

Specification-guided Verification and Abstraction Refinement of Mixed Monotone Stochastic Systems

Maxence Dutreix and Samuel Coogan

Abstract—This paper addresses the problem of verifying discrete-time stochastic systems against omega-regular specifications using finite-state abstractions. Omega-regular properties allow specifying complex behavior and encompass, for example, linear temporal logic. We focus on a class of systems with mixed monotone dynamics. This class has been shown to be amenable to efficient reachable set computation and models a wide-range of physically relevant systems. In general, finite-state abstractions of continuous state stochastic systems give rise to augmented Markov Chains wherein the probabilities of transition between states are restricted to an interval. We present a procedure to compute a finite-state Interval-valued Markov Chain abstraction of discrete-time, mixed monotone stochastic systems subject to affine disturbances given a rectangular partition of the state-space. Then, we suggest an algorithm for performing verification against omega-regular properties in IMCs. Specifically, we aim to compute bounds on the probability of satisfying a specification from any initial state in the IMC. This is achieved by solving a reachability problem on sets of so-called winning and losing components in the Cartesian product between the IMC and a Rabin automaton representing the specification. Next, the verification of IMCs may yield a set of states whose acceptance status is undecided with respect to the specification, requiring a refinement of the abstraction. We describe a specification-guided approach that compares the best and worst-case behaviors of accepting paths in the IMC and targets the appropriate states accordingly. Finally, we show a case study.

Index Terms—finite-state abstractions, interval-valued Markov chains, mixed monotone systems, stochastic systems.

I. INTRODUCTION

Recent years have witnessed a growing effort to design tools for verifying and optimizing the behavior of systems with respect to increasingly complex specifications. Many efficient verification techniques can be readily applied in the context of purely deterministic models. Yet, the study of stochastic systems generally requires a machinery of its own.

Manuscript submitted on ...

This project was supported in part by the NSF under project #1749357.

M. Dutreix is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA (e-mail: max-dutreix@gatech.edu).

S. Coogan is with the School of Electrical and Computer Engineering and the School of Civil and Environmental Engineering, Georgia Institute of Technology, Atlanta, GA, USA (e-mail: sam.coogan@gatech.edu).

For example, the model checking problem for discrete-time and continuous-time *Markov Chains* (MC) has been solved for the *Probabilistic Computation Tree Logic* (PCTL) and ω -regular properties [1] [2] [3]. Numerous off-the-shelf tools, such as PRISM [4], can efficiently verify MCs for a wide range of specifications.

However, these MC models assume that the probabilities of transition between states are known exactly. While computationally convenient, such representations may not be realistic in practice. In particular, the verification of continuous state systems is often performed by partitioning its domain into a discrete set of states. Because these states abstract a collection of continuous behaviors simultaneously, modeling a stochastic finite-state abstraction as a transition system can be challenging. One approach uses approximate Markov chains [5] as in the FAUST² model checker [6], where each discrete state is reduced to a single representative point. Another approach uses Markov set-chains [7] to abstract the evolution of such systems so that existing techniques are applicable. Similar automata-based techniques are found in [8] and [9] for linear time objectives. However, these techniques generally rely on a fine gridding of the continuous state-space which overlooks qualitative aspects that are important to certain specifications, e.g. the creation of absorbing states as pointed out in [8]. Abstraction-based methods have been employed as well for the verification of continuous-time stochastic hybrid models against safety specifications in [10] and [11].

Instead of approximate Markov chains, finite-state abstractions of continuous state-space, discrete-time stochastic systems are also amenable to Markovian models where the probabilities of transitions are restricted to some interval [12] [13], referred to as *Interval-valued Markov Chains* (IMC) [14]. We consider this approach in this work. Constructing interval-valued abstractions of stochastic systems with continuous state-spaces is often a computationally expensive process. Indeed, in the general case, calculating the exact lower and upper bounds of transitions between all discrete states involves numerical searches over the state-space, rendering this procedure highly time-inefficient. Nevertheless, we aim to show that this impractical computational blowup can be avoided in some cases by exploiting the inherent structure of the system's dynamics. Our proposed approach results in an IMC abstraction with conservative transition probability ranges, which remains sufficient for verification.

In particular, we consider a class of stochastic systems for

which the dynamics exhibits a *mixed monotone* property [15], [16]. Mixed monotonicity generalizes the property of monotonicity for dynamical systems for which trajectories maintain a partial ordering on states [17], [18], [19]. Many physical systems have been shown to be monotone or mixed monotone such as biological systems [20] and transportation networks [21], [22], [23]. This property enables efficient computation of reachable sets: a rectangular over-approximation of the one-step reachable set from any rectangular discrete state is determined by evaluating a certain decomposition function at the least and the greatest point of that state. In this paper, we study mixed monotone systems that are subject to an affine random disturbance vector whose components are mutually independent and for which the probability distribution for each component is unimodal and symmetric. For such systems, we show that an upper bound and a lower bound on the probability of transitions between states of a rectangular partition are found by evaluating only two integrals per dimension and per transition. We use these bounds to create an IMC abstraction of the original system that is suitable for verification.

Efficient verification algorithms for IMCs have been derived for the logic PCTL [12]. However, many specifications of interest cannot be expressed using PCTL, such as liveness, i.e., the infinitely repeated occurrence of an event [24]. This motivates the implementation of a machinery for handling such properties for the class of abstractions mentioned above. As a superset of the broadly used *Linear Temporal Logic* (LTL), ω -regular properties exhibit considerable expressiveness.

In this work, we treat the problem of verification for discrete-time, continuous-space stochastic systems against ω -regular properties. To this end, we use finite-state abstractions in the form of IMCs. We develop an algorithm for determining the best and worst-case outcome of the IMC by solving a reachability problem in the product between the IMC and a Rabin automaton corresponding to the specification of interest. Our approach can be decomposed into a graph search which identifies the *largest* and *permanent* so-called *winning* and *losing components* [25] in the product IMC, and the computation of bounds on the probability of reaching these components.

Such verification techniques may yield a set of states whose satisfiability status with respect to the desired specification is undecided. Previous works suggested a methodology to compute a gridding parameter that guarantees an upper bound on the size of the interval of satisfaction for all states in the resulting partition and for all specifications in the logic PCTL [5]. This parameter is a function of the system's properties only, such as Lipschitz constants. However, this method is often conservative and likely to provide very fine partitions and is therefore computationally expensive. Instead, we apply a specification-guided refinement procedure on an initial crude partition of the state-space and iteratively produce finer IMC abstractions until some precision criterion is attained. Verification is performed on each new partition and a specific set of states is selected for the next refinement step. Refinement heuristics were proposed for the PCTL framework in [12] and [13]. We present a technique for the ω -regular framework that accounts for the behavior of the paths generated by the best and worst case adversary of the IMC to refine regions of the

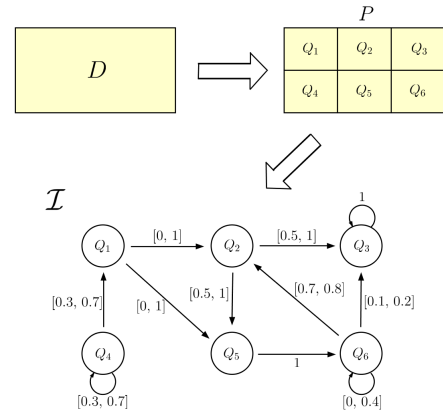


Fig. 1. A finite-state IMC abstraction \mathcal{I} over a continuous domain D . A partition P of D is generated and bounds on the transition probabilities between states are estimated.

state-space with the highest potential of reducing uncertainty.

In summary, the first contribution of this paper is an efficient procedure for constructing an IMC abstraction of stochastic mixed monotone systems subject to affine, unimodal and symmetrical disturbances. The second contribution is a technique for computing satisfiability bounds on ω -regular specifications in IMCs. The third contribution is a specification-guided refinement method that selectively and iteratively partitions certain regions of the continuous domain of stochastic systems for the purpose of verification against ω -regular properties. This technique improves on other specification-guided approaches by inspecting the accepting paths produced by the worst and best case adversary of the IMC instead of observing only the one-step transitions of each state and by considering the qualitative structure of the product IMC through its largest and permanent components to target the states accordingly.

The paper is organized as follows: Section II introduces preliminaries; Section III presents the problems investigated in this paper; Section IV derives an IMC abstraction techniques for affine-in-disturbance mixed monotone systems; Section V discusses the verification of IMCs against ω -regular properties; Section VI discusses state-space refinement; Section VII contains case studies; Section VIII summarizes the conclusions.

Preliminary results were reported in the conference papers [13] and [26]. The verification approach considered here is a significant improvement of these prior works. Specifically, the methodology proposed here does not require computing the complement of the system specification and presents a more advanced algorithm for partition refinement.

II. PRELIMINARIES

A *Deterministic Rabin Automaton (DRA)* [1] is a 5-tuple $\mathcal{A} = (S, \Sigma, \delta, s_0, Acc)$ where:

- S is a finite set of states,
- Σ is an alphabet,
- $\delta : S \times \Sigma \rightarrow S$ is a transition function,
- s_0 is an initial state,
- $Acc \subseteq 2^S \times 2^S$. An element $(E_i, F_i) \in Acc$, with $E_i, F_i \subset S$, is called a *Rabin Pair*.

A DRA \mathcal{A} reads an infinite string or *word* over alphabet Σ as an input and transitions from state to state according to δ .

The resulting sequence of states or *run* is an *accepting* run if it contains an infinite number of states belonging to F_i and a finite number of states in E_i for some i . A word is said to be accepted by \mathcal{A} if it produces an accepting run in \mathcal{A} . We call a set of words a *property*. The property *accepted* by \mathcal{A} is the set of all words accepted by \mathcal{A} .

A property over an alphabet Σ is ω -regular if and only if it is accepted by a Rabin Automaton with alphabet Σ (for more detailed definitions of ω -regular properties, see [1, Section 4.3.1]). In particular, all properties defined by a *Linear Temporal Logic* (LTL) formula are ω -regular. See [1] for a description of the semantics of LTL.

Throughout, all inequalities are interpreted elementwise so that, for $x, y \in \mathbb{R}^n$, $x \leq y$ means $x_i \leq y_i$ for $i = 1, \dots, n$, and similarly for $\geq, <$ and $>$.

$Q_j = \{x : a^j \leq x \leq b^j\}$ for some $a^j, b^j \in \mathbb{R}^n$ such that $a^j \leq b^j$ is a *compact rectangular set* and a^j, b^j are respectively called the *least point* and the *greatest point* of Q_j . For vectors, we reserve the subscript to index elements of the vector so that, e.g., a_i^j for $i \in \{1, \dots, n\}$ denotes the i -th element of $a^j \in \mathbb{R}^n$.

III. PROBLEM FORMULATION

We consider the discrete-time, continuous-state stochastic system

$$x[k+1] = \mathcal{F}(x[k], w[k]) \quad (1)$$

where $x[k] \in D \subset \mathbb{R}^n$ is the state of the system at time k , $w[k] \in W \subset \mathbb{R}^p$ is a random disturbance and $\mathcal{F} : D \times W \rightarrow D$ is a continuous map. At each time-step k , the random disturbance $w[k]$ is sampled from a probability distribution with density function $f_w : \mathbb{R}^p \rightarrow \mathbb{R}_{\geq 0}$ satisfying $f_w(z) = 0$ if $z \notin W$. Let $L : D \rightarrow \Sigma$ be a labeling function, where Σ is a finite alphabet. A random path $x[0]x[1] \dots$ satisfying (1) generates the word $L(x[0])L(x[1]) \dots$ over Σ . We assume throughout this work that the set of initial states of (1) is D .

We denote by Ψ an ω -regular property over alphabet Σ and define a probability operator $\mathcal{P}_{\bowtie p_{sat}}[\Psi]$ over ω -regular properties, with $\bowtie \in \{\leq, <, \geq, >\}$, $p_{sat} \in [0, 1]$. For any initial state $x \in D$, we define the satisfaction relation \models where

$$x \models \mathcal{P}_{\bowtie p_{sat}}[\Psi] \Leftrightarrow p_{\Psi}^x \bowtie p_{sat} ,$$

with p_{Ψ}^x being the probability that the word generated by a random path starting in x satisfies property Ψ (for a rigorous formalization, see, e.g., [5]). In this paper, we concentrate on formulas of the type

$$\phi = \mathcal{P}_{\bowtie p_{sat}}[\Psi] . \quad (2)$$

Our objective is to sort all initial states of system (1) into those that satisfy and those that do not satisfy specification (2), denoted by the sets Q_{ϕ}^{yes} and Q_{ϕ}^{no} respectively.

Problem: *Given a system of the form (1), find the sets of initial states $Q_{\phi}^{yes} \subseteq D$ and $Q_{\phi}^{no} \subseteq D$ that respectively satisfy and do not satisfy a formula ϕ of the form (2).*

The domain D of system (1) generally contains an uncountably infinite number of states and obtaining exact solutions to this problem may be infeasible for rich specifications. A common approximation approach consists in partitioning D into a finite collection of states P to obtain a finite abstraction of the stochastic dynamics. In this paper, we only consider partitions which are rectangular.

Definition 1 (Rectangular Partition): A *rectangular partition* P of the compact domain $D \subset \mathbb{R}^n$ is a collection of discrete states $P = \{Q_j\}_{j=1}^m$, $Q_j \subset D$, satisfying

- Q_j is a compact rectangular set $\forall j = 1, \dots, m$,
- $\bigcup_{j=1}^m Q_j = D$,
- $\text{int}(Q_j) \cap \text{int}(Q_\ell) = \emptyset \quad \forall j, \ell, j \neq \ell$,

where int denotes the interior. For any continuous state x belonging to a state Q_j , we write $x \in Q_j$.

Henceforth, we assume that D in system (1) is amenable to a rectangular partition.

Given a rectangular partition P of the domain D of (1), the likelihood of transitioning from a state Q_j of P to another state Q_ℓ generally varies with the continuous state abstracted by Q_j from which the transition is actually taking place. This prevents using the partition P to uniquely abstract (1) as a finite discrete-time MC. Instead, we produce an IMC abstraction of the system where the transition probabilities between states are constrained within some bounds. Fig. 1 depicts a schematic of a partition-based finite abstraction.

