# AUTOMATED VERIFICATION AND SYNTHESIS OF STOCHASTIC HYBRID SYSTEMS: A SURVEY

ABOLFAZL LAVAEI[1], SADEGH SOUDJANI[1], ALESSANDRO ABATE[2], AND MAJID ZAMANI[3,4]

Abstract. Stochastic hybrid systems have received significant attentions as a relevant modelling framework describing many systems, from engineering to the life sciences: they enable the study of numerous applications, including transportation networks, biological systems and chemical reaction networks, smart energy and power grids, and beyond. Automated verification and policy synthesis for stochastic hybrid systems can be inherently challenging: this is due to the heterogeneity of their dynamics (presence of continuous and discrete components), the presence of uncertainty, and in some applications the large dimension of state and input sets. Over the past few years, a few hundred articles have investigated these models, and developed diverse and powerful approaches to mitigate difficulties encountered in the analysis and synthesis of such complex stochastic systems. In this survey, we overview the most recent results in the literature and discuss different approaches, including *(in)finite abstractions*, *verification and synthesis for temporal logic specifications*, *stochastic similarity relations*, *(control) barrier certificates*, *compositional techniques*, and a selection of results on *continuous-time stochastic systems*; we finally survey recently developed *software tools* that implement the discussed approaches. Throughout the manuscript we discuss a few open topics to be considered as potential future research directions: we hope that this survey will guide younger researchers through a comprehensive understanding of the various challenges, tools, and solutions in this enticing and rich scientific area.

## 1. Introduction

Stochastic hybrid systems (SHS) concern complex dynamical models combining both digital-computation elements and physical components, tightly interacting with each other in feedback interconnections. SHS models thus comprise discrete dynamics modelling computational components including hardware and software, and continuous dynamics that model the physical system. Due to their broad real-world applications, such as (air) traffic networks, transportation systems, energy networks, process engineering, biological systems, and robotic manufacturing, (cf. Hu et al. (2004), Hespanha & Singh (2005), Prandini & Hu (2008), Singh & Hespanha (2010*b*), Vargas-García & Singh (2018), to name a few), over the past few years SHS have gained remarkable attention in the areas of control theory, formal verification, applied mathematics, and performance evaluation, among others. SHS applications have become more complex, with more digital components (e.g., for computation and communication) that interact with physical (analog) parts: this tight interaction causes major difficulties in designing and analyzing these complex systems. Accordingly, the ability to handle the interaction between continuous and discrete dynamics is a prerequisite for providing a rigorous formal framework for formal verification and synthesis of SHS.

Grown first within the area of hybrid systems and of stochastic control, SHS have been first comprehensively presented and widely discussed in the books by Blom et al. (2006) and by Cassandras & Lygeros (2006). The

historical roots in hybrid systems that underpin SHS research have brought to an inter-disciplinary look at these models, with emphasis split between classical dynamical analysis and control synthesis on the one hand, as well as computability and formal verification around rich, high-level specifications on the other. Automated verification and policy synthesis for SHS around high-level temporal requirements, *e.g.,* those expressed as (linear) temporal logic formulae (Pnueli 1977), are the core emphasis of this survey. Given a temporal property of interest for a dynamical model, formal verification is concerned to soundly check whether the desired specification is satisfied. If the underlying model is stochastic, the goal translates in formally quantifying the probability of satisfying the property of interest. A synthesis problem instead concerns dynamical models with the presence of control inputs: the goal is to formally design a controller (also known in different areas as policy, or strategy, or scheduler), which is by and large a state-feedback architecture, to enforce the property of interest. This procedure is also called "correct-by-construction control design", since every step in the controller synthesis procedure comes with a formal guarantee. In a stochastic setting, the key objective is to synthesize a controller that optimizes (e.g., maximizes) the probability of satisfying the given specification. As a result of their intrinsic soundness, formal methods approaches do not require any costly, exhaustive, and possibly unsuccessful post-facto validation, which is needed in many safety-critical, real-world applications.

The intrinsic complexity of SHS models, resulting from the aforementioned interaction of discrete and continuous components, as well as from the presence of uncertainty that is modelled via probability terms, makes it in general difficult - if at all possible - to obtain analytical results in their formal verification or for control synthesis tasks. Hence, verification and policy synthesis for SHS are generally addressed by techniques that either leverage model (finite) approximations, or the use of sufficient conditions for analysis. Accordingly, we categorize these two classes of approaches as (i) discretization-based and (ii) discretization-free techniques.

1.1. **Discretization-based Techniques.** In the analysis of SHS, it is often the case that quantities of interest, such as value functions, or the characterization of optimal policies, are in general not available in a closed (explicit, analytical) form. Therefore, a suitable approach for analyzing SHS is to approximate given ("concrete") SHS models by simpler ones endowed with finite state spaces, also known as "finite abstractions".[1] Finite abstractions of SHS are often in the form of Markov decision processes (MDP), where each discrete state corresponds to a set of continuous states of the concrete SHS model (similarly for inputs). In practice, such finite abstractions can be generated by partitioning state/input sets of the concrete models given some discretization parameters. If the underlying SHS is autonomous, *i.e.,* without control input, the finite MDP is then reduced to a finite Markov chain (MC). The discrete dynamics of the finite abstractions are similarly obtained from those of the concrete continuous models (cf. Fig. 1). Since the obtained abstractions are finite, many algorithmic machineries from computer science (Baier & Katoen 2008) are directly applicable to perform analysis, model checking, or to synthesize controllers maximizing rewards or enforcing complex properties, including those expressed as temporal logic formulae (to be discussed later). A crucial step related

[1]An alternative to the described approach is to generate simpler models that are still uncountable, as approximations of concrete SHS. As examples, these could be obtained by linearizing the (continuous) dynamics, by disregarding additive noise terms, or by reducing the dimension of the concrete models. We place less emphasis on this alternative set of approaches, since they are not in general automated, requiring instead manual solutions which are not generally applicable to SHS (as much as discretization-based techniques are), and since they seldom come with the guarantees that are instead typical of discretization-based techniques.

to these discretization-based techniques is to provide formal guarantees on the obtained abstractions, so that the verification or synthesis results on abstract models can be formally carried over to the original SHS: this is a key feature that characterizes the overall approach. Discretization-based techniques using finite abstractions are schematically illustrated in Fig. 1. As it can be observed, the original SHS is first approximated by a finite abstraction with discrete state and input sets. Then a discrete controller, in the form of a static lookup table or a dynamic controller (with finite memory), is synthesized over the constructed finite abstraction. Finally, the discrete controller is refined back over the original SHS via a *hybrid* interface map that contains states of both original and abstract systems, and the discrete input.

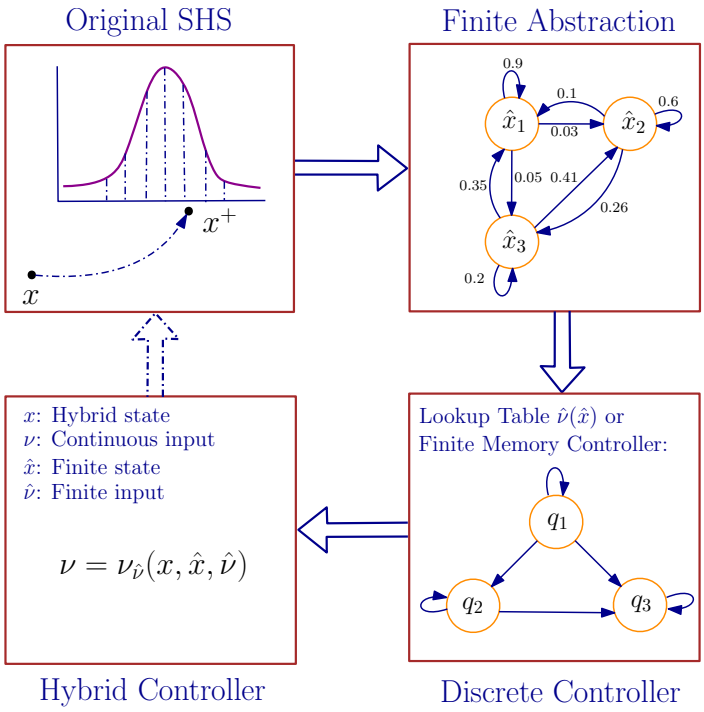


FIGURE 1. Illustration of the procedure underlying discretization-based techniques based on finite abstractions. The discrete controller can be a static lookup table or a dynamic controller (with finite memory).

**Remark 1.1.** *We remark a fundamental difference between the discussed discretization-based techniques for abstractions, which are focused on formally simplifying (SHS) models into abstract models that are amenable to be verified or subject to synthesis tasks; and standard approaches in literature that resort to (e.g., spatial) discretization to provide numerical implementations of algorithms for analysis or synthesis. The latter approaches deal with numerical solutions for quantities of interest, such as value functions or optimal policies. Beyond this fundamental difference, note that the latter approaches often do not come with correctness guarantees.*

1.2. **Discretization-free Techniques.** The techniques discussed in the setting of finite abstractions rely on the discretization (that is, partitioning, or gridding) of state and input/action sets; consequently, they can suffer from an issue known as the *curse of dimensionality*: the complexity of constructing the abstraction grows

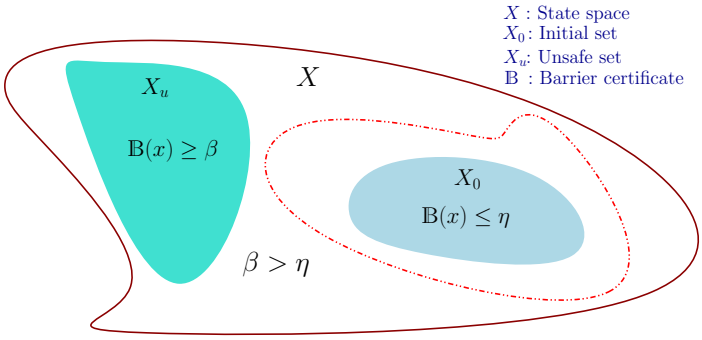


FIGURE 2. Discretization-free techniques can study probabilistic safety based on the construction of control barrier certificates. The (red) dashed line denotes the level set $\mathbb{B}(x) = \eta$.

exponentially with the state/input dimension of the SHS. This critical challenge motivates the development of discretization-free approaches, such as those based on *(control) barrier certificates*. These approaches, which have been recently introduced (over the last 15 years) for verification and/or controller synthesis of complex SHS, should again provide "sufficient results" for the analysis and/or synthesis over the given SHS models. Barrier certificates are Lyapunov-like functions defined over the state space of the system and satisfying a set of inequalities on both the function itself and the one-step transition (or the infinitesimal generator along the flow) of the system. An appropriate level set of a barrier certificate can separate an unsafe region from all system trajectories starting from a given set of initial conditions (cf. Fig. 2) with some probability lower bound. Consequently, the existence of such a function provides a formal probabilistic certificate for system safety. Notice that although barrier certificates are natively employed to ensure the safety of a SHS model, they have also been recently used in the literature to enforce alternative properties, such as temporal requirements (cf. Section 6).

It is worth mentioning that there exist other discretization-free techniques for analysis and controller synthesis in relevant literature, which are mainly based on optimization approaches, such as model predictive control (cf. Subsection 6.2). For instance, stochastic model predictive control (SMPC) (Mesbah 2016) is a widely investigated setup which, however, is not aimed at providing the formal guarantees on verification and controller synthesis of complex SHS on which this survey focuses. Model-order reductions are alternative types of discretization-free techniques, which originate from control literature in the frequency domain, the main goal of which is to establish a closeness relation between the transfer function of the original system and that of its reduced-order model: this is attained by providing a bound on the $\mathcal{H}_2$ norm of the error between transfer functions at given frequencies (Cheng et al. 2017, Yu et al. 2019, 2022). In Section 4 instead, we incorporate model-order reductions in the time domain within the development of *infinite-abstraction* techniques, which can thus handle high-level logic properties on model's trajectories, such as safety, reachability, etc. - this contrasts with the mentioned classical reduction techniques in the frequency domain, which are by and large exclusively developed for the analysis of input-output behaviour and stability.

In this survey paper, we discuss recent approaches grounded on both discretization-based and -free techniques. We should mention that the main focus of this survey is on *discrete-time, continuous-space* stochastic hybrid

systems, whereas we dedicate only one section (Sec. 9) to the otherwise interesting framework of *continuous-time, continuous-space* models, where we overview the corresponding major theoretical results. We should also stress that much of the presented work builds on the extensive theoretical and algorithmic background of *finite-space* Markov models, which is not overviewed here in view of length limitations: we refer the interested reader to (Baier & Katoen 2008, Kwiatkowska et al. 2011) for informative overviews.

1.3. **Different Types of Closeness Guarantees.** Earlier, we have emphasized the importance of providing "formal" abstractions: in this survey we discuss four different types of closeness guarantees (or error bounds) between original SHS and their finite abstractions, as introduced next. These guarantees allow to perform computations over the abstract models, and to formally refine them over the concrete SHS.

