4.1, one can readily observe that the state and input sets of augmented system \mathfrak{S}^p (i.e. X^p and W^p , respectively) are also continuous, and the function $f^p : X^p \times W^p \rightarrow X^p$ is a p -tuple of polynomial functions. Having this, one can then reformulate conditions (III.1)-(III.3) as an SOS optimization problem (cf. next lemma) to search for a polynomial ABC for augmented system \mathfrak{S}^p . In order to present the result below, we assume that the number of quantifiers “ \exists ” in $\phi = \mu_1 \pi_1 \dots \mu_p \pi_p \psi$ is equal to k and define $I_{\exists} = \{i \mid \mu_i = \exists, 1 \leq i \leq p\}$.

Lemma 4.2: Suppose Assumption 4.1 holds and sets X^p , A , B , and W^p are defined as $X^p = \{\tilde{x} \in \mathbb{R}^{np} \mid g(\tilde{x}) \geq 0\}$, $A = \{\tilde{x} \in \mathbb{R}^{np} \mid g_0(\tilde{x}) \geq 0\}$, $B = \{\tilde{x} \in \mathbb{R}^{np} \mid g_u(\tilde{x}) \geq 0\}$, and $W^p = \{\tilde{w} \in \mathbb{R}^{mp} \mid g_{in}(\tilde{w}) \geq 0\}$, where the inequalities are considered component-wise and functions g, g_0, g_u , and g_{in} are polynomials. Suppose there exist a polynomial $\mathbb{B}(\tilde{x})$ and k polynomials $h_j^i(\hat{x}_i, \hat{w}_i)$, $i \in I_{\exists}$, corresponding to the j^{th} entry of $w_i = (w_{i_1}, \dots, w_{i_m}) \in W \subseteq \mathbb{R}^m$, where \hat{x}_i refers to those components of the state with indices less than i and \hat{w}_i denotes the inputs associated with “ \forall ” quantifiers with indices less than i . In addition, suppose there exist sum-of-squares polynomials $\lambda(\tilde{x}, \tilde{w}), \lambda_0(\tilde{x}), \lambda_u(\tilde{x})$, and $\lambda_{in}(\tilde{x}, \tilde{w})$ of appropriate dimensions, such that the following expressions are sum-of-square polynomials:

$$-\mathcal{B}(\tilde{x}) - \lambda_0(\tilde{x})g_0^T(\tilde{x}), \quad (\text{IV.1})$$

$$\mathcal{B}(\tilde{x}) - \lambda_u(\tilde{x})g_u^T(\tilde{x}) - \varepsilon, \quad (\text{IV.2})$$

$$\begin{aligned} & -\mathcal{B}(f^p(\tilde{x}, \tilde{w})) + \mathcal{B}(\tilde{x}) - \lambda(\tilde{x}, \tilde{w})g^T(\tilde{x}) - \lambda_{in}(\tilde{x}, \tilde{w})g_{in}^T(\tilde{w}) \\ & - \sum_{i \in I_{\exists}} \sum_{j=1}^m (w_{i_j} - h_j^i(\hat{x}_i, \hat{w}_i)), \end{aligned} \quad (\text{IV.3})$$

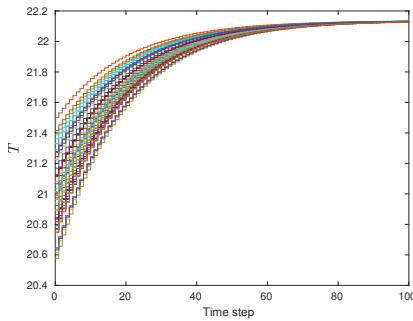


Fig. 5

Fig. 6: State runs of the system \mathfrak{S} starting from initial set X_3 .

where ε is a small positive number. Then, $\mathcal{B}(\tilde{x})$ is an ABC from set A to set B satisfying conditions (III.1)-(III.3).

Note that a small tolerance ε is needed in (IV.2) to ensure strict positivity of ABC as required in (III.2). The proof of Lemma 4.2 relies on the fact that if conditions (IV.1)-(IV.3) are rendered SOS by polynomials $\mathcal{B}(\tilde{x})$ and $h_j^i(\hat{x}_i, \hat{w}_i)$, then conditions (II.2)-(II.4) are satisfied. More details on the proof can be found in the Appendix. Existing tools such as SOSTools [43] can be used in conjunction with semi-definite programming solvers such as SeDuMi [44] to compute polynomial ABCs satisfying (IV.1)-(IV.3).

Remark 4.3: Note that the SOS approach for computing barrier certificates is restricted to systems with polynomial-type dynamics. One can also utilize different computational techniques when dynamics are not necessarily polynomial. For example, conditions (III.1)-(III.3) can be reformulated as a satisfiability problem and SMT solvers such as Z3 [45] and dReal [46] can be utilized to search for suitable barrier certificates using counterexample-guided inductive synthesis framework [47]. One can also train suitable barrier certificates using neural networks e.g. [48], [49].

Remark 4.4: The complexity analysis of our approach is described as follows. Suppose there exists a common ABC of degree $2d$ that verifies a system \mathfrak{S} against a HyperLTL formula $\phi = \mu_1, \pi_1 \dots, \mu_p, \pi_p \psi$, consisting of p trace quantifiers. To find this ABC, one requires to consider all the consecutive transition pairs of the NBA $\mathcal{A}_{\neg\psi}$ corresponding to the specification, and then select suitable consecutive pairs in every lasso such that they admit a common ABC. Let the number of states of $\mathcal{A}_{\neg\psi}$ be $|Q|$. There are $O(2^{|Q|^2})$ possible subsets of consecutive transition pairs for $\mathcal{A}_{\neg\psi}$. Then, for each such subset, the common ABC needs to be computed over the p -fold augmented system with dimension pn via SOS programming, resulting in a complexity given by $O\left(\binom{pn+d}{d}^2\right)$ [42], where $2d$ is the degree of the ABC. Therefore, the final complexity of our approach is polynomial in $O(2^{|Q|^2} \binom{pn+d}{d}^2)$. One observes that this complexity depends not just on the dimension of the original system n , but also on the number of quantifiers in the HyperLTL formula p . Even in the case of two quantifiers, this leads to a doubling in dimension. As the state dimension n and number of quantifiers p grow, we typically need to search for barrier certificates of higher degree (thus the value $2d$ needs to also grow). As a result, the value $\binom{pn+d}{d}^2$ grows at a rate that is worse than polynomial in the state dimension n . Thus,

our approach is not particularly scalable to high-dimensional systems. As future work we plan to investigate the use of compositional approaches such as in [34] to tackle this issue.