Definition 2 (Interval-Valued Markov Chain): An *Interval-Valued Markov Chain (IMC)* [13] is a 6-tuple $\mathcal{I} = (Q, \check{T}, \hat{T}, q_0, \Sigma, L)$ where:

- Q is a finite set of states,
- $\check{T} : Q \times Q \rightarrow [0, 1]$ maps pairs of states to a lower transition bound so that $\check{T}_{Q_j \rightarrow Q_\ell} := \check{T}(Q_j, Q_\ell)$ denotes the lower bound of the transition probability from state Q_j to state Q_ℓ , and
- $\hat{T} : Q \times Q \rightarrow [0, 1]$ maps pairs of states to an upper transition bound so that $\hat{T}_{Q_j \rightarrow Q_\ell} := \hat{T}(Q_j, Q_\ell)$ denotes the upper bound of the transition probability from state Q_j to state Q_ℓ ,
- $q_0 \subseteq Q$ is a set of initial states,
- Σ is a finite set of atomic propositions,
- $L : Q \rightarrow \Sigma$ is a labeling function from states to Σ ,

and \check{T} and \hat{T} satisfy $\check{T}(Q_j, Q_\ell) \leq \hat{T}(Q_j, Q_\ell)$ for all $Q_j, Q_\ell \in Q$ and

$$\sum_{Q_\ell \in Q} \check{T}(Q_j, Q_\ell) \leq 1 \leq \sum_{Q_\ell \in Q} \hat{T}(Q_j, Q_\ell) \quad (3)$$

for all $Q_j \in Q$.

Definition 3 (Markov Chain): A *Markov Chain (MC)* $\mathcal{M} = (Q, T, q_0, \Sigma, L)$ is defined similarly to an IMC with the difference that the transition probability function or transition matrix $T : Q \times Q \rightarrow [0, 1]$ satisfies $0 \leq T(Q_j, Q_\ell) \leq 1$ for all $Q_j, Q_\ell \in Q$ and $\sum_{Q_\ell \in Q} T(Q_j, Q_\ell) = 1$ for all $Q_j \in Q$.

A Markov Chain $\mathcal{M} = (Q, T, q_0, \Sigma, L)$ is said to be *induced* by IMC $\mathcal{I} = (Q, \tilde{T}, \hat{T}, q_0, \Sigma, L)$ if they share the same Q , q_0 , Σ and L , and for all $Q_j, Q_\ell \in Q$, $\tilde{T}(Q_j, Q_\ell) \leq T(Q_j, Q_\ell) \leq \hat{T}(Q_j, Q_\ell)$. A transition matrix T satisfying this inequality is said to be induced by \mathcal{I} .

The notation $\mathcal{P}_{\mathcal{M}}(Q_i \models \diamond U)$ for $U \subseteq Q$ denotes the probability of eventually reaching set U from initial state Q_i in a MC \mathcal{M} . With a slight abuse of notation, when the probability of reaching U is the same from all states in a set G , we write it as $\mathcal{P}_{\mathcal{M}}(G \models \diamond U)$.

An IMC \mathcal{I} is interpreted as an *Interval Markov Decision Process* (IMDP) [27] if, at each time step k , the environment non-deterministically chooses a transition matrix T_k induced by \mathcal{I} and the next transition occurs according to T_k . A mapping ν from a finite path $\pi = q_0 q_1 \dots q_k$ in \mathcal{I} to a transition matrix T_k is called an *adversary*. The set of all adversaries of \mathcal{I} is denoted by $\nu_{\mathcal{I}}$.

The probability of satisfying ω -regular property Ψ starting from initial state Q_i in IMC \mathcal{I} under adversary ν is denoted by $\mathcal{P}_{\mathcal{I}[\nu]}(Q_i \models \Psi)$. The greatest lower bound and least upper bound on the probability of satisfying property Ψ starting from initial state Q_i in IMC \mathcal{I} are denoted by $\tilde{\mathcal{P}}_{\mathcal{I}}(Q_i \models \Psi) = \inf_{\nu \in \nu_{\mathcal{I}}} \mathcal{P}_{\mathcal{I}[\nu]}(Q_i \models \Psi)$ and $\hat{\mathcal{P}}_{\mathcal{I}}(Q_i \models \Psi) = \sup_{\nu \in \nu_{\mathcal{I}}} \mathcal{P}_{\mathcal{I}[\nu]}(Q_i \models \Psi)$ respectively.

Definition 4 (IMC Abstraction): Given the system (1) evolving on a domain $D \subset \mathbb{R}^n$ and a partition $P = \{Q_j\}_{j=1}^m$ of D , an IMC $\mathcal{I} = (Q, \tilde{T}, \hat{T}, q_0, \Sigma, L)$ is an *abstraction* of (1) if:

- $P = Q$, i.e., the set of states of the IMC is the partition P ,
- For all $Q_j, Q_\ell \in P$,

$$\tilde{T}_{Q_j \rightarrow Q_\ell} \leq \inf_{x \in Q_j} Pr(\mathcal{F}(x, w) \in Q_\ell), \text{ and} \quad (4)$$

$$\hat{T}_{Q_j \rightarrow Q_\ell} \geq \sup_{x \in Q_j} Pr(\mathcal{F}(x, w) \in Q_\ell), \quad (5)$$

where $Pr(\mathcal{F}(x, w) \in Q_\ell)$ for fixed x is the probability that (1) transitions from x to a state $x' = \mathcal{F}(x, w)$ in Q_ℓ ,

- $P = q_0$, i.e., the set of initial states of the IMC is the partition P ,
- For all $Q_j \in P$ and for any two states $x_i, x_\ell \in Q_j$, it holds that $L(Q_j) := L(x_i) = L(x_\ell)$, that is, the partition respects the boundaries induced by the labeling function,
- \mathcal{I} is interpreted as an IMDP.

The fact that two continuous states within the same discrete state of the abstraction may engender different transition probabilities is captured by interpreting the IMC as an IMDP. One can always build a trivial IMC abstraction where all transitions are set to range from 0 to 1. The ease of finding tighter bounds on the transitions between states is dictated by the dynamics of the system and the geometry of the states in the partition. A contribution of this paper is to show that *affine-in-disturbance stochastic mixed monotone systems* evolving on a domain amenable to a rectangular partition can

be efficiently abstracted by IMCs. A formal definition of such systems is provided in the next section.

Subproblem 1: *Given an affine-in-disturbance stochastic mixed monotone system, construct a non-trivial IMC abstraction from a rectangular partition of its domain.*

Performing verification on an IMC abstraction provides probabilistic guarantees with respect to the original system's states. A consequence of model checking an IMC \mathcal{I} is that the probability of satisfying property Ψ from any of its initial states Q_j must be specified as an interval $I_j = [p_{min}^j, p_{max}^j]$, where $\mathcal{P}_{\mathcal{I}[\nu]}(Q_j \models \Psi) \in I_j$, $\forall \nu \in \nu_{\mathcal{I}}$. For any initial state Q_j in an IMC \mathcal{I} , we define the satisfaction relation \models for formulas of the type (2) where

$$Q_j \models \mathcal{P}_{\bowtie p_{sat}}[\Psi] \Leftrightarrow (p_{\Psi}^{Q_j})_{\nu} \bowtie p_{sat} \quad \forall \nu \in \nu_{\mathcal{I}},$$

with $(p_{\Psi}^{Q_j})_{\nu}$ being the probability that the word generated by a random path starting in Q_j satisfies property Ψ under adversary ν . We denote the set of initial states satisfying ϕ in \mathcal{I} by $(Q_{\phi}^{yes})_{\mathcal{I}}$, while states that do not satisfy ϕ are in $(Q_{\phi}^{no})_{\mathcal{I}}$. Note that any Q_j such that $p_{sat} \in]p_{min}^j, p_{max}^j[$ if $\bowtie \in \{\leq, \geq\}$, or $p_{sat} \in [p_{min}^j, p_{max}^j]$ if $\bowtie \in \{<, >\}$, is undecided with respect to ϕ in (2) and we write $Q_j \in (Q_{\phi}^?)_{\mathcal{I}}$. The remaining states either satisfy ϕ or do not satisfy ϕ .

Fact 1: Let \mathcal{I} be an IMC abstraction of (1) induced by a partition $P = \{Q_j\}_{j=1}^m$ of D . For any formula of the form (2), it holds that:

- $Q_j \in (Q_{\phi}^{yes})_{\mathcal{I}} \Rightarrow x \in Q_{\phi}^{yes} \quad \forall x \in Q_j$
- $Q_j \in (Q_{\phi}^{no})_{\mathcal{I}} \Rightarrow x \in Q_{\phi}^{no} \quad \forall x \in Q_j$.

Given an IMC abstraction \mathcal{I} of (1) generated from a partition P of D , our approach for addressing the verification problem is thus to implement a technique for determining non-trivial values of p_{min}^j and p_{max}^j and sort all states of P into the sets $(Q_{\phi}^{yes})_{\mathcal{I}}$, $(Q_{\phi}^{no})_{\mathcal{I}}$ and $(Q_{\phi}^?)_{\mathcal{I}}$.

Subproblem 2: *Given an IMC \mathcal{I} interpreted as an IMDP and an ω -regular property Ψ , find the greatest lower bound $\tilde{\mathcal{P}}_{\mathcal{I}}(Q_j \models \Psi)$ and the least upper bound $\hat{\mathcal{P}}_{\mathcal{I}}(Q_j \models \Psi)$ on the probability of satisfying Ψ from any initial state Q_j in \mathcal{I} .*

It may be the case that the total volume of uncertain states in $(Q_{\phi}^?)_{\mathcal{I}}$ is unsatisfactorily large due to a crude partition P of the continuous state-space. In such event, we aim to produce a finer partition P' from P which gives rise to a new IMC abstraction of the system with a greater number of states with tighter transition bounds between them. Verifying this finer IMC may yield a decreased volume of uncertain states. This refinement procedure is applied to subsequent partitions until some discretionary level of precision is reached.

Subproblem 3: *Given a system of the form (1) with an IMC abstraction \mathcal{I} of (1) and the corresponding sets $(Q_{\phi}^{yes})_{\mathcal{I}}$, $(Q_{\phi}^{no})_{\mathcal{I}}$ and $(Q_{\phi}^?)_{\mathcal{I}}$ obtained by solving Subproblem 2, refine the partition*

P of D until a predefined threshold of precision has been met.

In order to avoid state-space explosion, a careful choice of the states to be refined in partition P has to be made. In this work, we suggest a novel method for targeting the states that have the greatest potential of reducing the uncertainty of an abstraction with respect to a particular specification Ψ .

IV. IMC ABSTRACTION OF MIXED MONOTONE SYSTEMS AFFINE IN DISTURBANCE

In this section, we study a large class of stochastic systems that proves amenable to efficient computation. In particular, we consider affine-in-disturbance systems of the form

$$x[k+1] = \mathcal{F}(x[k]) + w[k] \quad (6)$$

with specific assumptions on \mathcal{F} and the distribution of the disturbance w . We introduce several definitions and specify these necessary assumptions for addressing Subproblem 1.

Definition 5 (Mixed monotone function): A function $\mathcal{F} : D \rightarrow D$ is *mixed monotone* if there exists a *decomposition function* $g : D \times D \rightarrow D$ satisfying [15], [16]:

- $\forall x \in D : \mathcal{F}(x) = g(x, x)$
- $\forall x^1, x^2, y \in D : x^1 \leq x^2$ implies $g(x^1, y) \leq g(x^2, y)$
- $\forall x, y^1, y^2 \in D : y^1 \leq y^2$ implies $g(x, y^2) \leq g(x, y^1)$

Mixed monotonicity generalizes the notion of monotonicity in dynamical systems, which is recovered when $g(x, y) = \mathcal{F}(x)$ for all x, y . Systems with monotone state update maps exhibit considerable structure useful for analysis [18], [19], [28]. Systems with mixed monotone state update maps have been shown to enjoy many of these same structural properties [15], [16]. For example, for mixed monotone \mathcal{F} with decomposition function g , for $x, y, z \in D$ satisfying $x \leq z \leq y$, we have $g(x, y) \leq \mathcal{F}(z) \leq g(y, x)$. This leads to the following proposition.

Proposition 1 ([16, Theorem 1]): Let $\mathcal{F} : D \rightarrow D$ be mixed monotone with decomposition function $g : D \times D \rightarrow D$, and let $a, b \in D$ satisfy $a \leq b$. Then

$$\{\mathcal{F}(x) : a \leq x \leq b\} \subseteq \{z : g(a, b) \leq z \leq g(b, a)\}.$$

Proposition 1 implies that the one-step reachable set from the rectangular region bounded between a and b is overapproximated by the rectangular region bounded by the two points $g(a, b)$ and $g(b, a)$. This property will prove key for efficient computation of IMC abstractions.

Definition 6 (Unimodal symmetric distribution): For a random disturbance $\omega \in \Omega \subset \mathbb{R}$ with Ω an interval, its probability density function $f_\omega : \mathbb{R} \rightarrow \mathbb{R}$ is *unimodal* if f_ω is differentiable on Ω and there exists a unique number $c \in \mathbb{R}$, referred as the mode of the distribution, such that, for $x \in \Omega$:

- $x < c \Rightarrow f'_\omega(x) \geq 0$,
- $x = c \Rightarrow f'_\omega(x) = 0$, and
- $x > c \Rightarrow f'_\omega(x) \leq 0$.

We only consider distributions without a “flat” peak, that is, unimodal distributions with a unique mode c . The probability density function f_ω is *symmetric* if there exists a number $d \in \mathbb{R}$ such that $f_\omega(d - x) = f_\omega(d + x)$ for all x .

Note that if f_ω is unimodal with mode c and symmetric, then it must be that $f_\omega(c - x) = f_\omega(c + x)$.

Henceforth, we make the following assumptions.

Assumption 1: $\mathcal{F}(x)$ in (6) is mixed monotone with decomposition function $g(x, y)$. Furthermore, the domain D of (6) can be partitioned into a rectangular partition.

Assumption 2: The random disturbance $w[k]$ in (6) is of the form $w[k] = [w_1[k] \ w_2[k] \ \dots \ w_n[k]]^T$, where each $w_i \in W_i \subset \mathbb{R}$ has probability density function $f_{w_i}(x_i)$, W_i is an interval, and the collection $\{w_i\}_{i=1}^n$ is mutually independent. Denote by $F_{w_i}(x) = \int_{-\infty}^x f_{w_i}(\sigma) d\sigma$ the cumulative distribution function for w_i .

Assumption 3: The probability density function f_{w_i} for each random variable w_i is symmetric and unimodal with mode c_i .