**Definition 1.2.** *Let $\Sigma$ be a concrete SHS and $\widehat{\Sigma}$ be its abstraction. For a given specification, the probabilistic closeness between $\Sigma$ and $\widehat{\Sigma}$ is defined according to one of the following:*

- *(i) the difference between probabilities of satisfaction of specifications over the original system $\Sigma$ and its corresponding abstraction $\widehat{\Sigma}$ (cf. equation* (3.3) *or* (3.4)*);*
- *(ii) the probability of the difference between the output trajectories of $\Sigma$ and $\widehat{\Sigma}$ being less than a given threshold (cf. equation* (3.6)*);*
- *(iii) the expectation (moment) of the difference between output trajectories of original system $\Sigma$ and those of its abstraction $\widehat{\Sigma}$ (cf. equation* (9.3)*);*
- *(iv) the probability of satisfaction of logic properties over the abstract system $\widehat{\Sigma}$ is either lower- or upper-bounds the satisfaction probability over original system $\Sigma$ (cf. equations* (3.7) *and* (3.8)*).*

It is worth mentioning that the proposed probabilistic closeness bounds in Definition 1.2 can be employed for abstractions that can be either specification-guided (i, iv) or specification-free (ii, iii). In general, abstractions that are specification-dependent are potentially less conservative as they are a-priori tailored to some given specifications. In comparison, specification-free abstractions are more general since, their corresponding closeness guarantees hold for classes of properties of interest, however this comes at the cost of an increase in their computational complexity. We shall further discuss closeness guarantees corresponding to these two types of abstractions in Sections 3 and 4.

1.4. **Contributions and Organization of this Survey.** This paper provides the first survey of literature on automated formal verification and synthesis of stochastic hybrid systems (SHS). While trying to be comprehensive, we focus on the most recent and sharpest results in the literature, and discuss related approaches in various sections in coarser detail. Besides the selection of the most relevant articles, this survey is intended to help researchers to gain an overall understanding of the many challenges and solution strategies related to the formal verification and the control synthesis of SHS, as well as the associated software tools that have been developed to support the theory. We discuss approaches in relevant literature via both discretization-based and -free techniques, categorizing them over four different closeness guarantees between the concrete SHS and their abstractions, according to Definition 1.2. We employ a running example and discuss approaches

under the lens of *(i) time complexity*, and *(ii) memory requirements*. We also discuss many open problems throughout this survey paper.

We remark that although the survey paper by Teel et al. (2014) also covers *stochastic* hybrid systems, its main focus is on stability analysis: different notions of stability are overviewed, including Lyapunov, Lagrange, asymptotic stability, and recurrence analysis. In contrast, here we focus on formal verification and synthesis goals, defined around complex properties including those expressed as temporal logic formulae (simple instances are safety and reachability specifications), as well as more general properties expressed via omega-regular languages (Baier & Katoen 2008). In addition, we zoom in on *algorithmic solutions* for verification and synthesis of SHS against temporal properties. An overview of the main developments in the area of stochastic model predictive control (SMPC) is in (Mesbah 2016): these results focus on constrained, optimal control synthesis, however they are not natively aimed at providing the formal guarantees on verification and controller synthesis of complex SHS, which are the core focus of our survey.

This survey paper is structured as follows. In Section 2, we formalize the models under study (syntax and semantics) and present preliminaries and main notations from control theory and computer science, which are widely employed throughout the survey. In Section 3, we present one of the pivotal theorems of the article, elaborating on different types of closeness guarantees between a discrete-time SHS and its abstraction. We discuss in depth the required assumptions, and present tools to compute such guarantees. Corresponding results on stochastic similarity relations to connect the probabilistic behavior of concrete models to that of their abstractions are also presented in the same section. Building on these notions of relations amongst models, work on the construction of *infinite* abstractions for SHS is discussed in Section 4, and corresponding results on the construction of *finite* abstractions are studied in Section 5. We also discuss existing abstraction algorithms, together with the assumptions and details underpinning them.

In Section 6, we first formally present the definition of control barrier certificates, as a discretization-free approach, for the analysis and synthesis of SHS. We then present another main theorem of this survey, which allows to quantify an upper bound on the probability that the given system reaches an unsafe region over both finite and infinite time horizons. We also briefly survey results on optimization-based methods for the analysis of SHS, as alternative discretization-free approaches that are, however, not core to this survey. Temporal logic verification and synthesis of SHS are studied in Section 7. Section 8 is devoted to compositional techniques as a potential direction for mitigating the curse of dimensionality. We present the definition of subsystem, together with the formal definition of interconnected systems. We then discuss the main compositionality results, based on two different techniques from literature.

Results for *continuous-time* SHS are briefly presented in Section 9. Section 10 is dedicated to surveying sample- and simulation-based analysis of SHS. Software tools on verification and synthesis of SHS are discussed in Section 11. In Section 12, we summarize the existing analysis methods and highlight relevant directions for future research. In particular, we discuss a few open problems including "formal analysis of SHS via learning and data-driven approaches", "formal synthesis of partially-observed SHS", "secured-by-construction controller synthesis for SHS", "(mixed)-monotonicity of SHS", "compositional construction of interval Markov processes", "compositional controller synthesis for SHS", and "potential extensions of software tools".

## 2. Notations, Preliminaries, and Models

The sets of non-negative and positive integers are denoted by $\mathbb{N} := \{0, 1, 2, \ldots\}$ and $\mathbb{N}_{\geq 1} := \{1, 2, 3, \ldots\}$, respectively. Moreover, the symbols $\mathbb{R}$, $\mathbb{R}_{>0}$, and $\mathbb{R}_{\geq 0}$ denote, respectively, the sets of real, positive and nonnegative real numbers. For any set $X$ we denote by $2^X$ the power set of $X$, namely the set of all subsets of $X$. Given $N$ vectors $x_i \in \mathbb{R}^{n_i}$, $n_i \in \mathbb{N}_{\geq 1}$, and $i \in \{1, \ldots, N\}$, we use $x = [x_1; \ldots; x_N]$ to denote the corresponding column vector of dimension $\sum_i n_i$. We denote by $\|\cdot\|$ and $\|\cdot\|_2$ the infinity and Euclidean norms, respectively. Given any $a \in \mathbb{R}$, $|a|$ denotes the absolute value of $a$. Symbols $\mathbb{I}_n$, $\mathbf{0}_n$, and $\mathbb{1}_n$ denote the identity matrix in $\mathbb{R}^{n\times n}$ and the column vector in $\mathbb{R}^{n\times 1}$ with all elements equal to zero and one, respectively. Given a matrix $P = \{p_{ij}\} \in \mathbb{R}^{n\times n}$, we denote the trace of $P$ by $\mathsf{Tr}(P)$, where $\mathsf{Tr}(P) = \sum_{i=1}^{n} p_{ii}$. We denote the disjunction ($\vee$) and conjunction ($\wedge$) of Boolean functions $f : \Gamma \to \{0, 1\}$ over a (possibly infinite) index set $\Gamma$ by $\bigvee_{\alpha\in\Gamma} f(\alpha)$ and $\bigwedge_{\alpha\in\Gamma} f(\alpha)$, respectively. Given functions $f_i : X_i \to Y_i$, for any $i \in \{1, \ldots, N\}$, their Cartesian product $\prod_{i=1}^{N} f_i : \prod_{i=1}^{N} X_i \to \prod_{i=1}^{N} Y_i$ is defined as $(\prod_{i=1}^{N} f_i)(x_1, \ldots, x_N) = [f_1(x_1); \ldots; f_N(x_N)]$. Given sets $X$ and $Y$, a relation $\mathscr{R} \subseteq X \times Y$ is a subset of the Cartesian product $X \times Y$ that relates $x \in X$ with $y \in Y$ if $(x, y) \in \mathscr{R}$, which is equivalently denoted by $x\mathscr{R}y$. A function $\gamma : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, is said to be a class $\mathcal{K}$ function if it is continuous, strictly increasing, and $\gamma(0) = 0$. A class $\mathcal{K}$ function $\gamma(\cdot)$ is said to be a class $\mathcal{K}_\infty$ if $\gamma(r) \to \infty$ as $r \to \infty$. A continuous function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is said to belong to class $\mathcal{KL}$ if, for each fixed $t$, the map $\beta(r, t)$ belongs to class $\mathcal{K}$ with respect to $r$, and for each fixed nonzero $r$, the map $\beta(r, t)$ is decreasing with respect to $t$, and $\beta(r, t) \to 0$ as $t \to \infty$.

We consider a probability space $(\Omega, \mathcal{F}_\Omega, \mathbb{P}_\Omega)$, where $\Omega$ is the sample space, $\mathcal{F}_\Omega$ is a sigma-algebra on $\Omega$ comprising subsets of $\Omega$ as events, and $\mathbb{P}_\Omega$ is a probability measure that assigns probabilities to events. We assume that the random variables introduced and discussed in this article are measurable functions of the form $X : (\Omega, \mathcal{F}_\Omega) \to (S_X, \mathcal{F}_X)$ (relevant literature contains details supporting these claims). As such, any random variable $X$ induces a probability measure on its space $(S_X, \mathcal{F}_X)$ as $Prob\{A\} = \mathbb{P}_\Omega\{X^{-1}(A)\}$ for any $A \in \mathcal{F}_X$. We often directly discuss the probability measure on $(S_X, \mathcal{F}_X)$ without explicitly mentioning the underlying probability space and the function $X$ itself.

A topological space $\mathcal{S}$ is called a Borel space if it is homeomorphic to a Borel subset of a Polish space (*i.e.,* a separable and completely metrizable topological space). Examples of a Borel space are Euclidean spaces $\mathbb{R}^n$, its Borel subsets endowed with a subspace topology as well as hybrid state spaces (Abate et al. 2008). Any Borel space $\mathcal{S}$ is assumed to be endowed with a Borel sigma-algebra, which is denoted by $\mathcal{B}(\mathcal{S})$. We say that a map $f : \mathcal{S} \to Y$ is measurable whenever it is Borel measurable.

### 2.1. Discrete-Time Stochastic Hybrid Systems.

In this survey, we consider stochastic hybrid systems models in discrete time (dt-SHS), first introduced by Amin et al. (2006), Abate et al. (2008), and defined formally as follows.

**Definition 2.1.** *A discrete-time stochastic hybrid system (dt-SHS) is characterized by the tuple*

$$\Sigma = (\mathcal{Q}, n, X, U, T_{\mathsf{x}}, Y, h), \text{ where} \tag{2.1}$$

- $\mathcal{Q} := \{q_1, \ldots, q_p\}$ *for some* $p \in \mathbb{N}_{\geq 1}$*, represents the discrete-state space;*

- $n : \mathcal{Q} \to \mathbb{N}_{\geq 1}$ *assigns to each discrete state value* $q \in \mathcal{Q}$ *the dimension of the continuous-state space* $\mathbb{R}^{n(q)}$*;*
- $X \subseteq \cup_{q\in\mathcal{Q}}\{q\}\times\mathbb{R}^{n(q)}$ *is a Borel space as the hybrid-state space of the system. We denote by* $(X, \mathcal{B}(X))$ *the measurable space, with* $\mathcal{B}(X)$ *being the Borel sigma-algebra over the state space;*
- $U \subseteq \mathbb{R}^m$ *is a Borel space as the input space of the system;*
- $T_{\mathsf{x}} : \mathcal{B}(X) \times X \times U \to [0, 1]$ *is a conditional stochastic kernel that assigns to any* $x \in X$*, and* $\nu \in U$*, a probability measure* $T_{\mathsf{x}}(\cdot|x, \nu)$ *on the measurable space* $(X, \mathcal{B}(X))$*. This stochastic kernel specifies probabilities over executions* $\{x(k), k \in \mathbb{N}\}$ *of the hybrid system, such that for any set* $\mathcal{A} \in \mathcal{B}(X)$ *and any* $k \in \mathbb{N}$*,*

$$\mathbb{P}(x(k+1) \in \mathcal{A}\big|x(k), \nu(k)) = \int_{\mathcal{A}} T_{\mathsf{x}}(\mathsf{d}x(k+1)\big|x(k), \nu(k));$$

- $Y \subseteq \mathbb{R}^q$ *is a Borel space as the output space of the system;*
- $h : X \to Y$ *is a measurable function as the output map that takes a state* $x \in X$ *to its output* $y = h(x)$*.*

An example of dt-SHS is discussed in the running example and equation (2.4).

This definition is general and describes numerous applications. As this general structure of the state space can be notationally heavy, for the scope of this survey and for the sake of an easier presentation, we will introduce definitions, algorithms, and theorems based on a specific class of SHS with a single discrete mode (*i.e.,* $X \subseteq \mathbb{R}^n$), called discrete-time stochastic control systems (dt-SCS) (Meyn & Tweedie 1993, Hernández-Lerma & Lasserre 1996). We emphasize that broadly the notions and approaches underlying the proposed results can be generalized to SHS endowed with *hybrid* state spaces. A schematic representation of dt-SCS $\Sigma$ is shown in Fig. 3.