Example 2.3 (Continued): We now utilize SOS programming to compute ABC for the transition pair $((a_2, a_3), \neg(a_1, a_1))$ obtained from Figure 2. We use the tools SOSTOOLS [43] and SeDuMi [44] on MATLAB to compute a polynomial ABC of degree 2 as $\mathcal{B}(T_1, T_2) = 1.2454T_1^2 - 1.6722T_1T_2 - 18.5791T_1 + 1.1656T_2^2 - 14.9555T_2 + 377.4684$ with a tolerance of $\varepsilon = 0.001$. The existence of the ABC proves that the safety controller designed for the system \mathfrak{S} , is indeed robust with respect to initial-state uncertainty with a robustness measure of $\delta = 0.5$. Figure 5 shows that the state runs obtained for \mathfrak{S} remains in the safe set $X_1 = [20, 25]$ when starting from the initial set $X_3 = [20.5, 21.5]$ which captures uncertainties in the initial state. We performed these computations on a machine running Linux Ubuntu OS (Intel i7-8665U CPU with 32GB RAM) and it took around 19 seconds for the computation of ABC.

In the following, we present another case study to demonstrate the effectiveness of our approach.

V. CASE STUDY

In this example, we consider the discrete-time, two-dimensional model of an autonomous vehicle on a single-lane road, with state variables as $x = [s, v]$, where s denotes the absolute position of the vehicle and v denotes the absolute velocity. The dynamics of the system are borrowed from [50] and governed by:

$$\mathfrak{S} : \begin{cases} s(t+1) = s(t) + \Delta\tau v(t) + \frac{\Delta\tau^2}{2}w(t), \\ v(t+1) = v(t) + \Delta\tau w(t), \end{cases} \quad (V.1)$$

where w is the exogenous input, *i.e.*, acceleration, and $\Delta\tau = 1$ is the sampling time. Here, we verify the δ -approximate initial state opacity property [50] for this system. The specification requires that, for any state run of the system that begins from a secret state, there must exist another state run that begins from a non-secret state such that both state runs render δ -close observations from the observer's (or intruder's) point of view. The significance of the specification can be motivated with the help of a simple example. Consider a scenario where the vehicle is assigned for a cash transit from a high-security bank to an ATM machine, and the initial locations of the vehicle must be kept secret. It is assumed that a malicious intruder is observing the velocity of the vehicle remotely and intends to gain access to the secret information and perform an attack. Therefore, it is critical to ensure that the secret states of the system are never revealed to the intruder. This security specification can be modeled as a δ -approximate initial-state opacity problem, where $\delta \geq 0$ captures the measurement precision of the intruder.

To express δ -approximate initial-state opacity as a HyperLTL specification, consider system (V.1) with state set $X = [0, 8] \times [0, 0.6]$ and exogenous input set $W = [-0.04, 0.04]$. The secret set is defined by $X_1 = [0, 1] \times [0, 0.6]$ and the non-secret set is consequently given by $X_2 = X \setminus X_1$. Here, we assume that the intruder can only observe the velocity of

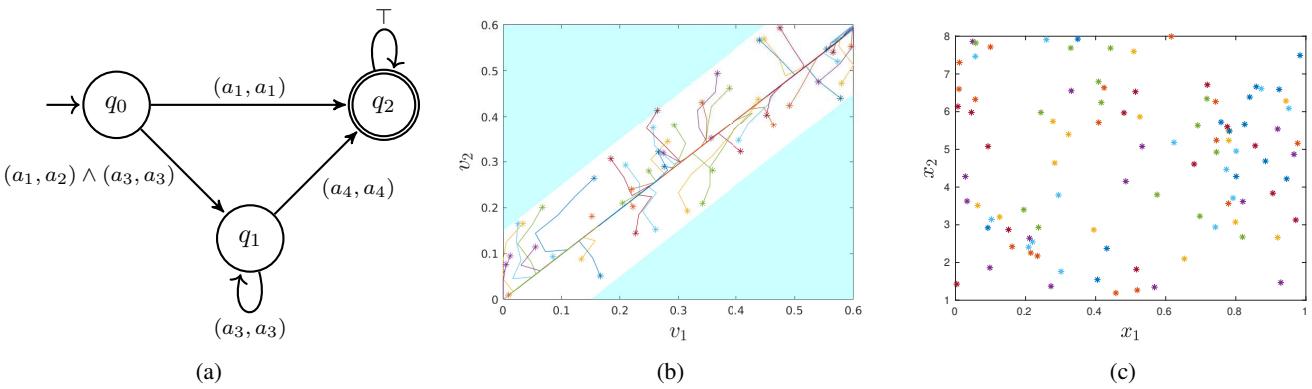


Fig. 7: (a) NBA $\mathcal{A}_{\neg\psi}$ corresponding to $\neg\psi$ (b) State runs of \mathfrak{S}^2 projected over the velocity coordinate. Region in blue indicates the unsafe set (c) The initial conditions of the state runs projected over the position coordinate marked by * which show that the first initial condition (i.e. x_1) is secret and the other one (i.e. x_2) is non-secret.