We now address Subproblem 1 for systems of the form (6) under Assumptions 1-3. We decompose our procedure for bounding the transition probability from a state Q_1 to a state Q_2 in two steps: first, we compute the rectangular over-approximation of the \mathcal{F} -reachable set from state Q_1 by taking advantage of the mixed monotonicity property. Next, we determine the positions of f_w within this rectangular region that respectively minimize and maximize its overlap with Q_2 . In the next section, we exploit the characteristics of w previously evoked to obtain an efficient computational procedure for computing these extremum points.

Proposition 2: Consider system (6) under Assumptions 1–3. Let $Q_1 = \{x : a^1 \leq x \leq b^1\}$ and $Q_2 = \{x : a^2 \leq x \leq b^2\}$ be two nonempty rectangular sets with least point a^j and greatest point b^j for $j = 1, 2$. Then

$$\begin{aligned} & \min_{x \in Q_1} Pr(\mathcal{F}(x) + w \in Q_2) \\ & \geq \prod_{i=1}^n \min_{z_i \in [g_i(a^1, b^1), g_i(b^1, a^1)]} \int_{a_i^2}^{b_i^2} f_{w_i}(x - z_i) dx \end{aligned} \quad (7)$$

and

$$\begin{aligned} & \max_{x \in Q_1} Pr(\mathcal{F}(x) + w \in Q_2) \\ & \leq \prod_{i=1}^n \max_{z_i \in [g_i(a^1, b^1), g_i(b^1, a^1)]} \int_{a_i^2}^{b_i^2} f_{w_i}(x - z_i) dx \end{aligned} \quad (8)$$

where g_i denotes the i -th element of $g(x, y)$, the decomposition function of \mathcal{F} .

All proofs are found in the Appendix. Before generalizing to higher dimensions, we treat a 1-dimensional version of our original problem. In Lemma 1, we prove that for a fixed interval $[a, b] \subset \mathbb{R}$, there exists a unique position for

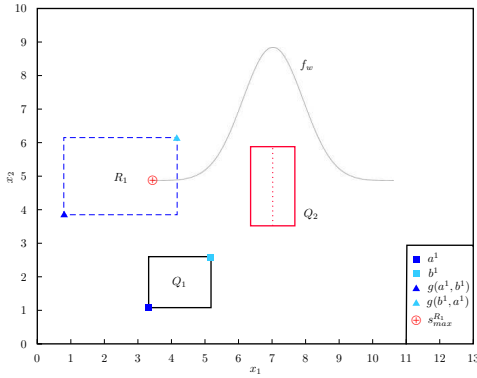


Fig. 2. Depiction of the procedure for computing an upper bound on the probability of transition from Q_1 to Q_2 . First, the one-step reachable set R_1 from Q_1 is over-approximated by evaluating the decomposition function at only two extremal points. Then, the distribution of $z + w$ is positioned as close to the center of Q_2 as possible under the restriction that $z \in R_1$. A lower bound on the transition probability is achieved by positioning the distribution as far from the center of Q_2 as possible.

a unimodal and symmetric distribution which maximizes its integral over $[a, b]$.

Lemma 1: Let $\omega \in \Omega \subset \mathbb{R}$ with Ω an interval be a random variable with symmetric and unimodal probability density function $f_\omega : \mathbb{R} \rightarrow \mathbb{R}$ and mode $c \in \mathbb{R}$. For any $a, b \in \mathbb{R}$ satisfying $a \leq b$ and any $r_1, r_2 \in \mathbb{R}$ satisfying $r_1 \leq r_2$, let

$$s_{max} = \frac{a + b}{2} - c$$

and define

$$s_{max}^r = \arg \min_{s \in [r_1, r_2]} |s_{max} - s| = \begin{cases} s_{max}, & \text{if } s_{max} \in [r_1, r_2] \\ r_2, & \text{if } s_{max} > r_2 \\ r_1, & \text{if } s_{max} < r_1, \end{cases}$$

$$s_{min}^r = \arg \max_{s \in [r_1, r_2]} |s_{max} - s| = \begin{cases} r_1, & \text{if } s_{max} > \frac{r_1 + r_2}{2} \\ r_2, & \text{otherwise.} \end{cases}$$

Then

$$\max_{s \in [r_1, r_2]} \int_a^b f_\omega(x - s) dx = \int_a^b f_\omega(x - s_{max}^r) dx \quad (9)$$

$$\min_{s \in [r_1, r_2]} \int_a^b f_\omega(x - s) dx = \int_a^b f_\omega(x - s_{min}^r) dx. \quad (10)$$

When $s_{max} \in [r_1, r_2]$ in Lemma 1, the lemma confirms the intuitive idea that the integral of a unimodal, symmetric distribution over some interval $I = [a, b]$ is maximized when the peak of its probability distribution lies at the center of I . However, for the type of systems considered in this work, the shift of such distributions will always be restricted to take values within a given rectangular set $[r_1, r_2]$ so that, when $s_{max} \notin [r_1, r_2]$, the shift $s \in [r_1, r_2]$ maximizing the overlap of the density function over I is the one closest to the global maximizing shift s_{max} . Conversely, a shift $s \in [r_1, r_2]$ minimizing this overlap is the one furthest from s_{max} .

Theorem 1 combines Lemma 1 and Proposition 2 to provide a procedure for constructing an IMC abstraction for

(6) given a rectangular partition of its domain D .

Theorem 1: Consider system (6) under Assumptions 1–3 and let $P = \{Q_j\}_{j=1}^m$ be a rectangular partition of D with each $Q_j = \{x : a^j \leq x \leq b^j\}$ for some $a^j, b^j \in \mathbb{R}^n$ satisfying $a^j \leq b^j$. For all $Q_j, Q_\ell \in P$, let

$$s_{i,max}^\ell = \frac{a_i^\ell + b_i^\ell}{2} - c_i \text{ for } i = 1, \dots, n,$$

$$\check{r}^j = g(a^j, b^j),$$

$$\hat{r}^j = g(b^j, a^j),$$

and define

$$\hat{T}_{Q_j \rightarrow Q_\ell} = \prod_{i=1}^n \int_{a_i^\ell}^{b_i^\ell} f_{w_i}(x_i - s_{i,max}^{j \rightarrow \ell}) dx_i, \quad (11)$$

$$= \prod_{i=1}^n \left(F_{w_i}(b_i^\ell - s_{i,max}^{j \rightarrow \ell}) - F_{w_i}(a_i^\ell - s_{i,max}^{j \rightarrow \ell}) \right),$$

$$\check{T}_{Q_j \rightarrow Q_\ell} = \prod_{i=1}^n \int_{a_i^\ell}^{b_i^\ell} f_{w_i}(x_i - s_{i,min}^{j \rightarrow \ell}) dx_i \quad (12)$$

$$= \prod_{i=1}^n \left(F_{w_i}(b_i^\ell - s_{i,min}^{j \rightarrow \ell}) - F_{w_i}(a_i^\ell - s_{i,min}^{j \rightarrow \ell}) \right)$$

where F_{w_i} is the cumulative distribution function for w_i and

$$s_{i,max}^{j \rightarrow \ell} = \begin{cases} s_{i,max}^\ell, & \text{if } s_{i,max}^\ell \in [\check{r}_i^j, \hat{r}_i^j] \\ \hat{r}_i^j, & \text{if } s_{i,max}^\ell > \hat{r}_i^j \\ \check{r}_i^j, & \text{if } s_{i,max}^\ell < \check{r}_i^j, \end{cases} \quad (13)$$

$$s_{i,min}^{j \rightarrow \ell} = \begin{cases} \check{r}_i^j, & \text{if } s_{i,max}^\ell > \frac{\check{r}_i^j + \hat{r}_i^j}{2} \\ \hat{r}_i^j, & \text{otherwise.} \end{cases} \quad (14)$$

Then $\mathcal{I} = (P, \check{T}, \hat{T})$ is an IMC abstraction of (6).

Theorem 1 provides the mathematical foundation for solving Subproblem 1. Given a system of the form (6) satisfying Assumptions 1 to 3, and a rectangular partition P , this theorem shows that an IMC abstraction of (6) can be computed efficiently. Specifically, for any state in P , we establish an over-approximation of its one-step reachable set by evaluating the system's decomposition function at only two points. Finding the maximizing and minimizing shifts inside the reachable sets decouples along each coordinate and involves a number of operations that is linear in the dimension n of the state-space, according to (13) and (14). Finally, we see in (11) and (12) that n integral evaluations are needed per transition bound. This last step requires two evaluations of F_{w_i} per bound for each i and thus amounts to $2n$ function evaluations per bound. Because the computed over-approximations of reachable sets for mixed monotone dynamics were shown to be tight in certain cases [16], our abstraction technique is generally not as conservative as those employing Lipschitz constants [6], while being computationally efficient.

V. VERIFICATION OF IMCS

In this section, we develop the machinery to address Subproblem 2. Let \mathcal{I} be an IMC abstraction of (1) obtained from,

e.g., the abstraction approach in Section IV for systems with specific form (6). For a formula ϕ of the type (2), our goal is to sort the initial states of \mathcal{I} into the sets Q_ϕ^{yes} , Q_ϕ^{no} and $Q_\phi^?$. To this end, for any initial state Q_j of \mathcal{I} , we require a lower bound and an upper bound on the probability of satisfying the ω -regular property Ψ for the probabilistic specification $\phi = \mathcal{P}_{\bowtie p_{sat}}[\Psi]$. We thus seek to compute the greatest lower bound $\check{\mathcal{P}}_{\mathcal{I}}(Q_j \models \Psi)$ and least upper bound $\widehat{\mathcal{P}}_{\mathcal{I}}(Q_j \models \Psi)$ such that, for any adversary $\nu \in \nu_{\mathcal{I}}$,

$$\check{\mathcal{P}}_{\mathcal{I}}(Q_j \models \Psi) \leq \mathcal{P}_{\mathcal{I}[\nu]}(Q_j \models \Psi) \leq \widehat{\mathcal{P}}_{\mathcal{I}}(Q_j \models \Psi).$$

Our approach draws from the verification of regular MCs against ω -regular properties using automata-based methods [1, Section 10.3]. First, we generate a DRA \mathcal{A} that recognizes the language induced by property Ψ . Such a DRA always exists. Several algorithms can generate a DRA for a large subset of ω -regular expressions [29] [30]. Then, we construct the product $\mathcal{I} \otimes \mathcal{A}$ as defined below.

Definition 7 (Product Interval-valued Markov Chain): Let $\mathcal{I} = (Q, \widetilde{T}, \widehat{T}, q_0, \Sigma, L)$ be an IMC and $\mathcal{A} = (S, \Sigma, \delta, s_0, Acc)$ be a DRA. The *product* $\mathcal{I} \otimes \mathcal{A} = (Q \times S, \widetilde{T}', \widehat{T}', q_0^\otimes, Acc', L')$ is an IMC where:

- $Q \times S$ is a set of states,
- $\widetilde{T}'_{\langle Q_j, s \rangle \rightarrow \langle Q_\ell, s' \rangle} = \begin{cases} \widetilde{T}'_{Q_j \rightarrow Q_\ell}, & \text{if } s' = \delta(s, L(Q_\ell)) \\ 0, & \text{otherwise} \end{cases}$
- $\widehat{T}'_{\langle Q_j, s \rangle \rightarrow \langle Q_\ell, s' \rangle} = \begin{cases} \widehat{T}'_{Q_j \rightarrow Q_\ell}, & \text{if } s' = \delta(s, L(Q_\ell)) \\ 0, & \text{otherwise} \end{cases}$
- $q_0^\otimes = \{(Q_j, s_0) : Q_j \in q_0\}$ is a set of initial states,
- $Acc' = \{E_1, E_2, \dots, E_k, F_1, F_2, \dots, F_k\}$ is a set of atomic propositions, where E_i and F_i are the sets in the Rabin pairs of Acc ,
- $L' : Q \times S \rightarrow 2^{Acc'}$ such that $H \in L'(\langle Q_j, s \rangle)$ if and only if $s \in H$, for all $H \in Acc'$ and for all j .

A MC induced by $\mathcal{I} \otimes \mathcal{A}$ is called a *product Markov Chain*, and we use the notation \mathcal{M}_{\otimes}^A to denote such an induced MC.

The probability of satisfying Ψ from initial state Q_j in a discrete-time MC equals that of reaching an accepting *Bottom Strongly Connected Component* (BSCC) from initial state $\langle Q_j, s_0 \rangle$ in the product MC with \mathcal{A} [1].

Definition 8 (Bottom Strongly Connected Component):

Given a MC \mathcal{M} with states Q , a set $B \subseteq Q$ is a *Bottom Strongly Connected Component* (BSCC) of \mathcal{M} if

- B is strongly connected: for each pair of states (q, t) in B , there exists a path $q_0 q_1 \dots q_n$ such that $T(q_i, q_{i+1}) > 0$, $i = 0, 1, \dots, n-1$, and $q_i \in B$ for $0 \leq i \leq n$ with $q_0 = q$, $q_n = t$,
- no proper superset of B is strongly connected,
- $\forall s \in B, \Sigma_{t \in B} T(s, t) = 1$.

In words, every state in a BSCC B is reachable from any state in B , and every state in B only transitions to another state in B . Moreover, B is accepting when at least one of

its states maps to the accepting set of a Rabin pair, while no state in B maps to the non-accepting set of that same pair.

Definition 9 (Accepting Bottom Strongly Connected Component): A Bottom Strongly Connected Component B of a product MC \mathcal{M}_{\otimes}^A is said to be *accepting* if:

$$\begin{aligned} \exists i : & \left(\exists \langle Q_j, s_\ell \rangle \in B . F_i \in L'(\langle Q_j, s_\ell \rangle) \right) \\ & \wedge \left(\forall \langle Q_j, s_\ell \rangle \in B . E_i \notin L'(\langle Q_j, s_\ell \rangle) \right). \end{aligned}$$

Definition 10 (Non-Accepting Bottom Strongly Connected Component): A Bottom Strongly Connected Component B of a product MC \mathcal{M}_{\otimes}^A is said to be *non-accepting* if it is not accepting.

We denote by U^A and U^N the sets of states that respectively belong to an accepting and a non-accepting BSCC in a product MC.