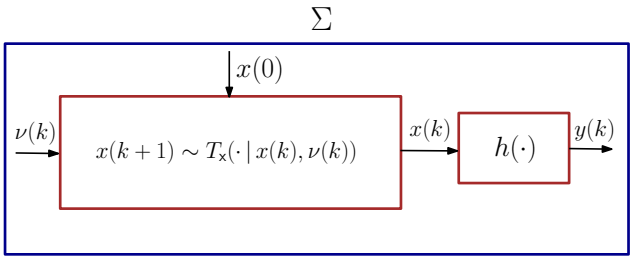


FIGURE 3. A schematic representation of a dt-SCS $\Sigma$.

As argued by Kallenberg (1997), any dt-SCS endowed with a stochastic transition kernel $T_{\mathsf{x}}$ as in Definition 2.1 can be *equivalently* represented by a dt-SCS with pair $(f, \varsigma)$, as formalized next. Note that this alternative representation is more common in control theory. It is often easier to show specific results of this paper based on the alternative representation.

**Definition 2.2.** *A discrete-time stochastic control system (dt-SCS) is represented by the tuple*

$$\Sigma = (X, U, \varsigma, f, Y, h), \text{ where} \tag{2.2}$$

- $X \subseteq \mathbb{R}^n$ *is a Borel space as the state space of the system;*

- $U \subseteq \mathbb{R}^m$ *is a Borel space as the input space of the system;*
- $\varsigma$ *is a sequence of independent and identically distributed (i.i.d.) random variables from a sample space* $\Omega$ *to the measurable space* $(\mathcal{V}_\varsigma, \mathcal{F}_\varsigma)$

$$\varsigma := \{\varsigma(k) : (\Omega, \mathcal{F}_\Omega) \to (\mathcal{V}_\varsigma, \mathcal{F}_\varsigma),\ k \in \mathbb{N}\};$$

- $f : X \times U \times \mathcal{V}_\varsigma \to X$ *is a measurable function characterizing the state evolution of the system;*
- $Y \subseteq \mathbb{R}^q$ *is a Borel space as the output space of the system;*
- $h : X \to Y$ *is a measurable function as the output map.*

*For given initial state* $x(0) \in X$ *and input sequence* $\nu(\cdot) : \mathbb{N} \to U$*, the evolution of the state of the dt-SCS* $\Sigma$ *can be written, with* $k \in \mathbb{N}$*, as*

$$\Sigma : \begin{cases} x(k+1) = f(x(k), \nu(k), \varsigma(k)), \\ y(k) = h(x(k)). \end{cases} \tag{2.3}$$

*We denote by* $\mathbb{U}$ *the collection of input sequences* $\{\nu(k) : \Omega \to U,\ k \in \mathbb{N}\}$*, in which* $\nu(k)$ *is independent of* $\varsigma(z)$ *for any* $k, z \in \mathbb{N}$ *and* $z \geq k$*. For any initial state* $a \in X$*, and input* $\nu(\cdot) \in \mathbb{U}$*, the random sequences* $x_{a\nu} : \Omega \times \mathbb{N} \to X$*, and* $y_{a\nu} : \Omega \times \mathbb{N} \to Y$ *that satisfy* (2.3) *are respectively called the* solution process *and the* output trajectory *of* $\Sigma$ *under an input* $\nu$ *and an initial state* $a$*. System* $\Sigma$ *is said to be finite if* $X$ *and* $U$ *are finite sets, and infinite otherwise.*

**Running Example.** To help the reader gain a better understanding of the details in this survey paper, we present a simple yet interesting running example. We tailor the models and apply the results presented in this survey to a temperature regulation problem for a room equipped with a heater. The model of this case study is borrowed from (Fehnker & Ivančić 2004, Meyer et al. 2017), but modified by including an additive noise, which is intended to capture the effect of uncertain weather- or user-dependent factors. The evolution of the temperature $T(\cdot)$ over time can be described by the following dt-SCS:

$$\Sigma : \begin{cases} T(k+1) = a(k)T(k) + \gamma T_h \nu(k) + \theta T_e + R\varsigma(k), \\ y(k) = T(k), \end{cases} \tag{2.4}$$

where the signal $a(k) := (1 - \theta - \gamma\nu(k))$ depends on the input $\nu(k)$, $R = 0.6$ is the noise coefficient, and $\theta = 0.4$, and $\gamma = 0.5$ are factors that affect the rate of heat conduction between the external environment and the room, and between the heater and the room. The parameter $T_e = -1\,^\circ C$ is the outside temperature, $T_h = 50\,^\circ C$ is the heater temperature, and $y$ is the (sensed, observed) output of the system, which in this instance corresponds to the temperature itself. Finally, $\varsigma$ is assumed to be i.i.d. with a normal distribution having zero mean and a covariance equal to 1.

The model in (2.4) can be alternatively (and equivalently) characterized via the tuple in (2.2): here $X, U$ are subsets of the real numbers, $f(x(k), \nu(k), \varsigma(k)) = a(k)T(k) + \gamma T_h \nu(k) + \theta T_e + R\varsigma(k)$, and the output map $h$ is identity (accordingly, the output space $Y = X$). Note that this system is a very special instance of SHS in (2.1) endowed with a single discrete mode, where the conditional stochastic kernel $T_{\mathsf{x}}$ is a normal distribution with mean $a(k)T(k) + \gamma T_h \nu(k) + \theta T_e$ and covariance $R^2$. Alternatively, if the input $\nu(k)$ is assumed to be a finite-valued function of the state, e.g. a binary function switching upon hitting the boundaries of a temperature interval, then the model can be interpreted as a two-mode SHS.

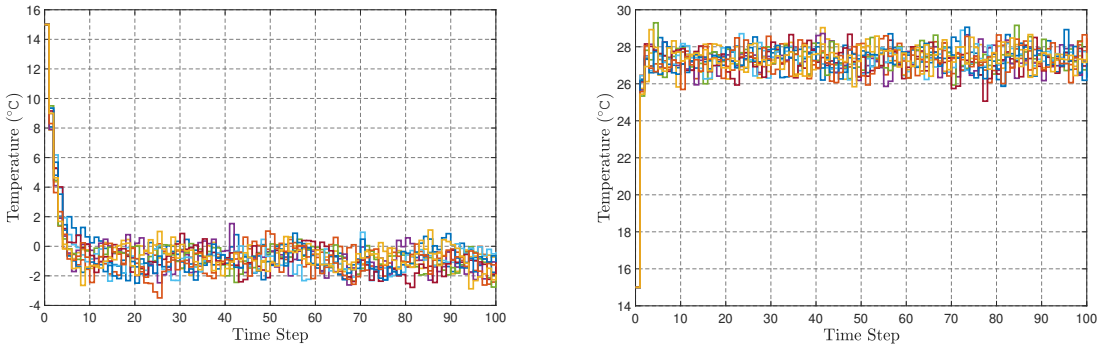


FIGURE 4. State trajectories, generated for the running example with 10 different noise realizations within the finite time horizon $T_d = 100$ from an initial condition $x_0 = 15$, and with $\nu = 0$ (top, heating off) and $\nu = 1$ (bottom, heating on), respectively.

In order to provide some intuitions on the evolution of temperature, we plot in Fig. 4 the state trajectories of the running example with 10 different noise realizations within the finite time horizon $T_d = 100$ from an initial condition $x_0 = 15$ and with inputs $\nu = 0$ and $\nu = 1$. □

Given the dt-SCS model in (2.2), we introduce *Markov policies* as follows.

**Definition 2.3.** *A Markov policy for the dt-SCS $\Sigma$ in (2.2) is a sequence $\mu = (\mu_0, \mu_1, \mu_2, \ldots)$ of universally measurable stochastic kernels $\mu_n$ (Bertsekas & Shreve 1996), each defined on the input space $U$ given $X$ and such that for all $x(n) \in X$, $\mu_n(U(x(n)) \,|\, x(n)) = 1$. The class of all Markov policies is denoted by $\mathcal{M}_p$.*

Informally, Markov policies are history-independent and the control input taken at the current time instance is selected, possibly randomly, with a distribution that depends only on the current state.

2.2. **Relations between Models.** We now define the notion of incremental input-to-state stability for $\Sigma$, as a pivotal assumption that will allow some of the results, in particular to provide closeness guarantees between output trajectories of concrete system $\Sigma$ and its abstraction $\widehat{\Sigma}$, as per *(ii)* and *(iii)* in Definition 1.2.

**Definition 2.4.** *A dt-SCS $\Sigma = (X, U, \varsigma, f, Y, h)$ is called* incrementally input-to-state stable ($\delta$-ISS) *if there exists a function $S : X \times X \to \mathbb{R}_{\geq 0}$ such that $\forall x, x' \in X$, $\forall \nu, \nu' \in U$, the following two inequalities hold:*

$$\underline{\alpha}(\|x - x'\|) \leq S(x, x') \leq \overline{\alpha}(\|x - x'\|), \tag{2.5}$$

*and*

$$\mathbb{E}\Big[S(f(x, \nu, \varsigma), f(x', \nu', \varsigma)) \,\Big|\, x, x', \nu, \nu'\Big] - S(x, x') \leq -\bar{\kappa}(S(x, x')) + \rho(\|\nu - \nu'\|), \tag{2.6}$$

*for some $\underline{\alpha}, \overline{\alpha}, \bar{\kappa} \in \mathcal{K}_\infty$, and $\rho \in \mathcal{K}_\infty \cup \{0\}$.*

Later we will show how one can use the $\delta$-ISS property to bound the distance between two solution processes starting from different initial conditions and under different input trajectories.

We now define the notion of stochastic simulation functions (SSF) between $\widehat{\Sigma}$ and $\Sigma$, which allows to provide closeness guarantees between the output trajectories of the two models, as per *(ii)* in Definition 1.2.

**Definition 2.5.** *Consider two dt-SCS* $\Sigma = (X, U, \varsigma, f, Y, h)$ *and* $\widehat{\Sigma} = (\hat{X}, \hat{U}, \varsigma, \hat{f}, \hat{Y}, \hat{h})$. *A function* $V : X \times \hat{X} \to \mathbb{R}_{\geq 0}$ *is called a* stochastic simulation function *(SSF) from* $\widehat{\Sigma}$ *to* $\Sigma$ *if*

- $\exists \alpha \in \mathcal{K}_\infty$ *such that*

$$\forall x \in X, \forall \hat{x} \in \hat{X}, \quad \alpha(\|h(x) - \hat{h}(\hat{x})\|) \leq V(x, \hat{x}),$$

- $\forall x \in X, \forall \hat{x} \in \hat{X}, \forall \hat{\nu} \in \hat{U}$, $\exists \nu \in U$ *such that*

$$\mathbb{E}\Big[V(f(x, \nu, \varsigma), \hat{f}(\hat{x}, \hat{\nu}, \varsigma)) \,\big|\, x, \hat{x}, \nu, \hat{\nu}\Big] \leq \kappa V(x, \hat{x}) + \rho_{\text{ext}}(\|\hat{\nu}\|) + \psi, \tag{2.7}$$

  *for some* $0 < \kappa < 1$, $\rho_{\text{ext}} \in \mathcal{K}_\infty \cup \{0\}$, *and* $\psi \in \mathbb{R}_{\geq 0}$.

*We denote by* $\widehat{\Sigma} \preceq \Sigma$ *if there exists an SSF* $V$ *from* $\widehat{\Sigma}$ *to* $\Sigma$, *and call the system* $\widehat{\Sigma}$ *an abstraction of the concrete (original) system* $\Sigma$. *Note that* $\widehat{\Sigma}$ *may be finite or infinite, depending on the cardinality of the sets* $\hat{X}$ *and* $\hat{U}$.

Informally, stochastic simulation functions are Lyapunov-like functions defined over the Cartesian product of the state spaces of two models, which relate their output trajectories and indeed guarantee that their mismatch (namely the difference between their outputs) remains within some guaranteed error bounds. This mismatch can be conceived as the abstraction error if one model is obtained as the simplification of a given concrete model. In particular, since SHS are in general complex and intractable, stochastic simulation functions are beneficial to connect the probabilistic behavior of concrete SHS to that of their abstractions: in particular, by providing closeness guarantees between output trajectories of two systems via the established stochastic simulation functions, one can perform formal analysis over the simplified abstractions and transfer the obtained results back over the original SHS.