the car with a precision of δ , *i.e.*, observations of two states $x_1 = [s_1, v_1] \in X$ and $x_2 = [s_2, v_2] \in X$ appear identical to the intruder if $\|v_1 - v_2\| \leq \delta$. We now construct atomic propositions as $\text{AP} = \{a_1, a_2, a_3, a_4\}$ where a_1 and a_2 are such that $L(x \in X_z) = a_z$ for $z = \{1, 2\}$. The atomic propositions a_3 and a_4 are constructed over the augmented state set such that we have $(a_3, a_3) := \{(L(x_1 = [s_1, v_1] \in X), L(x_2 = [s_2, v_2] \in X)) \mid \|v_1 - v_2\|^2 \leq \delta^2\}$ and $(a_4, a_4) := \{(L(x_1 = [s_1, v_1] \in X), L(x_2 = [s_2, v_2] \in X)) \mid \|v_1 - v_2\|^2 \geq \delta^2 + \epsilon\}$, where ϵ is a small positive number introduced to certify positivity using SOS programming. Note that atomic propositions for HyperLTL specifications are usually defined over the single system rather than the augmented one. On the other hand, the δ -approximate initial state opacity specification requires the atomic propositions to capture the δ -closeness between any two states of the augmented system. In a finite-state system, one could quantify δ -closeness by using finite conjuncts of atomic propositions defined over the original system, but in the infinite-state case such as ours, that is not possible. Therefore, to handle this non-trivial case, we modify atomic propositions slightly and define them over the augmented state set. Such modifications can be made without any loss of generality in our approach. Now, one can formulate the δ -approximate initial-state opacity specification as a HyperLTL formula given by $\phi = \forall \pi_1 \exists \pi_2 \psi$, where $\psi = a_{1\pi_1} \rightarrow (a_{2\pi_2} \wedge G(a_{3\pi_1} \wedge a_{3\pi_2}))$.

Consider the system $\mathfrak{S}^2 = \mathfrak{S} \times \mathfrak{S}$ with states $(x_1 = [s_1, v_1], x_2 = [s_2, v_2]) \in X^2$ and input $(w_1, w_2) \in W^2$, and the NBA $\mathcal{A}_{\neg\psi}$ corresponding to $\neg\psi$ that is obtained as shown in Figure 7a. We decompose $\mathcal{A}_{\neg\psi}$ to obtain transition pairs for all lassos. This is obtained as $((a_1, a_2) \wedge (a_3, a_3), (a_4, a_4))$, $((a_1, a_1), \top)$ and $((a_4, a_4), \top)$. The latter two do not admit ABC following Remark 3.5, and the transition pair $((a_1, a_1), \top)$ is ignored by assuming that the augmented system \mathfrak{S}^2 never starts from an initial condition corresponding to $\tilde{a}_1 = (a_1, a_1)$. Note that this assumption is only on the *virtual* copy of the system \mathfrak{S} and does not restrict the initial states of the original system \mathfrak{S} directly. For the transition pair $((a_1, a_2) \wedge (a_3, a_3), (a_4, a_4))$, we compute a suitable ABC by considering $\delta = 0.15$. Using SOSTOOLS and SeDuMi tools on MATLAB, and with

tolerance parameters $\epsilon = 0.01$ and $\varepsilon = 0.015$, we obtain ABC as follows.

$$\begin{aligned} \mathcal{B}((s_1, v_1), (s_2, v_2)) = & 85.03v_1^2 - 170.3v_1v_2 + 0.0048v_1s_1 \\ & - 0.0065v_1s_2 + 0.0413v_1 + 85.24v_2^2 - 0.004784v_2s_1 \\ & + 0.0063v_2s_2 - 0.0121v_2 + 0.0059s_1^2 - 0.0119s_1s_2 \\ & + 0.0241s_1 + 0.0061s_2^2 - 0.0825s_2 - 2.076, \end{aligned}$$

and the corresponding \exists quantifier on the input is fulfilled by $w_2(s_1, v_1, s_2, v_2, w_1) = 0.983v_1 - v_2 + w_1$. Therefore, we conclude that the system \mathfrak{S} satisfies the HyperLTL specification ϕ representing δ -approximate initial-state opacity problem with $\delta = 0.15$. Figure 7b shows the projection of a few state runs on the velocity coordinate of the augmented system \mathfrak{S} , with initial conditions in $A = L^{p-1}((a_1, a_2) \wedge (a_3, a_3))$. Figure 7c shows the initial conditions projected on the position coordinate. It follows that the state runs avoid reaching the unsafe regions, indicating that the original system is δ -approximate initial-state opacity. We should add that the computation of ABCs using the mentioned tools on MATLAB takes roughly 35 seconds on a machine running with Linux Ubuntu OS (Intel i7-8665U CPU with a 32 GB of RAM). The codes for our experiments may be found on our public repository: <https://github.com/mahathi-anand/CPS-Verification-against-HyperLTL.git>.

VI. DISCUSSION AND CONCLUSION

A. HyperLTL Synthesis

An interesting problem that follows HyperLTL verification for dynamical systems is the synthesis of controllers ensuring the satisfaction of HyperLTL specifications. In this case, for a system $\mathfrak{S} = (X, U, f)$ and a HyperLTL specification ϕ , one would view ν in (II.1) as control signal rather than exogenous one and design ν such that the corresponding traces T of \mathfrak{S} satisfy ϕ . Unfortunately, there are major challenges in synthesizing controllers even when the HyperLTL specification is a simple conditional invariance (CI). Let us consider a CI χ , and a controller $\mathcal{G} : X \rightarrow U$ such that $\nu(t) := \mathcal{G}(x(t))$. Then, condition (III.3) of the barrier certificate for CI χ in the context of synthesis can be reformulated as: for any

$$\tilde{x} = (x_1, x_2, \dots, x_p) \in X^p,$$

$$\mathcal{B}(f(x_1, \mathcal{G}(x_1)), f(x_2, \mathcal{G}(x_2)), \dots, f(x_p, \mathcal{G}(x_p))) - \mathcal{B}(\tilde{x}) \leq 0. \quad (\text{VI.1})$$

The above formulation ensures that the selection of the control input at any given state $x \in X$ according to map \mathcal{G} is independent of the previous traces selected by the players.