Note that each product MC \mathcal{M}_{\otimes}^A induced by $\mathcal{I} \otimes \mathcal{A}$ simulates the behavior of \mathcal{I} under some adversary $\nu \in \nu_{\mathcal{I}}$. Indeed, for any two states Q_j and Q_ℓ in \mathcal{I} and some states s, s', s'' and s''' in \mathcal{A} , we allow $T_{(Q_j, s) \rightarrow \langle Q_\ell, s' \rangle}$ and $T_{(Q_j, s'') \rightarrow \langle Q_\ell, s''' \rangle}$ to assume different values in \mathcal{M}_{\otimes}^A , which means that the transition probability between Q_j and Q_ℓ may change depending on the history of the path in \mathcal{I} as encoded in the state of \mathcal{A} .

Also, the adversary is history-independent or *memoryless* in the product automaton, that is, the adversary's chosen transition probability only depends on the current states of the IMC and the DRA \mathcal{A} . For reachability problems in IMCs, it was shown in [31] that memoryless adversaries yield the same bounds as the memory-dependent ones. The following facts establish that, therefore, such memoryless (in the product) adversaries are sufficient for IMC verification.

Fact 2: [1, p. 792, Theorem 10.56] [31] We denote the set of adversaries of \mathcal{I} that are memoryless in the product IMC $\mathcal{I} \otimes \mathcal{A}$ by $(\nu_{\mathcal{I}})_{\otimes}^A \subseteq \nu_{\mathcal{I}}$. It holds that

$$\begin{aligned} \inf_{\nu \in \nu_{\mathcal{I}}} \mathcal{P}_{\mathcal{I}[\nu]}(Q_i \models \Psi) &= \inf_{\nu \in (\nu_{\mathcal{I}})_{\otimes}^A} \mathcal{P}_{\mathcal{I}[\nu]}(Q_i \models \Psi) \\ \sup_{\nu \in \nu_{\mathcal{I}}} \mathcal{P}_{\mathcal{I}[\nu]}(Q_i \models \Psi) &= \sup_{\nu \in (\nu_{\mathcal{I}})_{\otimes}^A} \mathcal{P}_{\mathcal{I}[\nu]}(Q_i \models \Psi). \end{aligned}$$

Fact 3: For any adversary $\nu \in (\nu_{\mathcal{I}})_{\otimes}^A$ in \mathcal{I} , it holds that $\mathcal{P}_{\mathcal{I}[\nu]}(Q_i \models \Psi) = \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_{\nu}}(\langle Q_i, s_0 \rangle \models \diamond U^A)$, where $(\mathcal{M}_{\otimes}^A)_{\nu}$ denotes the product MC induced by $\mathcal{I} \otimes \mathcal{A}$ corresponding to adversary ν .

Consequently, computing $\check{\mathcal{P}}_{\mathcal{I}}(Q_i \models \Psi)$ and $\widehat{\mathcal{P}}_{\mathcal{I}}(Q_i \models \Psi)$ amounts to finding the product MCs induced by $\mathcal{I} \otimes \mathcal{A}$ that respectively minimize and maximize the probability of reaching an accepting BSCC from $\langle Q_i, s_0 \rangle$. Such reachability problems in IMCs were solved when the destination states are fixed for all induced MCs [12] [32].

However, in general, the sets U^A and U^N are not fixed in product IMCs and vary with the assumed values for each transition. Specifically, U^A and U^N are determined by transitions that are either “on” or “off”, i.e. those whose lower bound is zero and upper bound is non-zero, as seen in the example in Fig. 3: in the product MC $(\mathcal{M}_{\otimes}^A)_1$ induced by $\mathcal{I} \otimes \mathcal{A}$, U^A is




Fig. 3. Two product MCs $(\mathcal{M}_{\otimes}^A)_1$ and $(\mathcal{M}_{\otimes}^A)_2$ induced by a product IMC $\mathcal{I} \otimes \mathcal{A}$. Accepting BSCCs are shown in green; non-accepting BSCCs are red. In $(\mathcal{M}_{\otimes}^A)_1$, $U^A = \{Q_0, Q_1\}$ and $U^N = \{Q_2\}$; in $(\mathcal{M}_{\otimes}^A)_2$, $U^A = \emptyset$ and $U^N = \{Q_0, Q_1, Q_2\}$.

$\{Q_0, Q_1\}$ while $\{Q_2\}$ is non-accepting. In $(\mathcal{M}_{\otimes}^A)_2$, another product MC induced by $\mathcal{I} \otimes \mathcal{A}$, all states are in U^N .

The contribution of this paper to the verification of IMCs for ω -regular properties is twofold. First, we show that a product IMC always induces a largest *Losing Component* and *Winning Component*. These components contain states that reach a BSCC with probability 1. Upper and lower bounds on Ψ are computed by solving a reachability problem for these sets. We further introduce the notion of *Permanent Losing Components* and *Permanent Winning Components* which play a crucial role in the refinement algorithm derived in Section VI. Second, we describe a graph search algorithm to find these components.

A. Computation of Satisfiability Bounds in IMCs

Previous works highlighted the crucial role of BSCCs in product MCs [1, Theorem 10.56]. As the probability of reaching an accepting BSCC in a product MC determines the probability of satisfying some property in the original abstraction, we now further introduce the notions of winning and losing components. These components include states that may not belong to a BSCC but from which any path is bound to reach a BSCC.

Definition 11 (Winning/Losing Component): [25] A *winning (losing) component* WC (LC) of a product MC \mathcal{M}_{\otimes}^A is a set of states satisfying $\mathcal{P}_{\mathcal{M}_{\otimes}^A}(WC \models \diamond U^A) = 1$ ($\mathcal{P}_{\mathcal{M}_{\otimes}^A}(LC \models \diamond U^N) = 1$), where U^A (U^N) is the set of states belonging to an accepting (non-accepting) BSCC in \mathcal{M}_{\otimes}^A .

It naturally follows that the probability of eventually reaching a BSCC from some initial state is equal to that of reaching a winning or losing component.

Corollary 1: In any product MC \mathcal{M}_{\otimes}^A ,

$$\begin{aligned} \mathcal{P}_{\mathcal{M}_{\otimes}^A}(\langle Q_i, s_0 \rangle \models \diamond U^A) &= \mathcal{P}_{\mathcal{M}_{\otimes}^A}(\langle Q_i, s_0 \rangle \models \diamond WC) \\ \mathcal{P}_{\mathcal{M}_{\otimes}^A}(\langle Q_i, s_0 \rangle \models \diamond U^N) &= \mathcal{P}_{\mathcal{M}_{\otimes}^A}(\langle Q_i, s_0 \rangle \models \diamond LC). \end{aligned}$$

For any initial state in a product IMC, our goal is thus to find induced product MCs that minimize and maximize the probability of reaching a winning component.

We refer to the technical Appendix for all lemmas and proofs leading to the proposed solution. The key observation

is that any product IMC induces a *largest winning component* and a *largest losing component*. The largest winning component is the set of states of the product IMC belonging to a winning component for at least one induced product MC, while the largest losing component is the analogous set for losing components. Definitions of *permanent* and *potential* components follow directly from that of largest components.

Definition 12 (Largest Winning/Losing Components): A state $\langle Q_i, s_j \rangle \in Q \times S$ of a product IMC $\mathcal{I} \otimes \mathcal{A}$ is a member of the *Largest Winning (Losing) Component* $(WC)_L$ ($(LC)_L$) if there exists a product MC induced by $\mathcal{I} \otimes \mathcal{A}$ such that $\langle Q_i, s_j \rangle$ belongs to a winning (losing) component.

Definition 13 (Permanent Winning/Losing Components): A state $\langle Q_i, s_j \rangle \in Q \times S$ of a product IMC $\mathcal{I} \otimes \mathcal{A}$ is a member of the *Permanent Winning (Losing) Component* $(WC)_P$ ($(LC)_P$) of $\mathcal{I} \otimes \mathcal{A}$ if $\langle Q_i, s_j \rangle$ belongs to a winning (losing) component for all product MCs induced by $\mathcal{I} \otimes \mathcal{A}$.

Definition 14 (Potential Winning/Losing Components): A state $\langle Q_i, s_j \rangle \in Q \times S$ of a product IMC $\mathcal{I} \otimes \mathcal{A}$ is a member of the *Potential Winning (Losing) Component* $(WC)_?$ ($(LC)_?$) of $\mathcal{I} \otimes \mathcal{A}$ if $\langle Q_i, s_j \rangle \in (WC)_L \setminus (WC)_P$ ($\langle Q_i, s_j \rangle \in (LC)_L \setminus (LC)_P$).

Note that the sets $(WC)_?$ and $(LC)_?$ may intersect, and by extension $(WC)_L$ and $(LC)_L$, while $(WC)_P$ and $(LC)_P$ are disjoint. An important result established in this paper is that any product IMC induces a set of product MCs where all members of the largest winning component belong to a winning component simultaneously. A product IMC induces an analogous set of product MCs for the largest losing component. We provide proofs in Lemmas 5-7 of the Appendix.

We now state the main result of this section, which establishes that bounds on the probability of satisfying an ω -regular property in an IMC can be computed by solving a reachability maximization problem on a fixed set of states in a product IMC. These sets are the largest components of the product IMC. Furthermore, solving these problems induce sets of best and worst-case product MCs where the probabilities of reaching a winning component are respectively maximized and minimized for all initial states of the product IMC.

Theorem 2: Let \mathcal{I} be an IMC and \mathcal{A} be a DRA corresponding to ω -regular property Ψ . Let $(WC)_L$ and $(LC)_L$ be the largest winning and losing components and $(WC)_P$ and $(LC)_P$ be the permanent winning and losing components of the product IMC $\mathcal{I} \otimes \mathcal{A}$. Then for any initial state Q_i of \mathcal{I} ,

$$\tilde{\mathcal{P}}_{\mathcal{I}}(Q_i \models \Psi) = 1 - \hat{\mathcal{P}}_{\mathcal{I} \otimes \mathcal{A}}(\langle Q_i, s_0 \rangle \models \diamond (LC)_L) \quad (15)$$

$$\hat{\mathcal{P}}_{\mathcal{I}}(Q_i \models \Psi) = \hat{\mathcal{P}}_{\mathcal{I} \otimes \mathcal{A}}(\langle Q_i, s_0 \rangle \models \diamond (WC)_L). \quad (16)$$

Moreover, there exists a set of induced product MCs $(\mathcal{M}_{\otimes}^A)_{\text{worst}}$, where, $\forall \mathcal{M}_i \in (\mathcal{M}_{\otimes}^A)_{\text{worst}}$, the sets of all losing and winning components of \mathcal{M}_i are $(LC)_L$ and $(WC)_P$ respectively, and, $\forall \langle Q_i, s_0 \rangle \in (Q \times S)$, $\mathcal{P}_{\mathcal{M}_i}(\langle Q_i, s_0 \rangle \models \diamond (LC)_L) = \hat{\mathcal{P}}_{\mathcal{I} \otimes \mathcal{A}}(\langle Q_i, s_0 \rangle \models \diamond (LC)_L)$. Likewise, there

exists a set of induced product MCs $(\mathcal{M}_{\otimes}^A)_{\text{best}}$, where, $\forall \mathcal{M}_i \in (\mathcal{M}_{\otimes}^A)_{\text{best}}$, the sets of all losing and winning components of \mathcal{M}_i are $(LC)_P$ and $(WC)_L$ respectively, and, $\forall \langle Q_i, s_0 \rangle \in (Q \times S)$, $\mathcal{P}_{\mathcal{M}_i}(\langle Q_i, s_0 \rangle \models \diamond(WC)_L) = \widehat{\mathcal{P}}_{\mathcal{I} \otimes \mathcal{A}}(\langle Q_i, s_0 \rangle \models \diamond(WC)_L)$.

Proof: $\widehat{\mathcal{P}}_{\mathcal{I}}(Q_i \models \Psi)$ is equivalent to a lower bound on the probability of reaching an accepting BSCC from $\langle Q_i, s_0 \rangle$ in $\mathcal{I} \otimes \mathcal{A}$. Equation (15) follows from Lemma 3 and the following reasoning: assume (15) is not true. This implies that there exists an induced product MC where the probability of reaching a non-accepting BSCC from $\langle Q_i, s_0 \rangle$ is greater than the highest probability of reaching $(LC)_L$, which is a contradiction to Lemma 5 and 9. Next, we denote by \mathcal{D} the set of induced product MCs with set of winning components $(WC)_P$ and set of losing components $(LC)_L$ constructed in Lemma 10. Lemma 8 and Lemma 9 guarantee that the probability of reaching an accepting BSCC from all $\langle Q_i, s_0 \rangle$ is minimized in induced product MCs with the smallest set of winning components and the largest set of losing components respectively. Therefore, $(\mathcal{M}_{\otimes}^A)_{\text{worst}} \subseteq \mathcal{D}$. Equation (16) and the existence of $(\mathcal{M}_{\otimes}^A)_{\text{best}}$ are proved identically. ■

The equalities highlighted by this theorem are central to the elaboration of our verification procedure. We first solve a qualitative problem, which is to find the largest components of the product IMC. This can be achieved via graph search and will be the focus of the next section. Then, we compute upper and lower bound probabilities of reaching these components from all states in the product IMC using existing algorithms found in the literature [12] [32]. By doing so, we construct a best-case product MC $(\mathcal{M}_{\otimes}^A)_u \in (\mathcal{M}_{\otimes}^A)_{\text{best}}$ and a worst-case product MC $(\mathcal{M}_{\otimes}^A)_l \in (\mathcal{M}_{\otimes}^A)_{\text{worst}}$ which respectively maximizes and minimizes the probability of reaching an accepting BSCC from all initial states. Note that the transition values between states inside the components do not affect the reachability probabilities and do not need to be considered.

B. Winning and Losing Components Search Algorithm

We present a graph-based algorithm for finding $(WC)_P$, $(WC)_L$, $(LC)_P$ and $(LC)_L$ in a product IMC, divided into Algorithm 1 and Algorithm 2. We define the sets of potential and permanent BSCCs $(U^A)_?$, $(U^A)_P$, $(U^N)_?$ and $(U^N)_P$. Algorithm 1 takes a product IMC as input and returns its potential and permanent BSCCs. Algorithm 2 takes as inputs a product IMC and its permanent and potential BSCCs and outputs $(WC)_P$, $(WC)_L$, $(LC)_P$ and $(LC)_L$. The potential components $(WC)_?$ and $(LC)_?$ are the set differences between the largest and permanent components.