**Remark 2.6.** *The second condition in Definition 2.5 implies the existence of a function* $\nu = \nu_{\hat{\nu}}(x, \hat{x}, \hat{\nu})$ *for the satisfaction of* (2.7). *This function is called the "interface function" and will be used to refine a synthesized policy* $\hat{\nu}$ *for* $\widehat{\Sigma}$ *to a policy* $\nu$ *for* $\Sigma$ *(cf. Fig. 1), and will be discussed later in Sections 4 and 5.*

2.3. **Temporal Logic Specifications.** Formal specifications provide a rigorous and unambiguous formalism to express formal requirements over models. A common way to describe such formal requirements is utilizing specifications expressed as automata or in a temporal logic, *e.g.,* formulae expressed in linear temporal logic (LTL) (Pnueli 1977). Let us start with some basic properties. Consider the dt-SCS in (2.2) and measurable sets $\mathsf{A}, \mathsf{B} \subset Y$, named respectively "safe" and "target" set. (Later, in Def. 2.7, we shall encompass these sets through a labelling function $\mathsf{L}$.) We define the bounded-horizon safety property as $\Box^{\leq T_d}\mathsf{A}$, indicating that all output trajectories $\{y(k)\}_{k\geq 0}$ start from the safe set $\mathsf{A}$ and remain inside it over the finite-time horizon $k \in [0, T_d]$. Similarly, we say that an output trajectory $\{y(k)\}_{k\geq 0}$ reaches a target set $\mathsf{B}$ within the discrete time interval $[0, T_d] \subset \mathbb{N}$, if there exists a $k \in [0, T_d]$ such that $y(k) \in \mathsf{B}$: this bounded-horizon reachability property is denoted by $\Diamond^{\leq T_d}\{y \in \mathsf{B}\}$ or briefly $\Diamond^{\leq T_d}\mathsf{B}$. Extending the above requirements to infinite horizons by $T_d \to \infty$, we denote the corresponding safety and reachability properties as $\Box\mathsf{A}$ and $\Diamond\mathsf{B}$, which are colloquially

said "always $\mathsf{A}$" and "eventually $\mathsf{B}$", respectively. Additionally, we define *reach-avoid* specifications by the formula $\mathsf{A}\,\mathsf{U}\,\mathsf{B}$, requiring the output trajectories to reach the target set $\mathsf{B}$ while remaining in the safe set $\mathsf{A}$ - this property is also known as *constrained reachability*. More generally, all the described basic properties can be reframed as specifications in LTL.

We formally define syntax and semantics of linear temporal logic (LTL) as follows.

**Definition 2.7.** *Consider a set of atomic propositions $AP$ and the alphabet $\Sigma_{\mathsf{a}} := 2^{AP}$. Let $\omega = \omega(0), \omega(1), \omega(2), \ldots$ be an infinite word, that is, a string composed of letters from $\Sigma_{\mathsf{a}}$ (i.e. $\omega(i) \in \Sigma_{\mathsf{a}}, \forall i \in \mathbb{N}$). We are interested in those atomic propositions that are relevant to the dt-SCS via a measurable labeling function $\mathsf{L}$ from the (continuous) output space to the (finite) alphabet as $\mathsf{L} : Y \to \Sigma_{\mathsf{a}}$. Informally, the output space is "tagged with labels" that are relevant to the specifications of interest. Accordingly, output trajectories $\{y(k)\}_{k\geq 0} \in Y^{\mathbb{N}}$ of the dt-SCS can be readily mapped to the set of infinite words $\Sigma_{\mathsf{a}}^{\mathbb{N}}$, as*

$$\omega = \mathsf{L}(\{y(k)\}_{k\geq 0}) := \{\omega \in \Sigma_{\mathsf{a}}^{\mathbb{N}} \,|\, \omega(k) = \mathsf{L}(y(k))\}.$$

*We define the LTL syntax (Baier & Katoen 2008) as*

$$\varphi ::= \text{true} \,|\, p \,|\, \neg\varphi \,|\, \varphi_1 \wedge \varphi_2 \,|\, \bigcirc\varphi \,|\, \varphi_1 \,\mathsf{U}\, \varphi_2.$$

*Given a trace $\omega = \omega(0), \omega(1), \omega(2), \ldots$, let us denote the suffix of $\omega$ starting from $\omega(i)$ by*

$$(\omega, i) = \omega(i), \omega(i+1), \omega(i+2), \ldots.$$

*We denote by $(\omega, i) \models \varphi$ when the LTL formula $\varphi$ is true on the suffix $(\omega, i)$. This satisfaction is defined inductively as follows:*

- $(\omega, i) \models$ *true;*
- $(\omega, i) \models p$, *for $p \in AP$ iff $p \in \omega(i)$;*
- $(\omega, i) \models \neg\varphi$ *iff* $(\omega, i) \not\models \varphi$;
- $(\omega, i) \models \varphi_1 \wedge \varphi_2$ *iff* $(\omega, i) \models \varphi_1$ *and* $(\omega, i) \models \varphi_2$;
- $(\omega, i) \models \bigcirc\varphi$ *iff* $(\omega, i+1) \models \varphi$;
- $(\omega, i) \models \varphi_1 \,\mathsf{U}\, \varphi_2$ *iff for some $j$ such that $i \leq j$, we have $(\omega, j) \models \varphi_2$, and for all $k$ s.t. $i \leq k < j$, we have $(\omega, k) \models \varphi_1$.*

*Formula $\varphi$ is true on $\omega$, denoted by $\omega \models \varphi$, if and only if $(\omega, 0) \models \varphi$.*

Based on the above operators, we can also introduce other formulae, obtained via propositional or temporal manipulations. These can encode simple properties, such as the mentioned reachability and safety specifications, or more complicated requirements, obtained by composing arbitrary numbers of operators. For instance, $\varphi_1 \vee \varphi_2$, $\Diamond\varphi$, and $\Box\varphi$ have semantics

- $(\omega, i) \models \varphi_1 \vee \varphi_2$ *iff* $(\omega, i) \models \varphi_1$ *or* $(\omega, i) \models \varphi_2$;
- $(\omega, i) \models \Diamond\varphi$ *iff for some $j$ such that $i \leq j$, we have $(\omega, j) \models \varphi$;*
- $(\omega, i) \models \Box\varphi$ *iff for all $j$ such that $i \leq j$, we have $(\omega, j) \models \varphi$.*

Later, we shall be interested in quantifying the likelihood of verifying given LTL formulae by Markov models, such as MDPs or dt-SCS. Clearly, this requires reasoning about measurability of events associated to the LTL specifications introduced above; however, we shall not delve into measure-theoretical issues in the present survey.

We now introduce a fragment of LTL properties known as *syntactically co-safe* linear temporal logic (scLTL) (Kupferman & Vardi 2001); scLTL properties are popular since their satisfaction can be sufficiently witnessed by finite-length traces.

**Definition 2.8.** *An scLTL over a set of atomic propositions $AP$ is a fragment of LTL such that the negation operator ($\neg$) only occurs before atomic propositions, and it is characterized by the following grammar:*

$$\varphi ::= \text{true} \mid p \mid \neg p \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \bigcirc \varphi \mid \varphi_1 \,\mathsf{U}\, \varphi_2 \mid \Diamond \varphi,$$

*with $p \in AP$. The semantics of satisfaction follows from that of LTL.*

Another enticing aspect about scLTL is that it can be alternatively expressed by means of simple finite-state automata (Kupferman & Vardi 2001, Belta et al. 2017). This means that the set of words satisfying a given scLTL formula can be equivalently expressed as the set of words that are accepted by a proper (not necessarily unique) finite-state automaton. More specifically, we introduce a class of models known as deterministic finite-state automata (DFA).

**Definition 2.9.** *A DFA is a tuple $\mathcal{A} = (Q_\ell, q_0, \Sigma_a, F_a, \mathsf{t})$, where $Q_\ell$ is a finite set of locations (states), $q_0 \in Q_\ell$ is the initial location, $\Sigma_a$ is a finite set (a.k.a., alphabet), $F_a \subseteq Q_\ell$ is a set of accepting locations, and $\mathsf{t} : Q_\ell \times \Sigma_a \to Q_\ell$ is a transition function.*

Consider a set of atomic propositions $AP$ and the alphabet $\Sigma_{\mathsf{a}} := 2^{AP}$. A finite word composed of letters of the alphabet, *i.e.,* $\omega_f = (\omega_f(0), \ldots, \omega_f(n)) \in \Sigma_{\mathsf{a}}^{n+1}$, is accepted by a DFA $\mathcal{A}$ if there exists a finite run $q = (q(0), \ldots, q(n+1)) \in Q_\ell^{n+2}$ such that $q(0) = q_0$, $q(i+1) = \mathsf{t}(q(i), \omega_f(i))$ for all $0 \leq i \leq n$, and $q(n+1) \in F_{\mathsf{a}}$. The accepted language of $\mathcal{A}$, denoted $\mathcal{L}(\mathcal{A})$, is the set of all words accepted by $\mathcal{A}$. For every scLTL property $\varphi$, cf. Definition 2.8, there exists a DFA $\mathcal{A}_\varphi$ such that

$$\mathcal{L}_f(\varphi) = \mathcal{L}(\mathcal{A}_\varphi),$$

where $\mathcal{L}_f$ denotes the set of all words associated to an scLTL formula $\varphi$. In some parts of this article, we focus on the computation of probability of $\omega_f \in \mathcal{L}(\mathcal{A}_\varphi)$ over bounded intervals. In other words, we fix a time horizon $T_d$ and compute $\mathbb{P}(\omega_f(0)\omega_f(1)\ldots\omega_f(T_d) \in \mathcal{L}(\mathcal{A}_\varphi)$ s.t. $|\omega_f| \leq T_d + 1)$, with $|\omega_f|$ denoting the length of $\omega_f$.

The following example, borrowed from (Lavaei et al. 2019), provides an automaton associated with a reach-avoid specification

**Example 2.10.** *Consider two measurable sets $\mathsf{A}, \mathsf{B} \subset Y$ as the safe and target sets, respectively. We present the DFA for the specification ($\mathsf{A}\ \mathsf{U}\ \mathsf{B}$) which requires the output trajectories to reach the target set $\mathsf{B}$ while*

*remaining in the safe set* $\mathsf{A}$. *Note that we do not assume these two sets are disjoint. Consider the set of atomic propositions* $AP = \{\mathsf{A}, \mathsf{B}\}$ *and the alphabet* $\Sigma_{\mathsf{a}} = \{\emptyset, \{\mathsf{A}\}, \{\mathsf{B}\}, \{\mathsf{A}, \mathsf{B}\}\}$. *Define the labeling function as*

$$\mathsf{L}(y) = \begin{cases} \{\mathsf{A}\} =: a & \text{if } \ y \in \mathsf{A}\backslash\mathsf{B}, \\ \{\mathsf{B}\} =: b & \text{if } \ y \in \mathsf{B}, \\ \emptyset =: c & \text{if } \ y \notin \mathsf{A} \cup \mathsf{B}. \end{cases}$$

*As can be seen from the above definition of the labeling function* $\mathsf{L}$, *it induces a partition over the output space* $Y$ *as*

$$\mathsf{L}^{-1}(a) = \mathsf{A}\backslash\mathsf{B}, \quad \mathsf{L}^{-1}(b) = \mathsf{B}, \quad \mathsf{L}^{-1}(c) = Y\backslash(\mathsf{A} \cup \mathsf{B}).$$

*Note that we have indicated the elements of* $\Sigma_{\mathsf{a}}$ *with lower-case letters, for ease of notation. The specification* $(\mathsf{A}\ \mathsf{U}\ \mathsf{B})$ *can be equivalently written as* $(a\ \mathsf{U}\ b)$ *with the associated DFA depicted in Figure 5. This DFA has the set of locations* $Q_\ell = \{q_0, q_1, q_2, q_3\}$, *an initial location* $q_0$, *and an accepting location* $F_{\mathsf{a}} = \{q_2\}$. *Thus output trajectories of a dt-SCS* $\Sigma$ *satisfy the specification* $(a\ \mathsf{U}\ b)$ *if and only if their associated words are accepted by this DFA.*

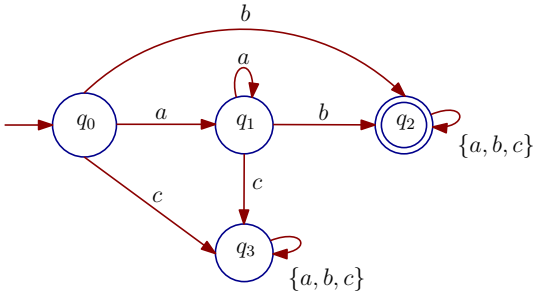


FIGURE 5. DFA $\mathcal{A}_\varphi$ of the reach-avoid specification $(a\ \mathsf{U}\ b)$.

Generalizing beyond scLTL (and corresponding DFAs), often we are interested in infinite paths through a system and thus in infinite words. $\omega$-regular languages generalize the definition of regular languages to encompass sets of infinite-length words. Correspondingly, $\omega$-regular properties are specifications expressed via $\omega$-regular languages and $\omega$-automata are finite-state models that can accept them. Two of the most commonly used automata to express $\omega$-regular properties are *Büchi* and *Rabin* automata (Baier & Katoen 2008). Indeed, it can be shown that non-deterministic[2] *Büchi* automata (NBA) encompass the entire LTL.