However, to satisfy this condition, one must simultaneously search for suitable functions \mathcal{B} and \mathcal{G} . This makes the above inequality non-convex in these unknown functions and unfortunately, one cannot leverage convex programming and correspondingly SOS and semi-definite programming to determine these functions even when they are assumed to be polynomials. However, given a map \mathcal{G} , one could search for a function \mathcal{B} such that condition (VI.1) is satisfied, which is technically a verification problem and not a synthesis one anymore. In general, even though one can verify whether a HyperLTL specification ϕ is realizable over a system, it is not possible to synthesize the control map \mathcal{G} that ensures the satisfaction of ϕ . In other words, it is not possible to find a solution to the HyperLTL synthesis problem.

This is due to the fact that the inputs obtained satisfying condition (III.3) in the case of verification may depend on the previously quantified traces, which is not possible when considering controller synthesis. Remark that the problem of HyperLTL verification coincides with HyperLTL synthesis when the specification is of the form $\phi = \exists \pi_1 \forall \pi_2 \dots \forall \pi_p \psi$. To verify such specifications, it is sufficient to synthesize a controller for the first trace. Since the remaining traces are controlled by Abelard, ABCs satisfying condition (III.3) implies the satisfaction of condition (VI.1) with the controller \mathcal{G} being the one synthesized for the first trace. We leave any further investigations for the synthesis problem as future work.

B. Conclusion

We proposed a discretization-free scheme for the verification of discrete-time uncertain dynamical systems against hyperproperties. Focusing on hyperproperties expressible as HyperLTL formulae, we presented an implicit automata-theoretic approach. In our approach the specifications are reduced to a collection of conditional invariance properties by utilizing an implicitly quantified Büchi automata corresponding to the complements of the specifications. Working with an augmented system, we were able to devise a notion of augmented barrier certificates over the self-composition of the original system as a certificate of conditional invariance. The existence of ABCs is a sufficient proof that the conditional invariance holds, this provides a verification guarantee over the satisfaction of the hyperproperty. For a general HyperLTL specification, we showed that a common ABC for at least one conditional invariance in every lasso is required to provide verification guarantees. However, for a HyperLTL specification in the $\forall^* \exists^*$ fragment, we provided a systematic algorithmic procedure that leverages the structure of the automata to allow for different ABCs for different lassos. We exploited a sum-of-squares approach to efficiently compute suitable ABCs. As future work, we plan on investigating approaches that allow

for multiple ABCs to guarantee the satisfaction of general HyperLTL specifications. We would also investigate verification problems for stochastic systems and synthesis problems against hyperproperties for continuous-state control systems. Moreover, we would also utilize compositionality approaches [34] to tackle scalability issues in computing ABCs for large-scale systems.

REFERENCES

- [1] P. Tabuada, *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media, 2009.
- [2] C. Belta, B. Yordanov, and E. Göll, *Formal methods for discrete-time dynamical systems*. Studies in Systems, Decision and Control, Springer, 2017.
- [3] K. Bae and J. Lee, "Bounded model checking of signal temporal logic properties using syntactic separation," *Proc. ACM Program. Lang.*, vol. 3, 2019.
- [4] C. Baier and J.-P. Katoen, *Principles of model checking*. MIT press, 2008.
- [5] M. R. Clarkson and F. B. Schneider, "Hyperproperties," *Journal of Computer Security*, vol. 18, no. 6, pp. 1157–1210, 2010.
- [6] L. Mazaré, "Using unification for opacity properties," in *Proceedings of the Workshop on Issues In The Theory of Security*, pp. 165–176, 2004.
- [7] Y. Wang, S. Nalluri, and M. Pajic, "Hyperproperties for robotics: Planning via hyperlTL," in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 8462–8468, 2020.
- [8] J. A. Goguen and J. Meseguer, "Security policies and security models," in *IEEE Symposium on Security and Privacy*, p. 11, 1982.
- [9] A. W. Roscoe, "CSP and determinism in security modelling," in *Proceedings 1995 IEEE Symposium on Security and Privacy*, pp. 114–127, 1995.
- [10] M. R. Clarkson, B. Finkbeiner, M. Koleini, K. K. Micinski, M. N. Rabe, and C. Sánchez, "Temporal logics for hyperproperties," in *Principles of Security and Trust, LNCS*, pp. 265–284, 2014.
- [11] T. Wongpiromsarn, U. Topcu, and A. Lamperski, "Automata theory meets barrier certificates: Temporal logic verification of nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 11, pp. 3344–3355, 2016.
- [12] P. Tabuada and G. J. Pappas, "Linear time logic control of discrete-time linear systems," *IEEE Transactions on Automatic Control*, vol. 51, no. 12, pp. 1862–1877, 2006.
- [13] M. Zamani, G. Pola, M. M. Jr., and P. Tabuada, "Symbolic models for nonlinear control systems without stability assumptions," *IEEE Transaction on Automatic Control*, vol. 57, pp. 1804–1809, July 2012.
- [14] B. Yordanov, J. Tumova, I. Cerna, J. Barnat, and C. Belta, "Temporal logic control of discrete-time piecewise affine systems," *IEEE Transactions on Automatic Control*, vol. 57, no. 6, pp. 1491–1504, 2012.
- [15] E. M. Wolff, U. Topcu, and R. M. Murray, "Automaton-guided controller synthesis for nonlinear systems with temporal logic," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4332–4339, 2013.
- [16] S. Prajna and A. Jadbabaie, "Safety verification of hybrid systems using barrier certificates," in *Hybrid Systems: Computation and Control, Lecture Notes in Computer Science*, pp. 477–492, 2004.
- [17] A. Bisoffi and D. V. Dimarogonas, "A hybrid barrier certificate approach to satisfy linear temporal logic specifications," in *Annual ACC*, pp. 634–639, 2018.
- [18] P. Jagtap, S. Soudjani, and M. Zamani, "Temporal logic verification of stochastic systems using barrier certificates," in *Automated Technology for Verification and Analysis, Lecture Notes in Computer Science*, pp. 177–193, 2018.
- [19] M. Anand, P. Jagtap, and M. Zamani, "Verification of switched stochastic systems via barrier certificates," in *IEEE 58th CDC*, pp. 4373–4378, 2019.
- [20] L. Lindemann and D. V. Dimarogonas, "Control barrier functions for signal temporal logic tasks," *IEEE Control Systems Letters*, vol. 3, no. 1, pp. 96–101, 2019.
- [21] G. Yang, C. Belta, and R. Tron, "Continuous-time signal temporal logic planning with control barrier functions," in *2020 American Control Conference (ACC)*, pp. 4612–4618, 2020.
- [22] B. Finkbeiner, M. N. Rabe, and C. Sánchez, "Algorithms for model checking HyperLTL and HyperCTL*," in *Computer Aided Verification*, 2015.
- [23] B. Finkbeiner, C. Hahn, and H. Torfah, "Model checking quantitative hyperproperties," in *Computer Aided Verification*, pp. 144–163, 2018.