We employ the following notations: a digraph G is said to be generated by an induced MC \mathcal{M}_{\otimes}^A with transition matrix T and states $Q \times S$ if G has a representative vertex for all states in \mathcal{M}_{\otimes}^A , and an edge exists between two such vertices if $T(Q_i, Q_j) > 0$, $Q_i, Q_j \in Q \times S$. $\text{Reach}(S, G)$ denotes the set of vertices in graph G from which there exists a path to the set of vertices S ; $\text{At}_?(S, G)$ denotes the set of vertices in G from which there exists a path to S for all graphs G' generated by an induced product MC of $\mathcal{I} \otimes \mathcal{A}$, where G and

Algorithm 1 Find Potential and Permanent BSCCs

```

1: Input: Product IMC  $\mathcal{I} \otimes \mathcal{A}$ 
2: Output: Potential and permanent BSCCs  $(U^A)_?$ ,  $(U^A)_P$ ,  $(U^N)_?$ ,  $(U^N)_P$ 
3: Initialize:  $(U^A)_?$ ,  $(U^A)_P$ ,  $(U^N)_?$ ,  $(U^N)_P := \emptyset$ 
4: Construct  $G = (V, E)$  with a vertex for each state in  $\mathcal{I} \otimes \mathcal{A}$  and an edge between states  $Q_i$  and  $Q_j$  if  $\widehat{T}(Q_i, Q_j) > 0$ 
5: Find all SCCs of  $G$  and list them in  $S$ 
6: for  $S_k \in S$  do
7:    $C_0 := \emptyset$ ,  $i := 0$ 
8:   repeat
9:      $R_i := S_k \setminus \cup_{\ell=0}^i C_\ell$ ;  $Tr_i := V \setminus R_i$ ;  $C_{i+1} := \text{At}_?(Tr_i, R_i)$ ;  $i := i + 1$ 
10:  until  $C_i = \emptyset$ 
11:  if  $i \neq 1$  then
12:    Find all SCCs of  $R_i$  and add them to  $S$ 
13:  else
14:    if  $S_k$  is accepting then
15:      In  $C$ , list all states in  $S_k$  mapping to some accepting set  $F_i$  if no other state in  $S_k$  maps to  $E_i$ . Find all SCCs of  $S_k \setminus \text{At}_?(C, S_k)$  and add them to  $S$ .
16:    else
17:      For all sets  $F_i$  to which at least one state in  $S_k$  is mapped, set  $S'_k = S_k$ , list all states mapping to  $E_i$  in  $C$ , find all SCCs of  $S'_k \setminus \text{At}_?(C, S'_k)$  and add them to  $S$ .
18:    end if
19:    if  $\text{At}_P(V \setminus S_k, S_k) \neq \emptyset$  then
20:       $(U^A)_? := (U^A)_? \cup \{S_k\}$  or  $(U^N)_? := (U^N)_? \cup \{S_k\}$  depending on the acceptance status of  $S_k$ .
21:    else
22:       $(U^A)_P := (U^A)_P \cup \{S_k\}$  or  $(U^N)_P := (U^N)_P \cup \{S_k\}$  depending on the acceptance status  $S_k$  and if no other state in  $S_k$  belongs to a potential BSCC of the opposite acceptance status. Else,  $(U^A)_? := (U^A)_? \cup \{S_k\}$  or  $(U^N)_? := (U^N)_? \cup \{S_k\}$ .
23:    end if
24:  end if
25: end for
26: return  $(U^A)_?$ ,  $(U^A)_P$ ,  $(U^N)_?$ ,  $(U^N)_P$ 

```

G' share the same set of vertices; and $\text{At}_P(S, G)$ denotes the set of vertices in G from which there exists a path to S for at least one graph G' generated by an induced product MC of $\mathcal{I} \otimes \mathcal{A}$. A detailed description of the algorithms is found below.

Algorithm 1:

Line 4: We first assume all transitions with a non-zero upper bound to be “on” and generate a graph $G = (V, E)$ with a vertex for all states and an edge for all transitions in $\mathcal{I} \otimes \mathcal{A}$.

Line 5: Next, we find all strongly connected components (SCC) of G and list them in S .

Line 6 to 10: For all SCC $S_k \in S$, we want to determine if there exists an induced MC where S_k is a BSCC. To this end, for all the states S_k^j in S_k , we check whether all outgoing transitions to states not in S_k can be turned “off” for some

induced product MC, that is if the transition lower bounds from S_k^j to states in $Tr_i = V \setminus R_i$ are 0 and the sum of the transition upper bounds from S_k^j to states in S_k is greater than 1, which is captured by the use of the function $At_?$. Otherwise, S_k^j is said to be *leaky* in all induced product MCs and S_k^j is added to the set C_i , which contains all leaky states of S_k found at iteration i . Note that $R_0 = S_k$ and that all leaky states previously found are removed from S_k at each iteration via variable R_i . The loop terminates when all states have been checked and no more leaky states are found, that is $C_i = \emptyset$.

Line 11 to 13: If S_k contained leaky states that were previously removed, we compute all SCCs formed by the remaining states in R_i and add them to the list of SCCs of G . If S_k did not contain any leaky state, it is a member of a largest set of BSCCs and the mapping of the states in S_k with respect to the Rabin Pairs decides whether S_k is accepting and $S_k \in (U^A)_L$ or non-accepting and $S_i \in (U^N)_L$.

Line 14 to 15: If $S_k \in (U^A)_L$, it could still contain potential non-accepting BSCCs, since $(U^A)_L$ and $(U^N)_L$ may comprise intersecting sets. Treat all states causing S_k to be accepting as leaky (states mapping to some F_i in the Rabin pairs when no states in S_k maps to E_i), remove from S_k all states that have a permanent path to the leaky states, compute all SCCs formed by the remaining states and add them to S .

Line 16 to 17: If $S_k \in (U^N)_L$, potential accepting BSCCs may lie inside S_k . For all sets F_i in the Rabin pairs to which at least one state in S_k is mapped, create a “copy” S'_k of S_k where all states causing S_k to be non-accepting are considered leaky (the states mapping to E_i), remove from S'_k all states that have a permanent path to the leaky states, compute all SCCs formed by the remaining states and add them to S .

Line 19 to 22: We check whether some state in B leaks outside of B for at least one induced MC. If so, the BSCC is not permanent. Otherwise, B is permanent if and only if no BSCC of the opposite acceptance status is found inside of B .

Algorithm 2:

Inspired by the Classical Algorithm for Buchi MDPs [33], we perform a graph search to find the permanent and largest winning and losing components of $\mathcal{I} \otimes \mathcal{A}$. The permanent components arise from permanent BSCCs, while the largest components stem from both potential and permanent BSCCs.

Line 3: We generate a graph $G = (V, E)$ where transitions with a non-zero upper bound are assumed to be “on”.

Line 4 to 11: To find the largest winning component $(WC)_L$ of the product IMC $\mathcal{I} \otimes \mathcal{A}$, we find the set R_i of all states from which there is a path to $(U^A)_? \cup (U^A)_P$ in G . Other states in G are “trap states” denoted by Tr_i . Then, we iteratively remove the set of states C_i from R_i that “leak” to Tr_i for all induced MCs of $\mathcal{I} \otimes \mathcal{A}$, and compute the new set R_{i+1} of states that have a path to $(U^A)_? \cup (U^A)_P$ once the leaky states are discarded. The iteration stops when no more leaky states are found, that is $C_i = \emptyset$. The remaining states belong to the largest winning component $(WC)_L$. The same procedure is applied with respect to $(U^N)_? \cup (U^N)_P$ to find the largest losing component $(LC)_L$.

Line 12 to 20: To find the permanent winning component $(WC)_P$ of the product IMC $\mathcal{I} \otimes \mathcal{A}$, we first discard all states

Algorithm 2 Find Largest and Permanent Components

- 1: **Input:** Product IMC $\mathcal{I} \otimes \mathcal{A}$ and its potential and permanent BSCCs $(U^A)_?, (U^A)_P, (U^N)_?, (U^N)_P$
 - 2: **Output:** Largest and permanent components $(WC)_L, (LC)_L, (WC)_P, (LC)_P$
 - 3: Construct $G = (V, E)$ with a vertex for each state in $\mathcal{I} \otimes \mathcal{A}$ and an edge between states Q_i and Q_j if $\hat{T}(Q_i, Q_j) > 0$
 - 4: **for** $B \in \{(U^A)_? \cup (U^A)_P, (U^N)_? \cup (U^N)_P\}$ **do**
 - 5: $C_0 := \emptyset, V_0 := V, i := 0$
 - 6: **repeat**
 - 7: $R_i := Reach(B \cap V_i, V_i); Tr_i := V_i \setminus R_i; C_{i+1} := At_?(Tr_i, V_i); V_{i+1} := V_i \setminus C_{i+1}; i := i + 1$
 - 8: **until** $C_i = \emptyset$
 - 9: $W := V \setminus \cup_{k=1}^i C_k$
 - 10: $(WC)_L := W, V_A := V_i$ if $B = (U^A)_? \cup (U^A)_P$ or $(LC)_L := W, V_N := V_i$ if $B = (U^N)_? \cup (U^N)_P$
 - 11: **end for**
 - 12: **for** $B \in \{(U^A)_P, (U^N)_P\}$ **do**
 - 13: In D , list the states of V belonging to the largest component of the opposite acceptance status from B
 - 14: $C_0 := \emptyset, V_0 := V_A \setminus D$ if $B = (U^A)_P$ or $V_0 := V_N \setminus D$ if $B = (U^N)_P, i := 0$
 - 15: **repeat**
 - 16: $R_i := Reach(B \cap V_i, V_i); Tr_i := V_i \setminus R_i; C_{i+1} := At_P(Tr_i, V_i); V_{i+1} := V_i \setminus C_{i+1}; i := i + 1$
 - 17: **until** $C_i = \emptyset$
 - 18: $W_P := V_0 \setminus \cup_{k=1}^i C_k$
 - 19: $(WC)_P := W_P$ if $B = (U^A)_P$ or $(LC)_P := W_P$ if $B = (U^N)_P$
 - 20: **end for**
 - 21: **return** $(WC)_L, (LC)_L, (WC)_P, (LC)_P$
-

which belong to the largest losing component $(LC)_L$ from the set of edges V_A to be analyzed. Then, we repeat the same procedure as in Algorithm 1 with respect to $(U^A)_P$ until no more leaky states are found, that is $C_i = \emptyset$, except that leaky states are now those which have a path to the trap states Tr_i in at least one induced MC of the product IMC $\mathcal{I} \otimes \mathcal{A}$. The remaining states upon completion of this iterative removal belong to $(WC)_P$. The same procedure is applied with respect to $(U^N)_P$ to find the permanent losing component $(LC)_P$.

To summarize, it is known that verification against temporal logic specifications in discrete-time MCs can be accomplished by solving a reachability problem on a product MC constructed from a Rabin automaton corresponding to the specification to be verified. The heart of this approach relies on analyzing winning and losing components of the product MC. These ideas do not directly extend to IMCs because BSCCs are not uniquely determined in this case; this is because some transitions can have a lower transition bound equal to 0 but an upper transition bound that is non-zero. Instead, we introduced the concepts of largest winning and losing components. In Theorem 2, we show that upper and lower bounds on the probability of satisfaction are obtained from these components. Algorithms 1 and 2 provide means for computing these components. Note that the proposed algorithm allows to perform verification

Algorithm 3 State-Space Refinement Scoring Procedure

```

1: Input: Worst-case product MC  $(\mathcal{M}_{\otimes}^A)_l$  and best-case
   product MC  $(\mathcal{M}_{\otimes}^A)_u$  induced by the product  $\mathcal{I} \otimes A$ 
2: Output: Scores  $\sigma = [\sigma_0, \dots, \sigma_N]$  for all states in  $\mathcal{I}$ 
3: Initialize:  $\sigma_i = 0$ , with  $\sigma_i$  the score of the  $i$ -th state of  $\mathcal{I}$ ,
    $p_{stop} \in (0, 1)$  user-defined probability threshold
4: for  $Q_\ell \in Q_\phi^?$  do
5:    $\pi := q_0 := \langle Q_\ell, s_0 \rangle$  in  $(\mathcal{M}_{\otimes}^A)_u$ 
6:   repeat
7:     if  $\mathcal{P}(\pi) < p_{stop}$  or  $Exp(\pi) = R(\pi)$  then
8:        $\pi := \pi^-$ 
9:     else
10:      if  $Exp(\pi) \neq \emptyset$  then
11:         $q_i \rightarrow Exp(\pi)$ , where  $q_i$  is any state in  $R(\pi) \setminus$ 
           $Exp(\pi)$ ;  $\pi := \pi^+(q_i)$ 
12:      else
13:         $Exp(\pi) := \cup_i \pi_i$ 
14:        if  $Last(\pi) \in (WC)_? \cup (LC)_?$  then
15:           $\sigma_j := \sigma_j + \mathcal{P}(\pi)(p_{max} - p_{min})$  for all states
             $\langle Q_j, s_i \rangle$  in the potential BSCC of  $Last(\pi)$ 
            with an outgoing transition which can be
            either zero or non-zero,  $p_{max}$  and  $p_{min}$  are
            the probabilities of reaching an accepting
            BSCC from  $Last(\pi)$  in  $(\mathcal{M}_{\otimes}^A)_u$  and  $(\mathcal{M}_{\otimes}^A)_l$ 
            respectively;  $\pi := \pi^-$ 
16:          else if  $Last(\pi) \in (WC)_P \cup (LC)_P$  then
17:             $\pi := \pi^-$ 
18:          else
19:             $\sigma_j := \sigma_j + \mathcal{P}(\pi)(p_{max} - p_{min})$ , where  $j$ 
              corresponds to  $\langle Q_j, s_i \rangle := Last(\pi)$ ,  $p_{max}$ 
              and  $p_{min}$  are as in line 15;
20:             $\pi := \pi^+(q_i)$  where  $q_i$  is any state in  $R(\pi)$ 
21:          end if
22:        end if
23:      end if
24:    until  $\pi = \emptyset$ 
25:  end for
26: return  $\sigma$ 

```

of IMCs without constructing an exponentially large Markov decision process, as done in [32].