**Definition 2.11.** *An NBA is a tuple* $\mathcal{A} = (Q_\ell, q_0, \Sigma_{\mathsf{a}}, F_{\mathsf{a}}, \mathsf{t})$, *where* $Q_\ell$ *is a finite set of locations,* $q_0 \subseteq Q_\ell$ *is the initial location,* $\Sigma_{\mathsf{a}}$ *is the finite alphabet,* $F_{\mathsf{a}} \subseteq Q_\ell$ *is a set of accept locations, and* $\mathsf{t} : Q_\ell \times \Sigma_{\mathsf{a}} \to 2^{Q_\ell}$ *is a transition relation.*

[2]In a deterministic automaton, each transition is *uniquely* determined by its source state and input symbol, whereas a non-deterministic automaton does not abide by this requirement. Non-determinism in Büchi automata is required to encompass general LTL properties.

An infinite word composed of letters of the alphabet, *i.e.,* $\omega = (\omega(0), \omega(1) \ldots) \in \Sigma_{\mathsf{a}}^{\omega}$, is accepted by the NBA $\mathcal{A}$ if there exists an infinite run $q = (q(0), q(1), \ldots) \in Q_{\ell}^{\omega}$ such that $q(0) \in q_0$, $q(i+1) \in \mathsf{t}(q(i), \omega(i))$ for all $0 \leq i$, and $q(i) \in F_{\mathsf{a}}$ infinitely often. Let us remark that the *non-deterministic* feature of NBA is necessary to express $\omega$-regular properties, and in particular the set of LTL specifications. Alternatively, later in this survey, we shall mention limit-deterministic *Büchi* automata. Similarly, deterministic *Rabin* automata can be utilized, which employ a different (and more involuted) acceptance semantics. However, for the sake of space we avoid to detail them and instead refer the readers to (Baier & Katoen 2008) for a comprehensive discussion.

For probabilistic models, properties of interest can be expressible in a different logic, which encompasses a probabilistic operator in its syntax and is such that its satisfaction is defined over states (branching semantics), as opposed to the case of LTL specifications whose satisfaction is defined over trajectories (linear semantics). Probabilistic computation tree logic (PCTL) (Ciesinski & Größer 2004) can be introduced as follows.

**Definition 2.12.** *The syntax of (PCTL) formulae is defined recursively using the following operators:*

$$\varphi ::= \text{true} \mid p \mid \varphi \wedge \varphi \mid \neg\varphi \mid \mathbb{P}_{\sim p}[\psi],$$
$$\psi ::= \bigcirc\varphi \mid \varphi \ \mathsf{U} \ \varphi,$$

*where $p \in AP$ and $\sim \in \{<, \leq, \geq, >\}$, and $p \in [0,1]$. The semantics for $\bigcirc$ and for $\mathsf{U}$ are as before, and the semantics for* true*, $p$, $\wedge$, and $\neg$ are also identical except for being defined with reference to a state $s = \omega(0)$, instead of the first element of a path $\omega$. The satisfaction semantics for the expression $s \models \mathbb{P}_{\sim p}[\psi]$ is defined as follows:*

$$\Pr(\{\omega \in \Sigma_{\mathsf{a}}^{\omega} \mid \omega(0) = s \text{ and } \omega, 0 \models \psi\}) \sim p,$$

*where* Pr *is the probability distribution over the infinite paths through $Y$ induced by the stochastic dynamics.*

Discussion of characterization and computation of basic PCTL specifications, such as safety and reachability, for dt-SCS is introduced by Amin et al. (2006) and generalized by Abate et al. (2008), is applied to invariance in (Pola & Pola 2006), and extended to reach-avoid properties in (Summers & Lygeros 2010). The characterization of properties in these works is based on a dynamic programming recursion, which is generalized by Ramponi et al. (2010), which connects the characterization of PCTL to dynamic programming. Further, (Tkachev et al. 2013) generalizes this work to regular and $\omega$-regular properties, respectively, leveraging DFA and deterministic NBA models from above, and a product construction further discussed in Section 7.

**Remark 2.13.** *Note that it has been shown in relevant literature that all random variables discussed in this survey are measurable functions of the form $X : (\Omega, \mathcal{F}_{\Omega}) \to (S_X, \mathcal{F}_X)$. Accordingly, all specifications discussed in this survey encompass measurable events under the system dynamics, so that one can properly assign probabilities to those events (Vardi 1985, Proposition 2.3).*

**Remark 2.14.** *Let us remark that literature on SHS deals with diverse types of logical specifications, including safety, reachability, reach-avoid, syntactically co-safe linear temporal logic (scLTL), bounded linear temporal logic (BLTL) (Maler & Nickovic 2004), signal temporal logic (STL) (Maler et al. 2008), probabilistic computation tree logic (PCTL), and metric interval temporal logic (MITL) (Maler & Pnueli 1995, Maler et al. 2005). However, for the sake of better readability, we mainly focus on simpler, more common requirements*

*that are widely employed in the literature and in the practice, including safety, reachability, reach-avoid, as well as LTL and PCTL properties.*

## 3. Stochastic Similarity Relations for Abstractions

In view of the generality of SHS and of the properties of interest, closed-form solutions for verification and synthesis problems, concerning the expression of value functions or of synthesized feedback policies, are in general not available explicitly, and thus require to be computed numerically. In this instance, an effective approach is to approximate a given SHS model by simpler abstract ones, for example models with lower dimensionality, simpler dynamics, or even a finite state space. In order to render this approximation "formal," it is desirable to provide guarantees on this approximation step, so that the analysis and/or the synthesis on the derived abstract models can be translated back to the original SHS. As anticipated earlier, stochastic similarity relations can indeed be employed to relate the probabilistic behavior of a concrete model (e.g., a given SHS) to that of its abstractions. They can be framed as stochastic simulation and bisimulation relations, in either exact or approximate form.

In the following, we present four theorems that summarize several results from relevant literature, and thus provide the four different types of closeness guarantees, as introduced earlier, between a concrete SHS and a derived abstraction. First, with focus on dt-SCS, we present bounds on the difference between the probability of satisfaction of logic properties over a given system $\Sigma$ and its corresponding finite abstraction $\widehat{\Sigma}$ (Soudjani 2014, Tkachev et al. 2013), as introduced in Sec. 1.1. This type of probabilistic closeness requires a Lipschitz continuity assumption on the stochastic kernel of the dt-SCS, as in the following.

**Definition 3.1.** *The dt-SCS in Definition 2.2 is Lipschitz continuous if the stochastic kernel $T_{\mathsf{x}}$ admits a density function $t_s(\bar{x}|x,u)$ satisfying the following inequality for some constant $\bar{\mathscr{H}} \geq 0$:*

$$|t_s(\bar{x}\,|\,x,u) - t_s(\bar{x}\,|\,x',u')| \leq \bar{\mathscr{H}}(\|x-x'\| + \|u-u'\|), \tag{3.1}$$

*for all $x, x', \bar{x} \in X$ and all $u, u' \in U$. If the policy for dt-SCS is given as $\nu : X \to U$, we define the Lipschitz constant of the stochastic kernel by $\mathscr{H}$, where*

$$|t_s(\bar{x}\,|\,x,\nu(x)) - t_s(\bar{x}\,|\,x',\nu(x'))| \leq \mathscr{H}\|x-x'\|, \tag{3.2}$$

*for all $x, x', \bar{x} \in X$.*

Now, we have all the ingredients to introduce the first approximation theorem, which is related to *(i)* in Definition 1.2. Note that a finite abstraction $\widehat{\Sigma}$ is obtained from the original system $\Sigma$ by first constructing finite partitions of state and input sets, and then selecting arbitrary "representative points" as abstract states and inputs. Transition probabilities in the finite abstraction $\widehat{\Sigma}$ are computed accordingly (cf. Section 5, Algorithm 1).

**Theorem 3.2.** *Let $\Sigma = (X, U, \varsigma, f, Y, h)$ be a continuous-space dt-SCS and $\widehat{\Sigma} = (\hat{X}, \hat{U}, \varsigma, \hat{f}, \hat{Y}, \hat{h})$ be its finite abstraction. Assume that the original system $\Sigma$ is Lipschitz continuous, as per Definition 3.1. For a given logic specification $\varphi$, and for any policy $\hat{\nu}(\cdot) \in \hat{\mathbb{U}}$ that preserves the Markov property for the closed-loop $\widehat{\Sigma}$*

*(i.e., system $\widehat{\Sigma}$ fed by input $\hat{\nu}(\cdot)$, which is denoted by $\widehat{\Sigma}_{\hat{\nu}}$), the probabilistic closeness between two systems is as follows:*

$$|\mathbb{P}(\Sigma_{\hat{\nu}} \vDash \varphi) - \mathbb{P}(\widehat{\Sigma}_{\hat{\nu}} \vDash \varphi)| \leq \lambda_1, \tag{3.3}$$

*with $\lambda_1 := T_d \delta \mathscr{H} \mathscr{L}_b$, where $T_d$ is the finite-time horizon, $\delta$ is the state discretization parameter, $\mathscr{H}$ is the Lipschitz constant of the stochastic kernel $T_{\mathsf{x}}$ under policy $\hat{\nu}$ as in* (3.2)*, and $\mathscr{L}_b$ is the Lebesgue measure of the state space. Moreover, the difference between the optimal probabilities of satisfying a given LTL specification $\varphi$ by the two models is bounded by*

$$\big|\sup_{\nu} \mathbb{P}(\Sigma_{\nu} \vDash \varphi) - \sup_{\hat{\nu}} \mathbb{P}(\widehat{\Sigma}_{\hat{\nu}} \vDash \varphi)\big| \leq \bar{\lambda}_1, \tag{3.4}$$

*with $\bar{\lambda}_1 := T_d \delta \bar{\mathscr{H}} \mathscr{L}_b$, where $\bar{\mathscr{H}}$ is the Lipschitz constant of the stochastic kernel $T_{\mathsf{x}}$ over the state $x$ and input $\nu$ as in* (3.1)*. Furthermore, for the optimal policy $\hat{\nu}^*$ that maximizes the satisfaction probability of $\varphi$ for the abstraction $\widehat{\Sigma}$, we have*

$$\big|\mathbb{P}(\Sigma_{\hat{\nu}^*} \vDash \varphi) - \mathbb{P}(\widehat{\Sigma}_{\hat{\nu}^*} \vDash \varphi)\big| \leq 2\bar{\lambda}_1. \tag{3.5}$$

**Remark 3.3.** *Note that the Lebesgue measure $\mathscr{L}_b$ (informally, the "volume") of the set of interest (within the state space) appears in the error formula, as per* (3.3)-(3.5)*, which makes them meaningful over bounded domain and, possibly, quite conservative. There exist techniques based on an adaptive and sequential gridding scheme (e.g., Soudjani & Abate (2013)) that mitigate both shortcomings.*

**Remark 3.4.** *Let us remark that, in general, the construction of the abstract system $\widehat{\Sigma}$ is performed in a way that allows a proper interpretation of the concrete specification $\varphi$ on the abstract model - as such, the abstraction will be property-dependent. In the closeness bounds* (3.3)-(3.5)*, the specification $\varphi$ is defined over the state space of both $\Sigma$ and $\widehat{\Sigma}$.*

**Remark 3.5.** *For a dt-SCS $\Sigma$ with linear dynamics $x(k+1) = Ax(k) + B\nu(k) + \varsigma(k)$, where $A = [a_{ij}]$, $B = [b_{ij}]$, and $\varsigma(k)$ is i.i.d. for $k = 0, 1, 2, \ldots$ with a normal distribution having zero mean and covariance matrix $\mathrm{diag}(\sigma_1, \ldots, \sigma_n)$, one can obtain $\mathscr{H} = \sum_{i,j} \frac{2|a_{ij}|}{\sigma_i \sqrt{2\pi}}$ and $\bar{\mathscr{H}} = \sum_{i,j} \frac{2|a_{ij}|}{\sigma_i \sqrt{2\pi}} + \sum_{i,j} \frac{2|b_{ij}|}{\sigma_i \sqrt{2\pi}}$. We refer the interested reader to (Soudjani & Abate 2013) for the computation of Lipschitz constant $\mathscr{H}$ and $\bar{\mathscr{H}}$ in the general case.*

In the next theorem, we present a condition such that the probabilistic distance between output trajectories of $\Sigma$ and $\widehat{\Sigma}$ is less than a given threshold, which is related to *(ii)* in Definition 1.2, as proposed by Lavaei et al. (2017). Let us note that this condition can work with either a finite abstraction or with an infinite abstraction with a lower-dimensional state space, which for instance can be constructed by means of a linear transformation of the state space, obtained with a rectangular matrix (cf. Theorem 4.2 and Figure 7).