[24] T.-H. Hsu, C. Sanchez, and B. Bonakdarpour, "Bounded model checking for hyperproperties," 2020. arXiv: 2009.08907.

[25] B. Finkbeiner, L. Haas, and H. Torfah, "Canonical representations of k-safety hyperproperties," in *IEEE 32nd Computer Security Foundations Symposium*, pp. 17–1714, 2019.

[26] N. Coenen, B. Finkbeiner, C. Sánchez, and L. Tentrup, "Verifying hyperliveness," in *Computer Aided Verification*, pp. 121–139, 2019.

[27] B. Finkbeiner and C. Hahn, "Deciding hyperproperties," in *27th International Conference on Concurrency Theory*, vol. 59 of *Leibniz International Proceedings in Informatics*, pp. 13:1–13:14, 2016.

[28] L. V. Nguyen, J. Kapinski, X. Jin, J. V. Deshmukh, and T. T. Johnson, "Hyperproperties of real-valued signals," in *15th ACM-IEEE International Conference on Formal Methods and Models for System Design*, p. 104–113, ACM, 2017.

[29] Y. Wang, M. Zarei, B. Bonakdarpour, and M. Pajic, "Statistical Verification of Hyperproperties for Cyber-Physical Systems," *ACM Trans. Embed. Comput. Syst.*, vol. 18, no. 5s, 2019.

[30] M. Anand, V. Murali, A. Trivedi, and M. Zamani, "Formal verification of hyperproperties for control systems," in *Proceedings of the Workshop on Computation-Aware Algorithmic Design for Cyber-Physical Systems*, Association for Computing Machinery, 2021.

[31] M. Y. Vardi and P. Wolper, "An automata-theoretic approach to automatic program verification," in *Proceedings of the First Symposium on Logic in Computer Science*, pp. 332–344, IEEE Computer Society, 1986.

[32] M. Y. Vardi, *An Automata-Theoretic Approach to Linear Temporal Logic*. Springer Berlin Heidelberg, 1996.

[33] E. Asarin, O. Maler, and A. Pnueli, "Reachability analysis of dynamical systems having piecewise-constant derivatives," *Theoretical Computer Science*, vol. 138, no. 1, pp. 35–65, 1995.

[34] P. Jagtap, A. Swikir, and M. Zamani, "Compositional construction of control barrier functions for interconnected control systems," in *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control*, HSCC '20, (New York, NY, USA), pp. 1–11, Association for Computing Machinery, 2020.

[35] K. Zhou and J. C. Doyle, *Essentials of robust control*. Prentice-Hall, 1998.

[36] K. Zhang, X. Yin, and M. Zamani, "Opacity of nondeterministic transition systems: A (bi)simulation relation approach," *IEEE Transactions on Automatic Control*, vol. 64, no. 12, pp. 5116–5123, 2019.

[37] V. Goranko, A. Kuusisto, and R. Rönnholm, "Game-theoretic semantics for alternating-time temporal logic," in *International Conference on Autonomous Agents & Multiagent Systems*, p. 671–679, 2016.

[38] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*. Prentice Hall, Dec. 2009.

[39] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, third edition. The MIT Press, 3rd edition ed., 2009.

[40] J. McLean, "A general theory of composition for trace sets closed under selective interleaving functions," in *Proceedings of 1994 IEEE Computer Society Symposium on Research in Security and Privacy*, pp. 79–93, 1994.

[41] J. Bochnak, M. Coste, and M.-F. Roy, *Real Algebraic Geometry*. Springer, 1998.

[42] P. A. Parrilo, "Semidefinite programming relaxations for semialgebraic problems," *Mathematical Programming*, vol. 96, no. 2, pp. 293–320, 2003.

[43] S. Prajna, A. Papachristodoulou, and P. Parrilo, "Introducing SOS-TOLLS: A general purpose sum of squares programming solver," in *Proceedings of the 41st IEEE Conference on Decision and Control*, vol. 1, pp. 741–746, 2002.

[44] J. F. Sturm, "Using sedumi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optimization methods and software*, vol. 11, no. 1-4, pp. 625–653, 1999.

[45] L. De Moura and N. Bjørner, "Z3: An efficient SMT solver," in *Proceedings of the International conference on Tools and Algorithms for the Construction and Analysis of Systems*, pp. 337–340, 2008.

[46] S. Gao, J. Avigad, and E. M. Clarke, " δ -complete decision procedures for satisfiability over the reals," in *Automated Reasoning*, LNCS, pp. 286–300, 2012.

[47] P. Jagtap, S. Soudjani, and M. Zamani, "Formal synthesis of stochastic systems via control barrier certificates," *IEEE Transactions on Automatic Control*, 2020.

[48] A. Peruffo, D. Ahmed, and A. Abate, "Automated and formal synthesis of neural barrier certificates for dynamical models," 2020.