VI. STATE-SPACE REFINEMENT

Given a partition P of the domain D and a specification ϕ as in (2), the verification procedure derived in Section V assigns each discrete state of P to one of the sets Q_ϕ^{yes} , Q_ϕ^{no} or $Q_\phi^?$. One aims to find a partition P that yields a low volume of undecided states in $Q_\phi^?$. To this end, we suggest a specification-guided iterative method. Specifically, we first generate a rough partition P of D and successively refine P into finer partitions by targeting the best candidate states for reducing the uncertainty in the abstraction with respect to ϕ . These states are chosen after comparing the behavior of the system in the best and worst-case scenarios computed during verification. The procedure stops when a user-defined criterion is reached. Here, we terminate when the fractional volume of

uncertain states is less than a threshold $V_{stop} \in [0, 1]$.

We seek to analyze the behavior of accepting paths in the best and worst-case product MCs $(\mathcal{M}_{\otimes}^A)_u$ and $(\mathcal{M}_{\otimes}^A)_l$ obtained at the time of verification and illustrated in Fig. 4. In particular, for every undecided state Q_j in $Q_\phi^?$, we look at all paths starting from $\langle Q_j, s_0 \rangle$ in $(\mathcal{M}_{\otimes}^A)_u$ and assign a score to the states encountered along them depending on how these states behave in $(\mathcal{M}_{\otimes}^A)_l$. We inspect a path until it reaches a state that belongs to either $(WC)_L$ or $(LC)_L$, or when its probability of occurrence in $(\mathcal{M}_{\otimes}^A)_u$ falls below a threshold p_{stop} . States with high scores are targeted for refinement.

We introduce some notation: for a finite path $\pi = q_0 q_1 \dots q_k$ in $(\mathcal{M}_{\otimes}^A)_u$, $Last(\pi)$ denotes the last state q_k of π ; π_i denotes the i -th state of π ; $\mathcal{P}(\pi) = T(q_0, q_1) \cdot T(q_1, q_2) \cdot \dots \cdot T(q_{k-1}, q_k)$, $\mathcal{P}(q_0) = 1$, is the probability of path π in $(\mathcal{M}_{\otimes}^A)_u$; $R(\pi)$ is the set of states that are one-step reachable from $Last(\pi)$ in $(\mathcal{M}_{\otimes}^A)_u$; $Exp(\pi)$ denotes all continuations of π from $R(\pi)$ that have been explored and is initialized to the empty set for all π ; π^- is the path obtained by removing the last state of π and $\pi^+(q_i)$ is the path with q_i appended to π . V_ϕ is the fractional volume of uncertain states and is equal to the sum of the volume of all states in $Q_\phi^?$ divided by the volume of the domain D . Our procedure is as follows:

- 1) Compute a refinement score for all states in \mathcal{I} according to Algorithm 3, which is described below:

Line 6 to 24: This loop terminates when $\pi = \emptyset$, that is, when all paths starting from $\langle Q_j, s_0 \rangle$ have been explored.

Line 7 to 8: If $\mathcal{P}(\pi) < p_{stop}$ or $Exp(\pi) = R(\pi)$, the probability of the path is below the pre-defined exploration threshold or all continuations of π have been explored. Thus, we return to the previous state in the path.

Line 10 to 13: If $Exp(\pi) \neq \emptyset$, add q_i to $Exp(\pi)$, where q_i is some unexplored state in $R(\pi)$ and extend the path to q_i . Else, π is a path fragment which has not been explored yet. Add all states in π to $Exp(\pi)$ to avoid loops.

Line 14 to 15: If $Last(\pi) \in (WC)_?$ or $Last(\pi) \in (LC)_?$, the path reached a state in a potential component. We want to target the states which can either confirm or refute that $Last(\pi)$ belongs to such a component. These states are the ones inside the potential BSCCs that $Last(\pi)$ belongs to (or makes a transition to with probability 1) that have outgoing transitions which can be either “on” or “off”, as depicted in Fig. 5. A potential “certainty gain” is added to the score of all such states and the path is returned to its previous state. If $Last(\pi)$ belongs to both $(WC)_?$ and $(LC)_?$, then the scoring scheme is applied to all intersecting potential BSCCs related to $Last(\pi)$. This heuristical gain quantifies a potential reduction in the width of the satisfaction interval of Q_ℓ in the scenario that the refinement of the considered states provides perfect information, i.e., the probability of reaching an accepting BSCC from $Last(\pi)$ becomes a fixed number.

Line 16 to 17: If $Last(\pi) \in (WC)_P$ or $Last(\pi) \in (LC)_P$, the path reached a region of the state-space that does not require refinement as it belongs to a permanent component. The path returns to its previous state.

Line 18 to 20: Else, $Last(\pi)$ does not belong to a




Fig. 4. Our IMC verification algorithm generates a best and worst-case product MC $(\mathcal{M}_{\otimes}^A)_u$ and $(\mathcal{M}_{\otimes}^A)_l$. Winning and losing components are respectively in red and green; permanent and potential components are circled in bold and dotted lines respectively. Comparing the behavior of the paths in the two scenarios is the basis of our refinement algorithm.




Fig. 5. For all undecided states Q_j of the IMC abstraction \mathcal{I} , we inspect all paths starting from $\langle Q_j, s_0 \rangle$ in $(\mathcal{M}_{\otimes}^A)_u$ to determine which states to refine. Above is an example of a path π_1 . A score is assigned to all states along π_1 as detailed in Section VI. In particular, if π_1 reaches a member of a potential BSCC, a score is assigned to the states which could possibly destroy the BSCC under refinement. These states are shown in blue and have outgoing transitions with lower bound 0.

winning or losing component for any refinement of the product IMC. The potential ‘‘certainty gain’’ one can hope for by refining $\langle Q_j, s_i \rangle = \text{Last}(\pi)$ is added to the score of Q_j . The path is continued to an unexplored state.

- 2) Refine the states in P_k with scores above a user-defined threshold to generate P_{k+1} .
- 3) Generate an IMC abstraction of the system with respect to P_{k+1} , perform model-checking and compute $V_?$.
- 4) If $V_? > V_{stop}$, return to step 1. Else, terminate.

It is not difficult to construct examples demonstrating that the volume of uncertain states $V_?$ need not decrease monotonically at each step of the refinement algorithm using our abstraction technique. This is because, when a parent state is refined to two children states, the sum of the upper transition bounds for the children states may be greater than the upper transition bound of the original parent state. Nevertheless, when \mathcal{F} is continuous, the size of the reachable sets, and consequently the error in the transitions, approaches zero as the grid size decreases. Thus, in the limit, the volume of uncertain states $V_?$ decreases to zero.

A common refinement approach consists in splitting the chosen states in the partition in half along their greatest dimension. As the scoring procedure in Algorithm 3 may select the entire state-space of IMC \mathcal{I} for refinement, the worst-case growth of the size of the product IMC $\mathcal{I} \otimes \mathcal{A}$ is exponential and scales in $\mathcal{O}(|S| \cdot 2^{|Q|})$, where $|S|$ and $|Q|$ are the number of states of automaton \mathcal{A} and IMC \mathcal{I} respectively.

However, because this path-based scoring procedure aims to target states which are most likely to reduce the volume

of undecided states in the partition with respect to the specification under consideration, our refinement algorithm tends to focus on specific regions of the state-space, as shown in the next section, with the effect of mitigating state explosion. In particular, we were able to achieve lower volumes V_{stop} before the number of states in the partition became prohibitive compared to our naive refinement method in [13] which systematically refined all uncertain states.

All algorithms discussed in this paper are implemented in a python package available at <https://github.com/gtfactslab/TACStochasticVerification>.

VII. CASE STUDY

We now apply our verification and refinement procedure in a case study. We consider a nonlinear, monotone bistable switch system with additive disturbance and governing equations

$$\begin{aligned} x_1[k+1] &= x_1[k] + (-ax_1[k] + x_2[k]) \cdot \Delta T + w_1 \\ x_2[k+1] &= x_2[k] + \left(\frac{(x_1[k])^2}{(x_1[k])^2 + 1} - bx_2[k] \right) \cdot \Delta T + w_2, \end{aligned} \quad (17)$$

where we assume w_1 and w_2 to be independent truncated Gaussian random variables sampled at each time step. $w_1 \sim \mathcal{N}(\mu = -0.3; \sigma^2 = 0.1)$ and is truncated on $[-0.4, -0.2]$; w_2 is identical. To keep the system self-contained in D , we assume that any time the disturbance would push the trajectory outside of D , it is actually maintained on the boundary of D . This assumption reflects the behavior of systems with bounded capacity where the state variables are restricted to some intervals. We choose $a = 1.3$, $b = 0.25$ and $\Delta T = 0.05$. The deterministic piece of the system has two stable equilibria at $(0, 0)$ and $(2.71, 3.52)$ and one unstable equilibrium. We seek to verify (17) on a domain D , with initial rectangular partition P depicted in Fig. 6 (Top) and Fig. 7 (Top), against the probabilistic LTL specifications

$$\begin{aligned} \phi_1 &= \mathcal{P}_{\geq 0.80}[\Box((\neg A \wedge \bigcirc A) \rightarrow (\bigcirc \bigcirc A \wedge \bigcirc \bigcirc \bigcirc A))] \\ \phi_2 &= \mathcal{P}_{\leq 0.90}[(\Diamond \Box A \rightarrow \Diamond B) \wedge (\Diamond C \rightarrow \Box \neg B)]. \end{aligned}$$

Specification ϕ_1 translates in natural language to ‘‘trajectories that have more than a 80% chance of remaining in an A state for at least 2 more time steps when entering an A state’’. Specification ϕ_2 translates to ‘‘trajectories that have less than a 90% chance of reaching a B state if it eventually always remain in A , and of always staying outside of B if it reaches a C state’’. Their Rabin automaton representations contain 5 and 7 states respectively. We perform verification with stopping criterion $V_{stop} = 0.13$ for ϕ_1 and $V_{stop} = 0.1$ for ϕ_2 . To construct IMC abstractions of this system, we use the technique shown in Section IV. Graph search is based on Section V-B and we compute reachability bounds applying the algorithm in [12]. Upon verification, we select states with an uncertainty score as defined in Section VI that is greater than 10% of the highest score for refinement. Selected states are split into two rectangles along their largest dimension to keep the new partition rectangular. The procedure was conducted on a 3.3 GHz Intel Core i7 with 8 GB of memory using Python.

For ϕ_1 , the refinement algorithm produced 3531 states and terminated in 1h56min after 12 refinement steps. For ϕ_2 , it

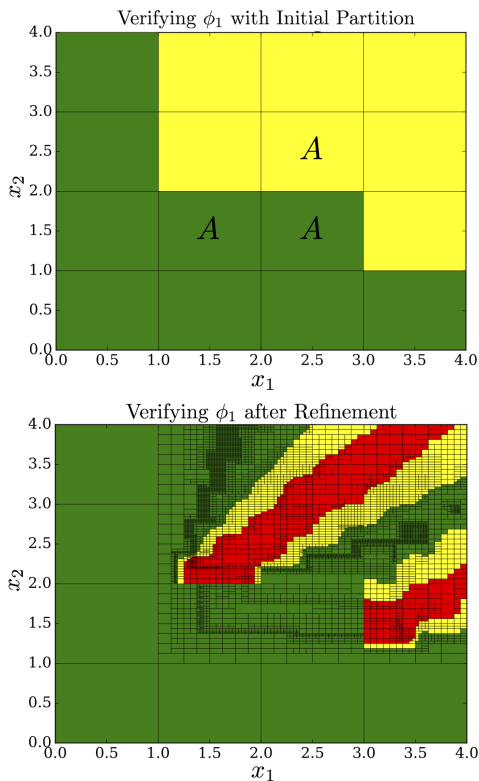


Fig. 6. Initial verification of a partition of domain D for specification ϕ_1 (Top), and verification of final partition (Bottom). States satisfying ϕ_1 are in green, states violating ϕ_1 are in red, undecided states are yellow.

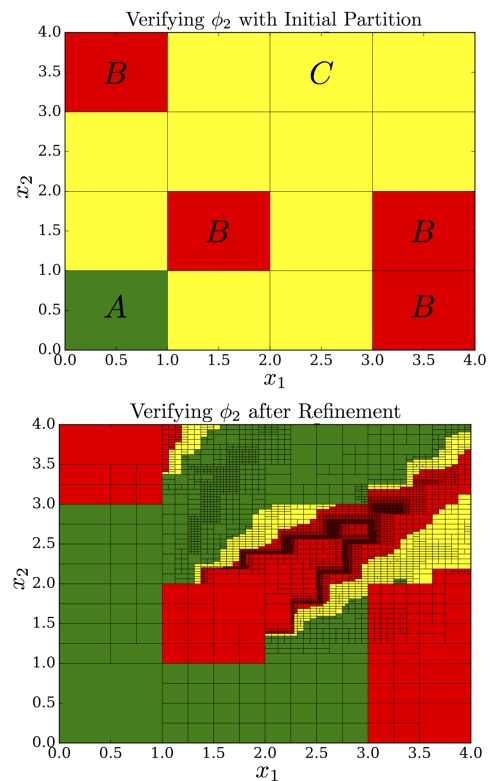


Fig. 7. Initial verification of a partition of domain D for specification ϕ_2 (Top), and verification of final partition (Bottom). States satisfying ϕ_2 are in green, states violating ϕ_2 are in red, undecided states are yellow.

generated 4845 states and terminated in 3h15min after 13 steps. The final partitions are shown in Fig. 6 (Bottom) and Fig. 7 (Bottom). Our new method outperforms the algorithm we propose in [13] which refines all undecided states at each refinement step: for instance, for ϕ_1 , [13] achieves $V_? = 0.2137$ in 2 hours 58 min and 11 steps. Our algorithm non-uniformly refined the initial partition across the state-space. In the first example, the boundary between regions which can and cannot reach an A state are heavily targeted, as well as boundaries between regions which could keep the system in an A state for one and two time steps. In the second example, the edges of a region leading to A via B are refined the most, as this region is critical with respect to ϕ_2 . Although these two examples share the same dynamics, our algorithm generates very different partitions depending on the specification. Therefore, specification-free gridding approaches are likely to perform conservatively for these examples.

Because this procedure produces different partitions for different temporal objectives, our algorithm is suited for the verification of systems against well-identified specifications which are known a priori, while FAUST², which provides error guarantees with respect to entire classes of temporal logic formulas, better accommodates situations requiring an analysis against numerous distinct specifications. Furthermore, unlike our work, FAUST² only permits verification for specifications with bounded-time temporal operators [34].