**Theorem 3.6.** *Let $\Sigma = (X, U, \varsigma, f, Y, h)$ be a continuous-space dt-SCS and $\widehat{\Sigma} = (\hat{X}, \hat{U}, \varsigma, \hat{f}, \hat{Y}, \hat{h})$ be its abstraction, which can be either with a lower dimension or defined over a finite state set. Suppose there exists an SSF $V : X \times \hat{X} \to \mathbb{R}_{\geq 0}$ from $\widehat{\Sigma}$ to $\Sigma$ as in Definition 2.5. For any input trajectory $\hat{\nu}(\cdot) \in \hat{\mathbb{U}}$ that preserves the Markov property for the closed-loop $\widehat{\Sigma}$, and for any random variables $a$ and $\hat{a}$ as the initial states of dt-SCS $\Sigma$ and $\widehat{\Sigma}$, respectively, one can construct an input trajectory $\nu(\cdot) \in \mathbb{U}$ for $\Sigma$ through an interface function associated with $V$ (cf. Def. 2.5) such that:*

$$\mathbb{P}\Big\{ \sup_{0\leq k\leq T_d} \|y_{a\nu}(k) - \hat{y}_{\hat{a}\hat{\nu}}(k)\| \geq \varepsilon \,|\, [a;\hat{a}] \Big\} \leq \lambda_2, \tag{3.6}$$

*where,*

$$\lambda_2 := \begin{cases} 1 - (1 - \frac{V(a,\hat{a})}{\alpha(\varepsilon)})(1 - \frac{\hat{\psi}}{\alpha(\varepsilon)})^{T_d}, & \text{if } \alpha(\varepsilon) \geq \frac{\hat{\psi}}{1-\kappa}, \\ (\frac{V(a,\hat{a})}{\alpha(\varepsilon)})\kappa^{T_d} + (\frac{\hat{\psi}}{(1-\kappa)\alpha(\varepsilon)})(1 - \kappa^{T_d}), & \text{if } \alpha(\varepsilon) < \frac{\hat{\psi}}{1-\kappa}, \end{cases}$$

*with $\hat{\psi} \geq \rho_{\text{ext}}(\|\hat{\nu}\|_\infty) + \psi$, where $\alpha \in \mathcal{K}_\infty, 0 < \kappa < 1$, $\rho_{\text{ext}} \in \mathcal{K}_\infty \cup \{0\}$, and $\varepsilon, \psi \in \mathbb{R}_{\geq 0}$ as introduced in Def. 2.5.*

**Remark 3.7.** *Note that the closeness bounds in* (3.3)-(3.5) *are specification-dependent (cf. presence of formula $\varphi$ in the statements), whereas the provided closeness guarantee in* (3.6) *is specification-free. As a result, we observe that on the one hand the closeness bound in* (3.6) *can be considered to be more general than in* (3.3)-(3.5)*, however on the other it is likely to come at the cost of being more complex to compute and of being less tight.*

In order to establish the presented closeness guarantee between output trajectories of $\Sigma$ and $\widehat{\Sigma}$ (as per (3.6)), some conditions are required (*cf.* Lavaei et al. (2019), Zamani, Rungger & Mohajerin Esfahani (2017)), namely asking that the concrete model $\Sigma$ is *incrementally input-to-state stable* ($\delta$-ISS), as per Definition 2.4. This relates to the nature of the guarantee, pertaining models' trajectories. In contrast, notice that the closeness guarantee in (3.3) does not require original systems to be $\delta$-ISS: instead, only the Lipschitz continuity of the associated stochastic kernels is required for such guarantee (Soudjani, Abate & Majumdar 2015). Accordingly, the nature of the obtained guarantee is different.

On the other hand, since the abstraction error presented in (3.3) depends on the Lipschitz constants of the stochastic kernel, the error grows to infinity when the standard deviation of the noise goes to zero, which is not the case for (3.6). Thus, whilst different in nature, the bound in (3.6) can practically outperform that in (3.3) for noises with a small standard deviation, as long as the $\delta$-ISS assumption is satisfied by the original model. In addition, recent works (Haesaert, Soudjani & Abate 2017, Haesaert & Soudjani 2020, Lavaei, Soudjani & Zamani 2020*b*) have proposed a closeness guarantee as a version of (3.3) by establishing an approximate probabilistic relation between $\Sigma$ and $\widehat{\Sigma}$ based on a notion called $\delta$-lifting. The proposed framework is based on constructing an $\varepsilon$-expansion or $\varepsilon$-contraction of the set of interest (cf. $\varepsilon$ in (3.6)) over the abstract system. Accordingly, the probability of satisfaction computed over the modified sets on the abstract system provides upper and lower bounds for the probability of satisfaction on the original model.

We now present a result, related to condition *(iv)* in Definition 1.2, where the probability of satisfaction of a temporal logic property over the abstract system $\widehat{\Sigma}$ is either a lower or upper bound for the probability of property satisfaction over the original system $\Sigma$.

**Theorem 3.8.** *Let $\Sigma = (X, U, \varsigma, f, Y, h)$ be a continuous-space dt-SCS and $\widehat{\Sigma} = (\hat{X}, \hat{U}, \varsigma, \hat{f}, \hat{Y}, \hat{h})$ be its finite abstraction. For a given LTL specification $\varphi$, and for any policy $\hat{\nu}(\cdot) \in \hat{\mathbb{U}}$ that preserves the Markov property for the closed-loop $\widehat{\Sigma}$, one can construct a policy of $\nu(\cdot) \in \mathbb{U}$ for $\Sigma$ such that:*

$$\mathbb{P}(\widehat{\Sigma}_{\hat{\nu}} \vDash \varphi) \leq \mathbb{P}(\Sigma_{\nu} \vDash \varphi). \tag{3.7}$$

**Remark 3.9.** *As* (3.7) *provides a lower bound for the probability of satisfaction over $\Sigma$, it is mainly useful when one is interested in maximizing the satisfaction probability. Conversely, if the goal is to minimize the probability of satisfaction, one would want to search for an upper bound of the satisfaction probability. Such an upper bound can be quantified from* (3.7) *using the negation of the specification (i.e., $\neg\varphi$) as the following:*

$$\mathbb{P}(\Sigma_{\nu} \vDash \varphi) \leq 1 - \mathbb{P}(\widehat{\Sigma}_{\hat{\nu}} \vDash \neg\varphi). \tag{3.8}$$

One can employ the same approach as in (Lavaei et al. 2019, Section 6) and transfer the proposed closeness bound of (3.6) to (3.3) to any specification that can be accepted by a DFA (Kupferman & Vardi 2001). In particular, any LTL property $\varphi$ over the concrete system can be seen as the union of events over the product output space (these events can be shown to be measurable — cf. Remark 2.13). For instance, for a given safe set $\mathbb{S} \subseteq Y$, the safety property over a finite-time horizon $T$ is a subset of $Y^{T+1}$ (with $Y^{T+1} = \prod_{i=0}^{T} Y$) indicated by the set $\mathbb{S}^{T+1}$. For all measurable events $\mathsf{A} \subset Y^{T+1}$, one can construct an $\epsilon$-expansion and $\epsilon$-contraction of $\mathsf{A}$ over the abstract model within a given finite-time horizon $T$, as

$$\mathsf{A}^{\epsilon} := \{\{y(k)\}_{0:T} \in Y^{T+1} \,\big|\, \exists \{\bar{y}(k)\}_{0:T} \in \mathsf{A} \text{ s.t. } \max_{k \leq T} \|\bar{y}(k) - y(k)\| \leq \epsilon\},$$

$$\mathsf{A}^{-\epsilon} := \{\{y(k)\}_{0:T} \in Y^{T+1} \,\big|\, \forall \{\bar{y}(k)\}_{0:T} \in Y^{T+1} \backslash \mathsf{A}, \max_{k \leq T} \|\bar{y}(k) - y(k)\| > \epsilon\},$$

where $\{y(k)\}_{0:T} = [y(0); \ldots; y(T)]$, whose probabilities of satisfaction give respectively upper and lower bounds for the probability of satisfaction in the concrete domain with some quantified error bounds in the form of (3.3).

For the sake of completeness, let us remark that closeness conditions in *(iii)* in Definition 1.2 will be covered in Section 9, in the context of continuous-time SHS.

3.1. **Literature on Similarity Relations for Stochastic Models.** There has been substantial work in the area of Formal Methods on different types of stochastic similarity relations, which are employed to relate the probabilistic behavior of a concrete model to that of its abstraction and have been more recently studied for continuous-space models (Panangaden 2009, Abate 2013). Early on, similarity relations over finite-state stochastic systems via exact notions of probabilistic bisimulation relations have been introduced by Larsen & Skou (1991). Leveraging probabilistic transition systems as the underlying semantic model, the article shows how a testing algorithm can distinguish, with a probability arbitrarily close to one, between processes that are not bisimilar. Similarity relations over finite-state probabilistic models via exact probabilistic simulation

relations are also presented by Segala & Lynch (1995). In general, similarities are based on simulation or bisimulation relations, and can be either exact or approximate. Whenever the relation between a concrete model and its abstraction is symmetric, it is called "bisimulation relation." Exact simulation relations require the outputs of related systems to be exactly the same, while approximate simulation relations relax this requirement by allowing the outputs to differ up to a given error term (Baier & Katoen 2008, Tabuada 2009, Belta et al. 2017).

Admittedly, exact bisimulation relations raise very strong requirements amongst models, and in practice very limited classes of models can admit abstractions with those types of relations (Desharnais et al. 2008, D'Innocenzo et al. 2012). This is particularly true for continuous-space models (Abate 2013). Similarity relations of probabilistic models via approximate versions of probabilistic (bi)simulation relations are provided by Desharnais et al. (2008). The proposed framework is based on two-player games: the existence of a winning strategy for one of the players induces the $\epsilon$-(bi)simulation, and furthermore letting $\epsilon = 0$ gives back the exact notion. The paper also proposes a polynomial time algorithm to compute a derived metric, where the distance between states $s$ and $t$ is defined as the smallest $\epsilon$ such that $s$ and $t$ are $\epsilon$-equivalent.

An approximate probabilistic bisimulation relation for discrete-time Markov chains is proposed by D'Innocenzo et al. (2012). The provided scheme exploits the structure and properties of the approximate probabilistic bisimulation and leverages the mathematical framework of Markov set-chains (Hartfiel 2006) (related to Interval MC in Section 5.3) in order to provide a quantified upper bound on a metric over probabilistic realizations for labeled Markov chains. It is shown that the existence of an approximate probabilistic bisimulation relation implies the preservation of robust PCTL formulae.

Similarity relations for models with general, uncountable state spaces have also been proposed in the more recent literature (Panangaden 2009, Abate 2013). These relations can depend on stability requirements, on model's dynamics via martingale theory (Hall & Heyde 2014), or on contractivity analysis (Zamani, Mohajerin Esfahani, Majumdar, Abate & Lygeros 2014). Notably, the work by Zamani, Mohajerin Esfahani, Majumdar, Abate & Lygeros (2014) argues that every stochastic control system satisfying a probabilistic variant of incremental input-to-state stability ($\delta$-ISS), and for every given precision $\varepsilon > 0$, a finite-state transition system can be constructed that is $\varepsilon$-approximately bisimilar to the original stochastic control system. It also provides a closeness bound between the $\delta$-ISS stochastic control system and its bisimilar finite abstraction (cf. closeness in (9.3)).

Similarity relations of dt-SCS via approximate (bi)simul-ation relations are proposed by Desharnais et al. (2004), in which the relations enforce structural abstractions of a model by exploiting continuity conditions on its probability laws. Approximation metrics of stochastic processes, in particular Markovian processes in discrete time evolving on general state spaces (which are again domains with infinite cardinality and endowed with proper measurability and metric structures), are based on the notion of probabilistic bisimulation.

Labelled Markov processes (LMP) as probabilistic versions of labelled transition systems with continuous state spaces are widely discussed by Panangaden (2009) and related to dt-SCS. This book covers basic probability

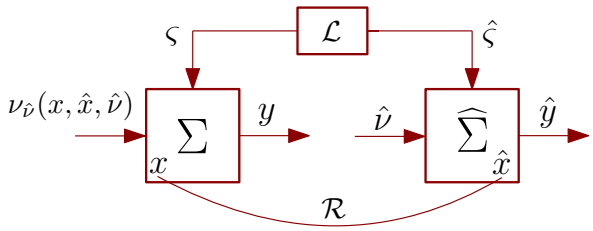


FIGURE 6. Notion of *lifting* for specifying the similarity between a dt-SCS $\Sigma$ and its abstraction $\widehat{\Sigma}$.

and measure theory on continuous state spaces and then develops the theory of LMPs. The main topics covered are bisimulation, the logical characterization of bisimulation, metrics and approximation theory.

Probabilistic model checking of dt-SCS via finite approximate bisimulations is proposed by Abate et al. (2014). The paper considers notions of (exact and approximate) probabilistic bisimulation and proposes a technique to compute an approximate probabilistic bisimulation of a dt-SCS, where the resulting abstraction is characterized as a finite-state Markov chain.

A notion of approximate similarity relation based on "lifted" probability measures is presented by Haesaert, Soudjani & Abate (2017), Lavaei, Soudjani & Zamani (2020*b*), which is inspired by notions of similarity relations proposed by Segala & Lynch (1995) for finite-state systems. The provided relation, underpinned by the use of metrics, allows in particular for a useful trade-off between deviations over probability distributions on states, and metric-based distances between model outputs. This new relation is inspired by a notion of simulation developed for finite-state models, and can be effectively employed over dt-SCS for both verification and synthesis purposes. The work also quantifies the distance in probability between the original system and its abstraction as a version of the closeness guarantee proposed in (3.3). The notion of lifting for specifying the similarity between a dt-SCS $\Sigma$ and its abstraction $\widehat{\Sigma}$ is schematically shown in Fig. 6. The relation $\mathcal{R}$ connects states of the two dt-SCS, and $\mathcal{L}$ specifies the relation between the two noises. The interface function $\nu_{\hat{\nu}}(x, \hat{x}, \hat{\nu})$ is used for refining a policy from the abstract system to the concrete one.