[49] H. Zhao, X. Zeng, T. Chen, and Z. Liu, "Synthesizing barrier certificates using neural networks," in *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control*, 2020.

[50] S. Liu and M. Zamani, "Verification of approximate opacity via barrier certificates," *IEEE Control Systems Letters*, vol. 5, no. 4, pp. 1369–1374, 2021.

APPENDIX

A. Proof of Lemma 3.1

If Eloise has a positional winning strategy in the turn-based evaluation game, then she can use the same strategy to choose traces in the stage-based evaluation game such that each index depends only on the states and actions at the current index in the traces quantified so far. Then, against an arbitrary policy chosen by Abelard, the resulting p -tuple of traces satisfy the LTL specification ψ .

B. Proof of Lemma 3.6

We prove this by contradiction. Suppose an ABC exists for χ , but χ is not valid. Then, regardless of Eloise's strategy, Abelard always has a strategy that allows him to win. Let the set of traces selected by the players be $T = \{\sigma_1, \dots, \sigma_p\}$. For Abelard to win, he must ensure that for some $j \in \mathbb{N}$, $\Pi[j, \infty] \models_T s_A$ and for some $k > j$, $\Pi[k, \infty] \models_T s_B$ (from the existence of an ABC and Remark 3.5 we can conclude that $k \neq j$) to falsify $G(s_A \rightarrow G(\neg s_B))$. We consider the case where the strategy of Eloise is to select inputs according to condition (III.3). We note that selecting such a strategy leads to a non-increase in the value of the barrier certificate for the corresponding state in the augmented system, regardless of Abelard's strategy. Let the set of traces at positions j and k correspond to states \tilde{x} and \tilde{x}' in the augmented system and let the corresponding input sequence that takes us from \tilde{x} to \tilde{x}' be $\tilde{\nu}$. From conditions (III.1) and (III.2), we have $\mathcal{B}(\tilde{x}) \leq 0$ and $\mathcal{B}(\tilde{x}') > 0$. For any $l \geq 0$, let $\tilde{w} = \tilde{\nu}(l)$, and $\tilde{x}_l = \mathbf{x}(l)$, then we have $\mathcal{B}(f^p(\tilde{x}_l, \tilde{w})) \leq \mathcal{B}(\tilde{x}_l)$ from condition (III.3) regardless of Abelard's strategy. By induction on this condition, we can infer that $\mathcal{B}(\tilde{x}') \leq 0$. This is a contradiction to condition (III.2). So, we infer that χ is valid in the turn-based game setting. Therefore, χ is also valid in the stage-based game setting according to Lemma 3.1.

C. Proof of Lemma 3.7

From Lemma 3.6, for a conditional invariance $\chi_j = \mu_1 \pi_1 \dots \mu_p \pi_p G(s_{A_j} \rightarrow G(\neg s_{B_j}))$, existence of an ABC \mathcal{B}_j implies that Eloise has a winning strategy to ensure that $T(\mathfrak{S}, L) \models \chi_j$. Therefore, for a set of conditional invariances χ_1, \dots, χ_k , Eloise may choose the same winning strategy corresponding to χ_j to ensure that at least one of the conditional invariances in the set holds. Therefore, we get $T(\mathfrak{S}, L) \models \chi$, where $\chi = \mu_1 \pi_1 \dots \mu_p \pi_p \bigvee_{1 \leq j \leq k} G(s_{A_j} \rightarrow G(\neg s_{B_j}))$.

D. Proof of Lemma 3.8

The proof once again follows from Lemma 3.6. The existence of a common ABC \mathcal{B} guarantees that condition (III.3) is satisfied for all conditional invariances χ_i , $1 \leq i \leq k$. This implies that Eloise may use the same strategy to disallow all the transition pairs (s_{A_i}, s_{B_i}) , $1 \leq i \leq k$. Therefore, we have $T(\mathfrak{S}, L) \models \chi$, where $\chi = \mu_1 \pi_1 \dots \mu_p \pi_p \bigwedge_{1 \leq i \leq k} G(s_{A_i} \rightarrow G(\neg s_{B_i}))$

E. Proof of Lemma 3.9

From Lemma 3.7, for the set of conditional invariances $\{\chi_{i,1}, \dots, \chi_{i,v_i}\}$ for some $1 \leq i \leq k$, the existence of ABC \mathcal{B} for some $\chi_{i,j}$ implies that

$T(\mathfrak{S}, L) \models \mu_1 \pi_1 \dots \mu_p \pi_p \bigvee_{1 \leq j \leq v_i} G(s_{A_{i,j}} \rightarrow G(\neg s_{B_{i,j}}))$. For each $1 \leq i \leq k$, if there exists a common ABC \mathcal{B} for some $\chi_{i,j}$, $1 \leq j \leq v_i$, then by Lemma 3.8, we have that $T(\mathfrak{S}, L) \models \mu_1 \pi_1 \dots \mu_p \pi_p \bigwedge_{1 \leq i \leq k} G(s_{A_{i,j}} \rightarrow G(\neg s_{B_{i,j}}))$. By combining these two results, for a family of set of conditional invariances $\{\{\chi_{1,1}, \dots, \chi_{1,v_1}\}, \dots, \{\chi_{k,1}, \dots, \chi_{k,v_k}\}\}$, one has $T(\mathfrak{S}, L) \models \chi$, where $\chi = \mu_1 \pi_1 \dots \mu_p \pi_p \bigwedge_{(1 \leq i \leq k)(1 \leq j \leq v_i)} G(s_{A_{i,j}} \rightarrow G(\neg s_{B_{i,j}}))$.