On the other hand, the toolbox StocHy [34] allows for the verification of certain unbounded-time operators through the use of IMC abstractions. However, because its verification

algorithm is based on [12], it solely accommodates specification belonging to a fragment of LTL and not all ω -regular properties. In addition, StocHy employs the IMC abstraction technique presented in [35] which applies only to affine-in-disturbance linear systems, while the abstraction method shown in this work and used in the case study involves the wide class of stochastic mixed monotone dynamics.

Finally, though similar in flavor, the objective of our verification technique is not identical to the one of FAUST² and StocHy. The latter aim to create partitions engendering a user-defined abstraction error, whereas the goal of this work is to obtain a small volume of uncertain abstraction states with respect to a probabilistic ω -regular specification. These uncertain abstraction states are hyperrectangular regions of the original system domain for which the satisfaction of the specification cannot be concluded given the current abstraction partitioning. On the other hand, note that low-error abstractions can generate a high volume of uncertain continuous states while high-error abstractions can produce a low volume of uncertain continuous states, depending on the threshold p_{sat} .

VIII. CONCLUSION

In this paper, we described an algorithm for performing verification against ω -regular properties in continuous state stochastic systems. The proposed approach relies on computing a finite-state abstraction of the original system in the form of an IMC and can accommodate classes of specifications not previously covered in the literature. We have developed an

efficient procedure for computing the IMC of a mixed monotone system with affine disturbance over rectangular partitions. Furthermore, we presented a specification-guided strategy for refining a finite partition of the continuous domain until a precision threshold has been met. Our technique resolves qualitative issues highlighted in the literature for abstraction-based methods by targeting states that are likely to confirm or destroy winning and losing components in a product IMC. We showed the practicality of this approach in two examples.

APPENDIX I PROOFS OF SECTION IV

A. Proof of Proposition 2

Proof: By Proposition 1, we observe

$$\{\mathcal{F}(x) : x \in Q_1\} \subseteq \{z : g(a^1, b^1) \leq z \leq g(b^1, a^1)\}. \quad (18)$$

To prove (7), we have

$$\begin{aligned} & \min_{x \in Q_1} Pr(\mathcal{F}(x) + w \in Q_2) \\ & \geq \min_{z: g(a^1, b^1) \leq z \leq g(b^1, a^1)} Pr(z + w \in Q_2) \end{aligned} \quad (19)$$

$$= \min_{z: g(a^1, b^1) \leq z \leq g(b^1, a^1)} \prod_{i=1}^n Pr(z_i + w_i \in [a_i^2, b_i^2]) \quad (20)$$

$$= \prod_{i=1}^n \min_{z_i: g_i(a^1, b^1) \leq z_i \leq g_i(b^1, a^1)} Pr(z_i + w_i \in [a_i^2, b_i^2]) \quad (21)$$

where (19) follows from (18), (20) follows from the mutual independence of all components of w in Assumption 2, and (21) holds because $g(a^1, b^1) \leq z \leq g(b^1, a^1)$ if and only if $g_i(a^1, b^1) \leq z_i \leq g_i(b^1, a^1)$ for all $i = 1, \dots, n$. Then (7) holds since $Pr(z_i + w_i \in [a_i^2, b_i^2]) = \int_{a_i^2}^{b_i^2} f_{w_i}(x - z_i) dx$. By a symmetric argument replacing min with max, (8) holds. ■

B. Proof Lemma 1

Proof: For $s \in \mathbb{R}$, define $H(s) = \int_a^b f_\omega(x - s) dx$. We claim

$$H(s_{max}) = \max_{s \in \mathbb{R}} H(s),$$

and that, for all $s_1, s_2 \in \mathbb{R}$ such that $|s_{max} - s_1| \geq |s_{max} - s_2|$, it holds that $H(s_1) \leq H(s_2)$, that is, $H(s)$ monotonically decreases as $|s_{max} - s|$ increases. Assuming the claim to be true, it follows that $\max_{s \in [r_1, r_2]} H(s) = H(s_{max}^r)$ and $\min_{s \in [r_1, r_2]} H(s) = H(s_{min}^r)$, i.e., (9) and (10), completing the proof. To prove the claim, we have, for all $s \in \mathbb{R}$,

$$\begin{aligned} & H(s_{max}) - H(s) \\ &= \int_a^b f_\omega(x - s_{max}) dx - \int_a^b f_\omega(x - s) dx \\ &= \int_{a-s_{max}}^{b-s_{max}} f_\omega(x) dx - \int_{a-s}^{b-s} f_\omega(x) dx \\ &= \int_{a-s_{max}}^{a-s} f_\omega(x) dx - \int_{b-s_{max}}^{b-s} f_\omega(x) dx \\ &= \int_0^{s_{max}-s} \left[f_\omega(x + \frac{a-b}{2} - c) - f_\omega(x + \frac{b-a}{2} - c) \right] dx. \end{aligned} \quad (22)$$

Moreover, because f_ω is symmetric and unimodal with mode c , $f_\omega(x + \frac{a-b}{2} - c) - f_\omega(x + \frac{b-a}{2} - c)$ is an odd function of x and is negative for $x > 0$ and positive for $x < 0$. Therefore, the integral in (22) is nonnegative and monotonically decreases as $|s_{max} - s|$ increases, thus proving the claim. ■

C. Proof of Theorem 1

Proof: For all $Q_j, Q_\ell \in P$ and $i = 1, \dots, n$,

$$\begin{aligned} & \min_{z_i \in [g_i(a^j, b^j), g_i(b^j, a^j)]} \int_{a_i^\ell}^{b_i^\ell} f_{w_i}(x - z_i) dx \\ &= \int_{a_i^\ell}^{b_i^\ell} f_{w_i}(x_i - s_{i, min}^{j \rightarrow \ell}) dx_i, \\ & \max_{z_i \in [g_i(a^j, b^j), g_i(b^j, a^j)]} \int_{a_i^\ell}^{b_i^\ell} f_{w_i}(x - z_i) dx \\ &= \int_{a_i^\ell}^{b_i^\ell} f_{w_i}(x_i - s_{i, max}^{j \rightarrow \ell}) dx_i, \end{aligned}$$

by Lemma 1. Then, by Proposition 2,

$$\min_{x \in Q_j} Pr(\mathcal{F}(x) + w \in Q_\ell) \geq \prod_{i=1}^n \int_{a_i^\ell}^{b_i^\ell} f_{w_i}(x_i - s_{i, min}^{j \rightarrow \ell}) dx_i \quad (23)$$

$$\max_{x \in Q_j} Pr(\mathcal{F}(x) + w \in Q_\ell) \leq \prod_{i=1}^n \int_{a_i^\ell}^{b_i^\ell} f_{w_i}(x_i - s_{i, max}^{j \rightarrow \ell}) dx_i, \quad (24)$$

so that (11)–(12) implies (4)–(5). Furthermore, (23)–(24) implies $\check{T}(Q_j, Q_\ell) \leq \hat{T}(Q_j, Q_\ell)$ and (4)–(5) implies (3) so that $\mathcal{I} = (P, \check{T}, \hat{T})$ is a valid IMC, concluding the proof. ■

APPENDIX II LEMMAS AND PROOFS OF SECTION V

Lemma 2: [1] For any infinite sequence of states $\pi = q_0 q_1 q_2 \dots$ in a Markov Chain, there exists an index $i \geq 0$ such that q_i belongs to a BSCC.

Corollary 2: For any initial state $\langle Q_i, s_0 \rangle$ in an induced product MC \mathcal{M}_\otimes^A ,

$$\mathcal{P}_{\mathcal{M}_\otimes^A}(\langle Q_i, s_0 \rangle \models \diamond U^A) + \mathcal{P}_{\mathcal{M}_\otimes^A}(\langle Q_i, s_0 \rangle \models \diamond U^N) = 1.$$

Lemma 3: For any initial state $\langle Q_i, s_0 \rangle$ in an induced product MC \mathcal{M}_\otimes^A ,

$$\mathcal{P}_{\mathcal{M}_\otimes^A}(\langle Q_i, s_0 \rangle \models \diamond WC) + \mathcal{P}_{\mathcal{M}_\otimes^A}(\langle Q_i, s_0 \rangle \models \diamond LC) = 1.$$

Proof: This lemma follows from Corollary 2. ■

Lemma 4: Let \mathcal{I} be an IMC and let \mathcal{M}_1 and \mathcal{M}_2 be two MCs induced by \mathcal{I} where the set B is a BSCC for both. If C_1 and C_2 are the sets of states such that $\mathcal{P}_{\mathcal{M}_1}(C_1 \models \diamond B) = 1$ and $\mathcal{P}_{\mathcal{M}_2}(C_2 \models \diamond B) = 1$, then there exists a MC \mathcal{M}_3 induced by \mathcal{I} such that $\mathcal{P}_{\mathcal{M}_3}((C_1 \cup C_2) \models \diamond B) = 1$.

Proof: Let T_1 and T_2 denote the transition matrices of \mathcal{M}_1 and \mathcal{M}_2 respectively, and Q denote the set of states in \mathcal{I} . Consider an induced MC \mathcal{M}_3 where $T_3(Q_i, Q_j) = T_1(Q_i, Q_j) \forall Q_i \in C_1$ and $\forall Q_j \in Q$, and $T_3(Q_i, Q_j) = T_2(Q_i, Q_j) \forall Q_i \in C_2 \setminus (C_1 \cap C_2)$ and $\forall Q_j \in Q$. By assumption, any state in C_1 reaches B with probability 1,

while all states in $C_2 \setminus (C_1 \cap C_2)$ reach $B \cup (C_1 \cap C_2)$ with probability 1. Since $\mathcal{P}_{\mathcal{M}_3}((C_1 \cap C_2) \models \diamond B) = 1$ by construction, we have $\mathcal{P}_{\mathcal{M}_3}((C_1 \cup C_2) \models \diamond B) = 1$. ■

Lemma 5: Let $\mathcal{I} \otimes \mathcal{A}$ be a product IMC and $(LC)_i$ be the losing components of any product MC $(\mathcal{M}_{\otimes}^A)_i$ induced by $\mathcal{I} \otimes \mathcal{A}$. There exists a set of product MCs induced by $\mathcal{I} \otimes \mathcal{A}$ with losing components $(LC)_L$ and such that $(LC)_i \subseteq (LC)_L$.

Proof: We proved in [26] that any product IMC induces a set of MCs with a largest set of non-accepting BSCCs. Lemma 5 is deduced from this fact and Lemma 4. ■

Lemma 6: Let $\mathcal{I} \otimes \mathcal{A}$ be a product IMC. Let $(\mathcal{M}_{\otimes}^A)_1$ and $(\mathcal{M}_{\otimes}^A)_2$ be two product MCs induced by $\mathcal{I} \otimes \mathcal{A}$ with sets of accepting BSCC U_1^A and U_2^A respectively. There exists a set of product MCs induced by $\mathcal{I} \otimes \mathcal{A}$ with winning components $(WC)_3$ and such that $(U_1^A \cup U_2^A) \subseteq (WC)_3$.

Proof: Let T_1 and T_2 denote the transition matrices of $(\mathcal{M}_{\otimes}^A)_1$ and $(\mathcal{M}_{\otimes}^A)_2$ respectively, and Q denote the set of states in $\mathcal{I} \otimes \mathcal{A}$. Assume $U_1^A \cap U_2^A = \emptyset$. There exists a set of product MCs induced by $\mathcal{I} \otimes \mathcal{A}$ such that $U_1^A \cup U_2^A$ are accepting BSCCs (see [26]), and thus winning components. If $U_1^A \cap U_2^A \neq \emptyset$, consider the set of all product MCs $(\mathcal{M}_{\otimes}^A)_i$ induced by $\mathcal{I} \otimes \mathcal{A}$ such that, for all transition matrices T_i of the product MCs in this set, $T_i(Q_i, Q_j) = T_1(Q_i, Q_j) \forall Q_i \in U_1^A$ and $\forall Q_j \in Q$, and $T_i(Q_i, Q_j) = T_2(Q_i, Q_j) \forall Q_i \in U_2^A \setminus (U_1^A \cap U_2^A)$ and $\forall Q_j \in Q$. Clearly, U_1^A is an accepting BSCC in all $(\mathcal{M}_{\otimes}^A)_i$. For all $(\mathcal{M}_{\otimes}^A)_i$, it holds that $\mathcal{P}_{(\mathcal{M}_{\otimes}^A)_i}((U_2^A \setminus (U_1^A \cap U_2^A)) \models \diamond(U_1^A \cap U_2^A)) = 1$, since U_2^A is a BSCC for the same probability assignments in $(\mathcal{M}_{\otimes}^A)_2$. Thus, $U_1^A \cup U_2^A$ are winning components with respect to all $(\mathcal{M}_{\otimes}^A)_i$. ■

Lemma 7: Let $\mathcal{I} \otimes \mathcal{A}$ be a product IMC and $(WC)_i$ be the winning components of any product MC $(\mathcal{M}_{\otimes}^A)_i$ induced by $\mathcal{I} \otimes \mathcal{A}$. There exists a set of product MCs induced by $\mathcal{I} \otimes \mathcal{A}$ with winning component $(WC)_L$ and such that $(WC)_i \subseteq (WC)_L$.