These notions and results are then generalized by Haesaert et al. (2018) to a larger class of temporal properties ((bounded) probabilistic reachability problems and co-safe LTL specifications) and by Haesaert & Soudjani (2020) to synthesize policies for a robust satisfaction of these properties, with applications in building automation systems (Haesaert, Cauchi & Abate 2017). An extension of these results to *networks* of dt-SCS is presented by Lavaei, Soudjani & Zamani (2020*b*), and will be discussed in more detail in Section 8.

A notion of approximate probabilistic trace equivalences for both finite-state Markov processes and dt-SCS, and its relation to approximate probabilistic bisimulation, is presented by Bian & Abate (2017). The proposed framework induces a tight upper bound on the approximation between finite-horizon traces, as expressed by a total variation distance. This bound can be employed to relate the closeness in satisfaction probabilities over bounded linear-time properties, such as in (3.3), and allows for probabilistic model checking of concrete models via their abstractions.

An approach for computing probabilistic bisimilarity distances for finite-state probabilistic automata has been proposed by van Breugel et al. (2021). The work proves that the bisimilarity distance bounds the difference in the maximal (or minimal) probability of two states to satisfy any arbitrary $\omega$-regular properties (*i.e.,* namely, the notion is specification-independent). As expected, since the proposed results should hold for any arbitrary $\omega$-regular specification, it can be much more conservative and difficult to be fulfilled or checked, compared to establishing the previously mentioned guarantees for a given specification.

We raise the following open challenge.

**Open Problem 1.** *Let $\Sigma_1$ and $\Sigma_2$ be two dt-SCSs. Develop an approach for computing the probabilistic bisimilarity distance between $\Sigma_1$ and $\Sigma_2$, satisfying any $\omega$-regular specification $\varphi$, as follows:*

$$|\mathbb{P}(\Sigma_1 \vDash \varphi) - \mathbb{P}(\Sigma_2 \vDash \varphi)| \leq \lambda_3, \quad \forall \varphi.$$

It is worth concluding this section emphasizing again that establishing stochastic similarity relations is crucial to connect the probabilistic behavior of an original SHS, which can be complex, to that of its abstraction. Consequently, by providing closeness guarantees between the output trajectories of two systems via the established stochastic similarity relations, one can perform formal analysis over the simpler abstraction and transfer the obtained results back to the original SHS.

## 4. Infinite Abstractions

The computational complexity associated to verifying or to synthesizing controllers for dt-SCS (and thus for SHS) models can be alleviated leveraging abstractions in two consecutive stages. In the first phase, the original complex systems can be abstracted by models either with simpler dynamics (e.g. linear, noiseless, etc.) or lower-dimensional state spaces (this is also known in the control literature as "model-order reduction" (Antoulas 2005, Ionescu & Astolfi 2015). Then one can employ those simpler models (a.k.a. infinite abstractions) as a replacement of original systems, perform analysis and synthesis over those models, and finally refine the results back (via an interface map) over the original models. Since the mismatch between outputs of original systems and those of their infinite abstractions is formally quantified, one can guarantee that concrete systems also satisfy the specifications as abstract ones with some guaranteed error bounds. An example of infinite abstractions is schematically depicted in Fig. 7. In comparison with Fig. 1, which focuses on discretization-based techniques to obtain finite abstractions, Fig. 7 focuses on infinite abstractions with lower-dimensional systems.

**Remark 4.1.** *Infinite abstractions can take numerous shapes and forms: they can for instance be linearized versions of the original models, they can be obtained via polynomial truncation, they can be models with different noises (e.g., stochastic realizations (van Schuppen 1989)), or models obtained by disregarding the noise terms (Zamani, Mohajerin Esfahani, Majumdar, Abate & Lygeros 2014, Zamani & Abate 2014*a*) (cf. Theorem 4.2 and Figure 7). The main focus of this section is placed on infinite abstractions with lower-dimensional state spaces, which are in practice compact representations of the concrete models.*

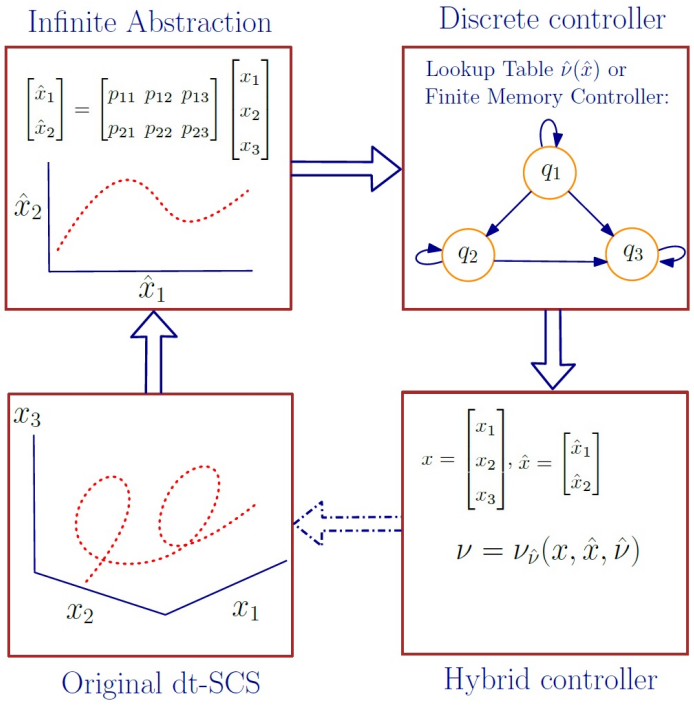


FIGURE 7. **Infinite abstractions.** The original dt-SCS has a 3-dimensional state set while its abstraction has a 2-dimensional state set. This model reduction can be performed via a transformation matrix $P$ satisfying conditions (4.4) and (4.5).

Note that one can construct finite abstractions directly, without going through infinite abstractions first. However, constructing finite abstractions for high-dimensional systems can result in large, finite state spaces, which might not be practically viable with limited computational and memory resources. One of the main benefits of infinite abstractions is thus to help reducing dimensions or complexity of concrete systems, which can then allow leveraging finite abstractions for the reduced-order models, while still providing the probabilistic closeness guarantees.

Developed earlier for continuous-time models (Julius & Pappas 2009, Abate 2009, Zamani, Rungger & Mohajerin Esfahani 2017) and further discussed in Section 9, the construction of infinite abstractions for discrete-time stochastic control systems is proposed by Lavaei et al. (2017) and by Lavaei, Soudjani & Zamani (2020*d*) and summarized in the results below. The abstraction framework is based on notions of stochastic simulation functions, introduced earlier (Def. 2.5). These functions relate output trajectories of an abstract system to those of the original one, such that the mismatch between the output trajectories of two systems remains within some guaranteed error bound. Through these stochastic simulation functions it is possible to quantify the probabilistic distance between the original stochastic system and its abstraction, based on the closeness in (3.6). The aforementioned work also focuses on a class of discrete-time *linear* stochastic control systems, as in (4.1) and further detailed next, and proposes a computational scheme to construct infinite abstractions together with their corresponding stochastic simulation functions.

Consider the class of discrete-time linear stochastic control system (a special instance of dt-SCS), as

$$\Sigma : \begin{cases} x(k+1) = Ax(k) + B\nu(k) + R\varsigma(k), \\ y(k) = Cx(k), \end{cases} \tag{4.1}$$

where the additive noise $\varsigma(k)$ is a sequence of independent random vectors with multivariate standard normal distributions. We use the tuple $\Sigma = (A, B, C, R)$ to refer to the class of linear systems in (4.1). In the next theorem, we establish a formal relation between $\Sigma$ and its reduced-order model $\widehat{\Sigma}$, by constructing corresponding matrices $\hat{A}, \hat{B}, \hat{C}, \hat{R}$.

**Theorem 4.2.** *Let $\Sigma = (A, B, C, R), \widehat{\Sigma} = (\hat{A}, \hat{B}, \hat{C}, \hat{R})$ be two linear dt-SCS with independent additive noises. Suppose there exist a matrix $K$ and a positive-definite matrix $M$ such that the following matrix inequalities*

$$C^T C \preceq M, \tag{4.2}$$

$$\big((1+\pi)(A+BK)^T M(A+BK) - M\big) \preceq -\hat{\kappa} M, \tag{4.3}$$

*hold for some constants $0 < \pi$ and $0 < \hat{\kappa} < 1$. If further*

$$AP = P\hat{A} - BQ, \tag{4.4}$$

$$CP = \hat{C}, \tag{4.5}$$

*hold for some matrices $Q$ and $P$ of appropriate dimension, then there exists a quadratic SSF $V(x, \hat{x})$ (Lavaei et al. 2017) between $\Sigma$ and $\widehat{\Sigma}$ as*

$$V(x, \hat{x}) = (x - P\hat{x})^T M(x - P\hat{x}), \tag{4.6}$$

*where $P \in \mathbb{R}^{n \times \hat{n}}$ is a matrix of an appropriate dimension with $\hat{n}$ being the dimension of the reduced-order model $\widehat{\Sigma}$.*

The stochastic simulation function $V(x, \hat{x})$ in (4.6) gives a probabilistic closeness guarantee between the original dt-SCS $\Sigma$ and its infinite abstraction $\widehat{\Sigma}$, as per (3.6).

**Remark 4.3.** *Condition* (4.4) *holds as long as condition (V.18) in (Zamani & Arcak 2018) is satisfied. In addition, notice that the results in Theorem 4.2 do not impose any condition on the matrix $\hat{B}$, which thus can be chosen arbitrarily. As an example, one can select $\hat{B} = \mathbb{I}_{\hat{n}}$, which renders the abstract system $\widehat{\Sigma}$ fully actuated and, hence, can facilitate a subsequent synthesis task.*

*Notice further that the matrix $\hat{R}$ can be also chosen arbitrarily. In this case, the probabilistic closeness between two systems $\Sigma$ and $\widehat{\Sigma}$ can be quantified as $\lambda_2$ in* (3.6)*, where*

$$\psi = \mathsf{Tr}\big(R^T M R + \hat{R}^T P^T M P \hat{R}\big).$$

*One can readily verify that selecting $\hat{R} = 0$ results in a tighter relationship between the original system $\Sigma$ and its infinite abstraction $\widehat{\Sigma}$. However, observe that this is not the case when the noises of the concrete system and of its infinite abstraction are the same, as assumed in (Zamani 2014, Zamani, Rungger & Mohajerin Esfahani 2017), in where $\hat{R}$ can be chosen appropriately to minimize the error term.*

The construction of infinite abstractions for dt-SCS is also discussed by Lavaei et al. (2019). The proposed approach employs the notion of stochastic storage function (a variant of the stochastic simulation function in Definition 2.5) between a concrete system and its abstraction, which allows to provide a closeness guarantee as in (3.6). This work also focuses on a specific class of discrete-time *nonlinear* stochastic systems by adding $E\tilde{\varphi}(Fx(k))$ to (4.1) in which $E \in \mathbb{R}^{n\times 1}$, $F \in \mathbb{R}^{1\times n}$, and $\Upsilon : \mathbb{R} \to \mathbb{R}$ is the nonlinearity term satisfying a slope restriction as

$$0 \leq \frac{\Upsilon(c) - \Upsilon(d)}{c - d} \leq b, \tag{4.7}$$

for any $c, d \in \mathbb{R}, c \neq d$, for some $b \in \mathbb{R}_{>0} \cup \{\infty\}$, and proposes a construction scheme for building infinite abstractions together with their corresponding stochastic storage functions.

It is worth mentioning that the contributions in (Lavaei et al. 2017, 2019) do not raise any restrictions on the sources of uncertainty in the concrete and abstract systems (*i.e.,* the noise of the abstraction can be completely independent of that of the concrete system). In particular, the results provided by Lavaei et al. (2017, 2019) are more general than (Zamani, Rungger & Mohajerin Esfahani 2017), where the noises in the concrete and abstract systems are assumed to be the same, which practically means the abstraction has access to the noise of the concrete system. The results in (Lavaei et al. 2017, 2019) provide a closeness guarantee between output trajectories of $\Sigma$ and $\widehat{\Sigma}$ as in (3.6).

To provide a broader context, approximations of large-scale dynamical systems in the context of model-order reduction are studied by Antoulas (2005) by combining system theory with numerical linear algebra. A notion of moment matching is presented by Ionescu & Astolfi (2015), which discussed a family of (nonlinear) parametrized reduced-order models that achieve moment matching.