F. Proof of Theorem 3.10

Let the NBA $\mathcal{A}_{\neg\psi}$ corresponding to $\neg\psi$ have k lassos, such that the i^{th} lasso has v_i pairs of consecutive transitions. Let the pair $(s_{A_{i,j}}, s_{B_{i,j}})$ correspond to the j^{th} pair of consecutive transitions along the i^{th} lasso. Let $\chi_{i,j} = \mu_1 \pi_1 \dots \mu_p \pi_p G(s_{A_{i,j}} \rightarrow G(\neg s_{B_{i,j}}))$ denote a conditional invariance specification and let us consider the set of conditional invariances $\{\chi_{1,1}, \dots, \chi_{1,v_1}, \chi_{2,1} \dots \chi_{2,v_2}, \dots, \chi_{k,1}, \dots, \chi_{k,v_k}\}$. Then the existence of a common ABC satisfying Lemma 3.9 implies that if the augmented system lands on a state satisfying $s_{A_{i,j}}$, Eloise has a strategy to ensure that it never reaches a state satisfying $s_{B_{i,j}}$ for every $1 \leq i \leq k$, and some $1 \leq j \leq v_j$. However, to satisfy $\neg\psi$, Abelard must have a strategy that allows him to visit a state satisfying some $s_{A_{i,j}}$ and then later visit a state satisfying $s_{B_{i,j}}$, for some $1 \leq i \leq k$ and every $1 \leq j \leq v_i$, to follow the transitions along the i^{th} lasso. Since this is not possible due to the existence of the ABC, we infer that ϕ is satisfied.

G. Proof of Algorithms 2 and 3

Algorithm 2 is a sound way of obtaining suitable common ABCs in the presence of two or more outgoing edges from any state in DBA $\mathcal{A}_{\neg\psi}$, and different ABCs otherwise. It calls on function *ABC_FIND*, presented in Algorithm 3, that takes a state r in $\mathcal{A}_{\neg\psi}$ as input and returns a set of lassos that pass through r and can be denied by means of a common ABC. The correctness of Algorithm 3 can be established by ensuring Eloise has a strategy to disallow all the lassos in \mathcal{R}_d . Initially, we have $\mathcal{R}_d = \emptyset$ and therefore Eloise may use any strategy. Now, at some iteration of the outermost loop, let \mathcal{R}_d consist of the lassos that have already been disallowed. Furthermore, consider an outgoing label $s_A \in \Sigma^p$ from r to r' . The first inner loop then identifies the labels s_B such that $\Delta'(r', s_B) \neq \emptyset$, and constructs a set S that consists of all such pairs (s_A, s_B) . The algorithm then finds a common ABC for the set $S_a \subseteq S$ and, by Lemma 3.8, we can infer that Eloise has a strategy of selecting inputs such that she can always avoid those states that satisfy s_B in the augmented system after satisfying s_A . Therefore, at any iteration of the second inner loop, Eloise has a strategy to avoid the lassos already present in \mathcal{R}_d , or the lassos corresponding to the pairs $(s_A, s_B) \in S_a$. Moreover, due to the specification being in the $\forall^* \exists^*$ fragment and the determinism in the automata, Eloise can uniquely determine the state of the augmented system as well as the current state in the automaton. This allows Eloise to select a unique strategy to avoid lassos r that pass through the

states (r, r', r'') such that $r' = \Delta'(r, s_A)$ and $r'' = \Delta'(r', s_B)$, and add them to \mathcal{R}_d . Lastly, we note that if we can successfully find common ABCs then the algorithm terminates as there are finitely many labels in Σ^p .

We now prove the correctness of Algorithm 2 by using the fact that \mathcal{R}_d consists of all the lassos that can be denied by Eloise. Initially, \mathcal{R}_d is empty, and therefore, Eloise may use any strategy. Now, at some iteration of the outermost loop, consider a state r and the set \mathcal{R}_d that consists of the lassos that have already been denied. On calling the function *ABC_FIND* for a state r , we find a common ABC for those lassos $\mathbf{r} \in \mathcal{R}$ such that $(r, r', r'') \in \mathbf{r}$, where $\Delta'(r, s_A) = r'$, $\Delta'(r', s_B) = r''$ and $(s_A, s_B) \in S_a$. Then, from the proof of Algorithm 3, we have that Eloise has a unique strategy to deny those lassos. If $\mathcal{R}_m = \emptyset$, then $\mathcal{R}_d = \mathcal{R}$. Thus all lassos that pass through some state in F have been denied and the states in F cannot be visited infinitely often. Then, we can infer that Eloise has a strategy to never satisfy $\neg\psi$ or that $T(\mathfrak{S}, L) \models \phi$. However, if we traverse all the states and $\mathcal{R}_m \neq \emptyset$, the algorithm is inconclusive.

H. Proof of Lemma 4.2

Since $\lambda_0(\tilde{x})$ is an SOS polynomial, we have that $\lambda_0(\tilde{x})g_0^T(\tilde{x})$ is non-negative over A . Therefore, if condition (IV.1) is an SOS polynomial, and therefore non-negative, it would directly imply condition (III.1). Similarly, the SOS constraint (IV.2) implies condition (III.2). Now we show that condition (IV.3) implies (III.3). By selecting inputs $w_{i_j} = h_j^i(\hat{x}_i, \hat{w})$, the last term in (IV.3) vanishes. Since the expression $\lambda(\tilde{x}, \tilde{w})g^T(\tilde{x})$ is non-negative over X^p and $\lambda_{in}(\tilde{x}, \tilde{w})g_{in}^T(\tilde{w})$ is non-negative over W^p , we have that for all $\tilde{x} \in X^p$, $-\mathcal{B}(f^p(\tilde{x}, \tilde{w})) + \mathcal{B}(\tilde{x}) \geq 0$. This implies that condition (III.3) holds, thus concluding the proof.