Proof: This lemma follows from Lemmas 4 and 6. ■

Lemma 8: Let $\mathcal{I} \otimes \mathcal{A}$ be a product IMC. Let $(\mathcal{M}_{\otimes}^A)_1$ and $(\mathcal{M}_{\otimes}^A)_2$ be two product MCs induced by $\mathcal{I} \otimes \mathcal{A}$ with winning components $(WC)_1$ and $(WC)_2$ respectively, and such that $(WC)_2 \subseteq (WC)_1$. Also, their losing components $(LC)_1$ and $(LC)_2$ are such that $(LC)_1 = (LC)_2 = LC$ and their respective transition matrices T_1 and T_2 satisfy $T_1(Q_i, Q_j) = T_2(Q_i, Q_j) \forall Q_i \in (Q \times S) \setminus ((WC)_1 \cup LC)$ and $\forall Q_j \in (Q \times S)$. The sets of accepting BSCCs of $(\mathcal{M}_{\otimes}^A)_1$ and $(\mathcal{M}_{\otimes}^A)_2$ are denoted by U_1^A and U_2^A respectively. For any initial state $\langle Q_i, s_0 \rangle$, it holds that

$$\mathcal{P}_{(\mathcal{M}_{\otimes}^A)_1}(\langle Q_i, s_0 \rangle \models \diamond U_1^A) \geq \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_2}(\langle Q_i, s_0 \rangle \models \diamond U_2^A).$$

Proof: For any initial state $\langle Q_i, s_0 \rangle \in LC$, it holds that $\mathcal{P}_{(\mathcal{M}_{\otimes}^A)_1}(\langle Q_i, s_0 \rangle \models \diamond U_1^A) = \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_2}(\langle Q_i, s_0 \rangle \models \diamond U_2^A) = 0$. For any initial state $\langle Q_i, s_0 \rangle \in ((WC)_1 \cap (WC)_2)$, it holds that $\mathcal{P}_{(\mathcal{M}_{\otimes}^A)_1}(\langle Q_i, s_0 \rangle \models \diamond U_1^A) = \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_2}(\langle Q_i, s_0 \rangle \models \diamond U_2^A) = 1$. For any initial state $\langle Q_i, s_0 \rangle \in ((WC)_1 \setminus (WC)_2)$, it holds that $\mathcal{P}_{(\mathcal{M}_{\otimes}^A)_1}(\langle Q_i, s_0 \rangle \models \diamond U_1^A) = 1 \geq \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_2}(\langle Q_i, s_0 \rangle \models \diamond U_2^A)$. For any initial state $\langle Q_i, s_0 \rangle \in$

$(Q \times S) \setminus ((WC)_1 \cup LC)$ (denoted by H for clarity), we have

$$\begin{aligned} & \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_1}(\langle Q_i, s_0 \rangle \models \diamond U_1^A) = \\ & T_1(\langle Q_i, s_0 \rangle, (WC)_1 \setminus (WC)_2) \cdot \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_1}((WC)_1 \setminus (WC)_2 \models \diamond U_1^A) + T_1(\langle Q_i, s_0 \rangle, (WC)_2) \cdot \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_1}((WC)_2 \models \diamond U_1^A) \\ & + \sum_{Q_j \in H} T_1(\langle Q_i, s_0 \rangle, Q_j) \cdot \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_2}(Q_j \models \diamond U_2^A) \\ & \geq \\ & \sum_{Q_j \in (WC)_1 \setminus (WC)_2} T_2(\langle Q_i, s_0 \rangle, Q_j) \cdot \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_2}(Q_j \models \diamond U_2^A) \\ & + T_2(\langle Q_i, s_0 \rangle, (WC)_2) \cdot \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_2}((WC)_2 \models \diamond U_2^A) \\ & + \sum_{Q_j \in H} T_2(\langle Q_i, s_0 \rangle, Q_j) \cdot \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_2}(Q_j \models \diamond U_2^A) \\ & = \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_2}(\langle Q_i, s_0 \rangle \models \diamond U_2^A) \end{aligned}$$

based on the transition matrices assumptions. ■

Lemma 9: Let $\mathcal{I} \otimes \mathcal{A}$ be a product IMC. Let $(\mathcal{M}_{\otimes}^A)_1$ and $(\mathcal{M}_{\otimes}^A)_2$ be two product MCs induced by $\mathcal{I} \otimes \mathcal{A}$ with losing components $(LC)_1$ and $(LC)_2$ respectively, and such that $(LC)_2 \subseteq (LC)_1$. Also, their winning components $(WC)_1$ and $(WC)_2$ are such that $(WC)_1 = (WC)_2 = WC$ and their respective transition matrices T_1 and T_2 satisfy $T_1(Q_i, Q_j) = T_2(Q_i, Q_j) \forall Q_i \in (Q \times S) \setminus ((LC)_1 \cup WC)$ and $\forall Q_j \in (Q \times S)$. The sets of non-accepting BSCCs of $(\mathcal{M}_{\otimes}^A)_1$ and $(\mathcal{M}_{\otimes}^A)_2$ are denoted by U_1^N and U_2^N respectively. For any initial state $\langle Q_i, s_0 \rangle$, it holds that

$$\mathcal{P}_{(\mathcal{M}_{\otimes}^A)_1}(\langle Q_i, s_0 \rangle \models \diamond U_1^N) \geq \mathcal{P}_{(\mathcal{M}_{\otimes}^A)_2}(\langle Q_i, s_0 \rangle \models \diamond U_2^N).$$

Proof: The proof is identical to the one of Lemma 8. ■

Lemma 10: Let $\mathcal{I} \otimes \mathcal{A}$ be a product IMC with permanent and largest sets $(WC)_P$, $(LC)_P$, $(WC)_L$ and $(LC)_L$ as previously defined. There exists a set of induced MCs of $\mathcal{I} \otimes \mathcal{A}$ whose sets of winning and losing components are $(WC)_P$ and $(LC)_L$, and a set of induced MCs whose sets of losing and winning components are $(LC)_P$ and $(WC)_L$.

Proof: Consider the set \mathcal{C} of all induced MCs of $\mathcal{I} \otimes \mathcal{A}$ whose set of losing components is $(LC)_L$. For any $\langle Q_i, s_0 \rangle \in ((WC)_? \setminus (LC)_L)$, consider an induced product MC $\mathcal{M} \in \mathcal{C}$ with transition matrix T such that $\langle Q_i, s_0 \rangle$ is not a winning component of \mathcal{M} . Such an induced product MC always exists by the definition of $(WC)_?$ and Lemma 8. Denote by $(WC)_?^{\mathcal{M}}$ the winning components of \mathcal{M} which also belong to $(WC)_?$. There exists an induced product MC \mathcal{M}' with transition matrix T' such that, for all $q_i \in (WC)_?^{\mathcal{M}}$, $\mathcal{P}_{\mathcal{M}'}(q_i \models \diamond \neg((WC)_?^{\mathcal{M}} \cup (WC)_P)) > 0$, otherwise $q_i \in (WC)_P$, which is a contradiction. Consider the induced product MC $\mathcal{M}'' \in \mathcal{C}$ with transition matrix T'' such that $T''(q_i, q_j) = T'(q_i, q_j)$ for all $q_i \in (WC)_?^{\mathcal{M}}$ and $q_j \in (Q \times S)$, and $T'' = T$ for all other transitions. The sets of winning and losing components of \mathcal{M}'' are $(WC)_P$ and $(LC)_L$, proving the claim. The proof with respect to $(LC)_P$ and $(WC)_L$ is identical. ■

REFERENCES

- [1] C. Baier, J.-P. Katoen, and K. G. Larsen, *Principles of model checking*. MIT press, 2008.

- [2] A. Aziz, K. Sanwal, V. Singhal, and R. Brayton, "Model-checking continuous-time markov chains," *ACM Transactions on Computational Logic (TOCL)*, vol. 1, no. 1, pp. 162–170, 2000.
- [3] C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen, "Model-checking algorithms for continuous-time Markov chains," *IEEE Transactions on software engineering*, vol. 29, no. 6, pp. 524–541, 2003.
- [4] M. Kwiatkowska, G. Norman, and D. Parker, "Prism 4.0: Verification of probabilistic real-time systems," in *International conference on computer aided verification*. Springer, 2011, pp. 585–591.
- [5] A. Abate, A. D'Innocenzo, and M. D. Di Benedetto, "Approximate abstractions of stochastic hybrid systems," *IEEE Transactions on Automatic Control*, vol. 56, no. 11, pp. 2688–2694, 2011.
- [6] S. Soudjani, C. Gevaerts, and A. Abate, "Faust 2: Formal abstractions of uncountable-state stochastic processes," *arXiv preprint arXiv:1403.3286*, 2014.
- [7] A. Abate, A. D'Innocenzo, M. D. Di Benedetto, and S. S. Sastry, "Markov set-chains as abstractions of stochastic hybrid systems," in *International Workshop on Hybrid Systems: Computation and Control*. Springer, 2008, pp. 1–15.
- [8] A. Abate, J.-P. Katoen, and A. Mereacre, "Quantitative automata model checking of autonomous stochastic hybrid systems," in *Proceedings of the 14th international conference on Hybrid systems: computation and control*. ACM, 2011, pp. 83–92.
- [9] I. Tkachev and A. Abate, "Formula-free finite abstractions for linear temporal verification of stochastic hybrid systems," in *Proceedings of the 16th international conference on Hybrid systems: computation and control*. ACM, 2013, pp. 283–292.
- [10] E. M. Hahn, A. Hartmanns, H. Hermanns, and J.-P. Katoen, "A compositional modelling and analysis framework for stochastic hybrid systems," *Formal Methods in System Design*, vol. 43, no. 2, pp. 191–232, 2013.
- [11] M. Fränzle, E. M. Hahn, H. Hermanns, N. Wolovick, and L. Zhang, "Measurability and safety verification for stochastic hybrid systems," in *Proceedings of the 14th international conference on Hybrid systems: computation and control*. ACM, 2011, pp. 43–52.
- [12] M. Lahijanian, S. B. Andersson, and C. Belta, "Formal verification and synthesis for discrete-time stochastic systems," *IEEE Transactions on Automatic Control*, vol. 60, no. 8, pp. 2031–2045, 2015.
- [13] M. Dutreix and S. Coogan, "Efficient Verification for Stochastic Mixed Monotone Systems," in *International Conference on Cyber-Physical Systems*, 2018.
- [14] I. O. Kozine and L. V. Utkin, "Interval-valued finite Markov chains," *Reliable computing*, vol. 8, no. 2, pp. 97–113, 2002.
- [15] H. Smith, "Global stability for mixed monotone systems," *Journal of Difference Equations and Applications*, vol. 14, no. 10-11, pp. 1159–1164, 2008.
- [16] S. Coogan and M. Arcak, "Efficient finite abstraction of mixed monotone systems," in *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*. ACM, 2015, pp. 58–67.
- [17] M. W. Hirsch, "Systems of differential equations that are competitive or cooperative II: Convergence almost everywhere," *SIAM Journal on Mathematical Analysis*, vol. 16, no. 3, pp. 423–439, 1985.
- [18] H. L. Smith, *Monotone dynamical systems: An introduction to the theory of competitive and cooperative systems*. American Mathematical Society, 1995.
- [19] D. Angeli and E. Sontag, "Monotone control systems," *IEEE Transactions on Automatic Control*, vol. 48, no. 10, pp. 1684–1698, 2003.
- [20] E. D. Sontag, "Monotone and near-monotone biochemical networks," *Systems and Synthetic Biology*, vol. 1, no. 2, pp. 59–87, 2007.
- [21] G. Gomes, R. Horowitz, A. A. Kurzhanskiy, P. Varaiya, and J. Kwon, "Behavior of the cell transmission model and effectiveness of ramp metering," *Transportation Research Part C: Emerging Technologies*, vol. 16, no. 4, pp. 485–513, 2008. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0968090X0700085X>
- [22] E. Lovisari, G. Como, and K. Savla, "Stability of monotone dynamical flow networks," in *Proceedings of the 53rd Conference on Decision and Control*, 2014, pp. 2384–2389.
- [23] S. Coogan, M. Arcak, and A. A. Kurzhanskiy, "Mixed monotonicity of partial first-in-first-out traffic flow models," in *IEEE Conference on Decision and Control*, 2016, pp. 7611–7616.
- [24] K. Y. Rozier, "Linear temporal logic symbolic model checking," *Computer Science Review*, vol. 5, no. 2, pp. 163–203, 2011.
- [25] C. Baier, M. Größer, M. Leucker, B. Bollig, and F. Ciesinski, "Controller synthesis for probabilistic systems," in *Exploring New Frontiers of Theoretical Informatics*. Springer, 2004, pp. 493–506.
- [26] M. Dutreix and S. Coogan, "Satisfiability Bounds for ω -regular Properties in Interval-valued Markov Chains," in *Proceedings of the 57th IEEE Conference on Decision and Control*. (to appear), 2018.
- [27] K. Sen, M. Viswanathan, and G. Agha, "Model-checking Markov chains in the presence of uncertainties," in *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 2006, pp. 394–410.
- [28] M. Hirsch and H. Smith, "Monotone dynamical systems," *Handbook of differential equations: Ordinary differential equations*, vol. 2, pp. 239–357, 2005.
- [29] J. Klein and C. Baier, "Experiments with deterministic ω -automata for formulas of linear temporal logic," *Theoretical Computer Science*, vol. 363, no. 2, pp. 182–195, 2006.
- [30] T. Babiak, F. Blahoudek, M. Křetínský, and J. Strejček, "Effective translation of LTL to deterministic Rabin automata: Beyond the (F, G)-fragment," in *International Symposium on Automated Technology for Verification and Analysis*. Springer, 2013, pp. 24–39.
- [31] T. Chen, T. Han, and M. Kwiatkowska, "On the complexity of model checking interval-valued discrete time Markov chains," *Information Processing Letters*, vol. 113, no. 7, pp. 210–216, 2013.
- [32] K. Chatterjee, K. Sen, and T. Henzinger, "Model-checking ω -regular properties of interval markov chains," *Foundations of Software Science and Computational Structures*, pp. 302–317, 2008.
- [33] L. De Alfaro, "Computing minimum and maximum reachability times in probabilistic systems," in *International Conference on Concurrency Theory*. Springer, 1999, pp. 66–81.
- [34] N. Cauchi, K. Degiorgio, and A. Abate, "Stochy: automated verification and synthesis of stochastic processes," *arXiv preprint arXiv:1901.10287*, 2019.
- [35] N. Cauchi, L. Laurenti, M. Lahijanian, A. Abate, M. Kwiatkowska, and L. Cardelli, "Efficiency through uncertainty: scalable formal synthesis for stochastic hybrid systems," in *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control*. ACM, 2019, pp. 240–251.



Maxence Dutreix received the B.S. degree in Engineering Physics from the University of California, San Diego, CA, USA, in 2015. He is currently a Ph.D. student in Electrical and Computer Engineering at the Georgia Institute of Technology, Atlanta, GA, USA. His research interest is in the area of verification and synthesis for stochastic dynamical systems.



Samuel Coogan is an Assistant Professor at Georgia Tech in the School of Electrical and Computer Engineering and the School of Civil and Environmental Engineering. Prior to joining Georgia Tech in July 2017, he was an assistant professor in the Electrical Engineering Department at UCLA from 2015–2017. He received the B.S. degree in Electrical Engineering from Georgia Tech and the M.S. and Ph.D. degrees in Electrical Engineering from the University of California, Berkeley. His research is in the area of dynamical systems and autonomy and focuses on developing scalable tools for verification and control of networked, cyber-physical systems with an emphasis on transportation systems. He received a CAREER award from NSF in 2018, a Young Investigator Award from the Air Force Office of Scientific Research in 2018, and the Outstanding Paper Award for the IEEE Transactions on Control of Network Systems in 2017.