I. Suspect-Pair Decomposition

Our verification approach relies on reducing the complex HyperLTL specification into a collection of conditional invariance guarantees over consecutive transition pairs. This was originally proposed in [47] but in the context of verification for LTL specifications. Consider the NBA $\mathcal{A}_{\neg\psi} = (Q, q_0, \Sigma^p, \Delta, F)$ compiling the negation of the desired specification. A *lasso* for the NBA $\mathcal{A}_{\neg\psi}$ is a pair (\tilde{r}, \tilde{r}') such that $\tilde{r} = (r_0, r_1, \dots, r_{n-1}, r_n)$ represents a simple path, and $\tilde{r}' = (r'_0, r'_1, \dots, r'_{n'})$ represents a simple cycle, with $r_n = r'_0 = r'_{n'}$.

We define the set \mathcal{R} to be the set of all *lasso runs* that start at the initial state and cycle on an accepting state, i.e.,

$$\mathcal{R} := \{ \mathbf{r}_l = (r_0, \dots, r_{n-1}, r'_0, \dots, r'_{n'}) \mid r_0 = q_0, r'_0 = r'_{n'} \in F \}.$$

For $\tilde{a} \in \Sigma^p$, we define the set $\mathcal{R}_{\tilde{a}}$ as the subset of \mathcal{R} starting from \tilde{a} , i.e.

$$\mathcal{R}_{\tilde{a}} := \{ \mathbf{r}_l = (r_0, r_1, \dots, r_{n-1}, r'_0, \dots, r'_{n'}) \in \mathcal{R} \mid (r_0, \tilde{a}, r_1) \in \Delta, \tilde{a} \in \Sigma^p \}.$$

For each lasso run $\mathbf{r}_l = (r_0, r_1, \dots, r_{n+n'+1}) \in \mathcal{R}_{\tilde{a}}$, we define a set $\mathcal{P}_{\tilde{a}}(\mathbf{r}_l)$ consisting of runs of length 3 as

$$\mathcal{P}_{\tilde{a}}(\mathbf{r}_l) := \{ (r_i, r_{i+1}, r_{i+2}) \in \mathbf{r}_l \mid 0 \leq i \leq n + n' - 1 \}.$$

We correspondingly define the set $\mathcal{S}_{\tilde{a}}(\mathbf{r}_l)$ consisting of consecutive transition pairs $s_A, s_B \in \Sigma^p$ such that

$$\mathcal{S}_{\tilde{a}}(\mathbf{r}_l) = \{ (s_A, s_B) \mid (r, s_A, r') \in \Delta \text{ and } (r', s_B, r'') \in \Delta, (r, r', r'') \in \mathcal{P}_{\tilde{a}}(\mathbf{r}_l) \}.$$

We let the set $\mathcal{S}_{\mathcal{A}_b} = \bigcup_{\tilde{a} \in \Sigma^p} \bigcup_{\mathbf{r}_l \in \mathcal{R}_{\tilde{a}}} \mathcal{S}_{\tilde{a}}(\mathbf{r}_l)$ be the set of all such consecutive transition pairs obtained from NBA $\mathcal{A}_{\neg\psi}$. These transition pairs correspond to different conditional invariance specifications for which a suitable ABC is synthesized.



Mahathi Anand is a postdoctoral researcher at the Institute of Systems Theory and Automatic Control, University of Stuttgart, Germany since August 2023. Previously, In 2023, she obtained her doctoral degree in computer science from LMU Munich, Germany, with the thesis titled 'Formal Analysis of Control Systems via Inductive Approaches'. Prior to that, she graduated with a Bachelors in Technology in Electrical and Electronics Engineering from SRM Institute of Science and Technology, India in 2016 and completed her Masters of Technology in Electrical Engineering with specialization in System and Control from Indian Institute of Technology Roorkee, India in 2019. She also spent some time at the Hybrid Control Systems Lab in Technical University of Munich to work on her Masters thesis. Her research interests mainly include formal methods in control with a focus in nonlinear, stochastic, and learning-enabled systems.



Vishnu Murali is a PhD Student in the Department of Computer Science at the University of Colorado Boulder, USA since January 2021. He received a B.E. degree in Computer Science and Engineering from Visvesvaraya Technological University, India in 2019 and an M.S. degree in Computer Science at the University of Colorado Boulder, USA in 2023. His research interests are mainly in formal methods with a focus on logic, automata and cyber-physical systems.



Ashutosh Trivedi is an Associate Professor of Computer Science at the University of Colorado Boulder, specializing in the development of formal methods for safety-critical learning-enabled systems. He obtained his doctoral degree in computer science from the University of Warwick, with a specialization in game theory and optimization. Prior to his current position, Ashutosh held the position of Assistant Professor at the IIT Bombay. Before that, he worked as a Postdoctoral Research Associate at the University of Pennsylvania and the University of Oxford.

Ashutosh is a recipient of the 2022 NSF CAREER award and a Liverpool fellowship. His research interests lie at the intersection of computer science, control theory, and machine learning, with a focus on formal methods, reinforcement learning, and software security.



Majid Zamani (M'12–SM'16) is an Associate Professor in the Computer Science Department at the University of Colorado Boulder, USA. He is also a guest professor in the Computer Science Department at the Ludwig Maximilian University of Munich. He received a B.Sc. degree in Electrical Engineering in 2005 from Isfahan University of Technology, Iran, an M.Sc. degree in Electrical Engineering in 2007 from Sharif University of Technology, Iran, an MA degree in Mathematics and a Ph.D. degree in Electrical Engineering

both in 2012 from University of California, Los Angeles, USA. Between September 2012 and December 2013, he was a postdoctoral researcher at the Delft Center for Systems and Control, Delft University of Technology, Netherlands. From May 2014 to January 2019, he was an Assistant Professor in the Department of Electrical and Computer Engineering at the Technical University of Munich, Germany. From December 2013 to April 2014, he was an Assistant Professor in the Design Engineering Department, Delft University of Technology, Netherlands. He received the George S. Axelby Outstanding Paper Award from the IEEE Control Systems Society in 2023, the NSF Career award in 2022 and the ERC Starting Grant and Proof of Concept Grant from the European Research Council in 2018 and 2023, respectively.

His research interests include verification and control of hybrid systems, embedded control software synthesis, networked control systems, and incremental properties of nonlinear control systems.