A general framework for structure-preserving model reduction of a second-order network system based on graph clustering is studied by Cheng et al. (2017), where the dissimilarities of vertices are quantified by the $\mathcal{H}_2$-norms of the transfer function discrepancies. An $\mathcal{H}_2$ sub-optimal model reduction for second-order network systems is proposed by Yu et al. (2019), and an extension is recently presented by Yu et al. (2022), in which the main objective is to find a reduced-order model that not only approximates the input-output mapping of the original system but also preserves crucial model structure.

**Remark 4.4.** *Note that the model-order reduction techniques in (Yu et al. 2019, 2022, Cheng et al. 2017) deal with models in the frequency domain, and their main goal is to establish a closeness relation between the transfer function of the original system and of its reduced-order model by providing closeness guarantees based on the $\mathcal{H}_2$ norm. Since studies in the frequency domain are mainly developed for stability and input-output behaviour, handling more complex logical properties (such as the discussed safety, reachability, etc.) via those techniques is not straightforward. In comparison, the discussed infinite-abstraction techniques concerning models in the time domain can readily be employed to study verification and synthesis problems over logical specifications.*

**Running example (continued).** We consider the above running example, concerning a two-dimensional model, and now aim at constructing a one-dimensional infinite abstraction (*i.e.,* a proper reduced-order model),

by satisfying conditions (4.2)-(4.5). The two-dimensional room temperature regulation model is given by

$$\Sigma : \begin{cases} T(k+1) = AT(k) + \gamma T_h \nu(k) + \theta T_E + R\varsigma(k), \\ y(k) = CT(k), \end{cases}$$

where:

$$A = \begin{bmatrix} 1-2\sigma-\theta & \sigma \\ \sigma & 1-2\sigma-\theta \end{bmatrix}, \quad T_E = [T_{e_1}; T_{e_2}],$$
$$T(k) = [T_1(k); T_2(k)], \quad \nu(k) = [\nu_1(k); \nu_2(k)], \quad \varsigma(k) = [\varsigma_1(k); \varsigma_2(k)].$$

Moreover, $R = 0.01\mathbb{I}_2$, $C = \mathbb{1}_2^T$, $T_{e_i} = -1\,^\circ C$, $i \in \{1,2\}$, $T_h = 50\,^\circ C$, $\theta = 0.4$, $\gamma = 0.5$, and $\sigma = 0.1$ (the latter is a conduction factor between the two rooms). The goal is to construct a one-dimensional infinite abstraction $\widehat{\Sigma}$ from $\Sigma$ by satisfying conditions (4.2)-(4.5), which can be met by synthesizing

$$M = \mathbb{I}_2,\ P = \mathbb{1}_2,\ K = \mathbf{0}_{2\times 2},\ Q = \mathbb{1}_2,$$
$$\hat{A} = 25.5,\ \hat{C} = 1,\ \pi = 1,\ \hat{\kappa} = 0.34.$$

Then, there exists a quadratic SSF $V(x,\hat{x})$ between $\Sigma$ and $\widehat{\Sigma}$, as in (4.6). By taking $\hat{R} = 0.01$ and the initial states of the two models $\Sigma$ and $\widehat{\Sigma}$ to be equal to 20, and using the bound in (3.6), one can guarantee that the distance between the outputs of $\Sigma$ and $\widehat{\Sigma}$ does not exceed $\varepsilon = 1$ over the time horizon $T_d = 100$, with a probability of at least 95%, *i.e.,*

$$\mathbb{P}\Big\{\|y(k) - \hat{y}(k)\| \leq 1,\ \forall k \in [0, 100]\Big\} \geq 0.95.$$

One can utilize the obtained reduced-order model and construct a finite abstraction for later verification and synthesis purposes - this goal will be further discussed in the next section. □

## 5. Finite Abstractions

In the second phase of the abstraction procedure (cf. Fig. 1), one can construct finite abstractions usually in the form of finite Markov decision processes (MDPs). These abstractions are approximate descriptions of (reduced-order) systems, in which each discrete state corresponds to a set of continuous states of the (reduced-order) systems. Since the obtained abstractions are finite, one can employ algorithmic machinery and existing software tools to automatically synthesize controllers, which can then be applied (refined) over the concrete models, thus enforcing complex properties, including specifications expressed as temporal logical formulae.

A concrete model $\Sigma$ is approximated by a *finite* $\widehat{\Sigma}$ using Algorithm 1. For the sake of an easier presentation, we present the construction algorithm just for dt-SCS, however we refer the interested reader to (Zamani, Tkachev & Abate 2017, Tkachev et al. 2017, Haesaert, Soudjani & Abate 2017) for related, more general SHS, and to (Zamani & Abate 2014*b*, Zamani et al. 2015, Lavaei, Soudjani & Zamani 2020*a*) for the construction of finite MDPs for a class of SHS namely stochastic *switched* systems. To construct such a finite approximation, the state and input sets (over which one is interested to perform analysis and synthesis) of the dt-SCS $\Sigma$

are restricted to be compact.[3] The rest of the state space can be considered as a single absorbing state. Algorithm 1 first constructs a finite partition of the state set $X = \cup_i \mathsf{X}_i$ and the input set $U = \cup_i \mathsf{U}_i$. Then arbitrary "representative points" $\bar{x}_i \in \mathsf{X}_i$ and $\bar{\nu}_i \in \mathsf{U}_i$ are selected as abstract states and inputs. Transition probabilities in the finite MDP $\widehat{\Sigma}$ are computed according to (5.1). The output map $\hat{h}$ is the same as $h$ with its domain restricted to the finite set $\hat{X}$ (cf. Step 7) and the output set $\hat{Y}$ is the image of $\hat{X}$ under $h$ (cf. Step 6).

**Algorithm 1** Approximation of a dt-SCS $\Sigma$ by a finite MDP $\widehat{\Sigma}$

**Require:** Input dt-SCS $\Sigma = (X, U, T_{\mathsf{x}}, Y, h)$

1: Select finite partitions of sets $X, U$ as $X = \cup_{i=1}^{n_{\bar{x}}} \mathsf{X}_i$, $U = \cup_{i=1}^{n_{\bar{\nu}}} \mathsf{U}_i$

2: For each $\mathsf{X}_i$, and $\mathsf{U}_i$, select single representative points $\bar{x}_i \in \mathsf{X}_i$, $\bar{\nu}_i \in \mathsf{U}_i$

3: Define $\hat{X} := \{\bar{x}_i, i = 1, ..., n_{\bar{x}}\}$ as the finite state set of MDP $\widehat{\Sigma}$ with the finite input set $\hat{U} := \{\bar{\nu}_i, i = 1, ..., n_{\bar{\nu}}\}$

4: Define the map $\Xi : X \to 2^X$ that assigns to any $x \in X$, the corresponding partition set it belongs to, *i.e.,* $\Xi(x) = \mathsf{X}_i$ if $x \in \mathsf{X}_i$ for some $i \in \{1, 2, \ldots, n_x\}$

5: Compute the discrete *transition probability matrix* $\hat{T}_{\mathsf{x}}$ for $\widehat{\Sigma}$ as:

$$\hat{T}_{\mathsf{x}}(x'|x, \nu) = T_{\mathsf{x}}(\Xi(x')|x, \nu), \tag{5.1}$$

for all $x, x' \in \hat{X}, \nu \in \hat{U}$

6: Define the output space $\hat{Y} := h(\hat{X})$

7: Define the output map $\hat{h} := h|_{\hat{X}}$

**Ensure:** Output finite MDP $\widehat{\Sigma} = (\hat{X}, \hat{U}, \hat{T}_{\mathsf{x}}, \hat{Y}, \hat{h})$

Given a dt-SCS $\Sigma = (X, U, \varsigma, f, Y, h)$, the finite MDP $\widehat{\Sigma}$ constructed in Algorithm 1 can be represented as

$$\widehat{\Sigma} = (\hat{X}, \hat{U}, \varsigma, \hat{f}, \hat{Y}, \hat{h}), \tag{5.2}$$

where $\hat{f} : \hat{X} \times \hat{U} \times \mathcal{V}_\varsigma \to \hat{X}$ is defined as

$$\hat{f}(\hat{x}, \hat{\nu}, \varsigma) = \Pi_x(f(\hat{x}, \hat{\nu}, \varsigma)), \tag{5.3}$$

and $\Pi_x : X \to \hat{X}$ is the map that assigns to any $x \in X$, the representative point $\bar{x} \in \hat{X}$ of the corresponding partition set containing $x$. The initial state of $\widehat{\Sigma}$ is also selected according to $\hat{x}_0 := \Pi_x(x_0)$, with $x_0$ being the initial state of $\Sigma$.

The dynamical representation of the abstract finite MDP $\widehat{\Sigma}$ employs the map $\Pi_x : X \to \hat{X}$, which satisfies the inequality

$$\|\Pi_x(x) - x\| \leq \delta, \quad \forall x \in X, \tag{5.4}$$

where $\delta := \sup\{\|x - x'\|,\ x, x' \in \mathsf{X}_i,\ i = 1, 2, \ldots, n_x\}$ is the state discretization parameter.

[3]This compactness assumptions can be relaxed, albeit at the cost of additional (but quantifiable) approximation errors, as discussed in (Soudjani & Abate 2014*a*, 2015).

**Remark 5.1.** *Observe that the state discretization parameter $\delta$ appears in the probabilistic closeness quantified in* (3.3)-(3.6)*: thus, one can decrease the error by reducing the state discretization parameter, namely by aptly refining the state partitions. Notice that there is no requirement on the shape of the partition elements in constructing the finite MDPs. For the sake of an easier implementation, one can for instance consider partition sets as hyper-boxes, and representative points as centers of each box (cf. Fig. 8). The errors and guarantees derived above provide flexibility and have been embedded in a few software tools generating finite abstractions (cf. Soudjani, Gevaerts & Abate (2015), Lavaei, Khaled, Soudjani & Zamani (2020)).*

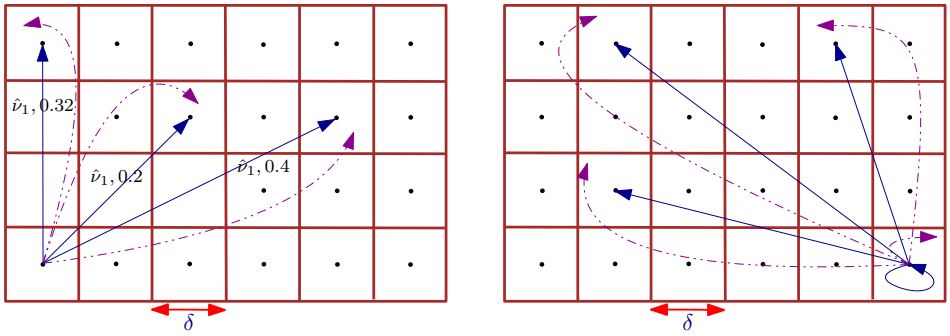


FIGURE 8. Construction of finite MDPs: A grid is first overlaid on the state and input sets. The center of each cell is considered as a representative point, and the transition probability (*i.e.,* probability of jumping from each representative point in a cell to all other cells) for all possible discrete inputs is computed. By repeating the process and storing the probabilities in a matrix called *transition probability matrix*, the corresponding finite MDP is accordingly constructed.

5.1. **Abstractions for Finite-Horizon Specifications.** The construction of finite abstractions of SHSs has been initially proposed by Abate et al. (2010) and used for formal verification and synthesis. This work investigates probabilistic safety and reachability over a finite-time horizon for a general class of discrete-time SHS with control inputs. The proposed framework characterizes the set of initial conditions providing a certain probabilistic guarantee that the system will keep evolving within a desired 'safe' region of the state space in terms of a value function, and determines 'maximally safe' Markov policies via dynamic programming over the finite abstract MDP. An improved gridding scheme, which is adaptive and sequential, for the abstraction and verification of stochastic processes is proposed by Soudjani & Abate (2013) (cf. Remark 5.1). The abstract model is constructed as a Markov chain using an adaptive gridding algorithm that conforms to the underlying dynamics of the model and thus mitigates the curse of dimensionality unavoidably related to the partitioning procedure. The work focuses on the study of a particular specification (probabilistic safety or invariance, over a finite horizon) and the results are then extended to SHS models with hybrid state spaces. The closeness guarantee between the original SHS and their finite abstractions is in the form of (3.3).

The above results in general rely on Lipschitz continuity of the stochastic kernel associated with the system. The works (Soudjani & Abate 2012*b*, 2014*b*) extend the method for systems with discontinuous stochastic

kernels and provide error bounds for inequalities of the form (3.3). On the other hand, if the kernel admits higher-order derivatives, refined computations are proposed by Soudjani & Abate (2012*a*), with errors that naturally depend on higher orders of the discretization parameter $\delta$ and that can show faster convergence to zero.

Among the many applications of these formal abstractions, the approach is employed to aggregate modeling and control of thermostatically controlled loads, as in (Soudjani et al. 2014), and in building automation (Cauchi & Abate 2018). In the aggregation procedure, each thermostatically controlled load model in the population, in principle similar to the Running Example in this paper, is formally abstracted as a finite MDP, and the cross product of these MDPs is lumped into its coarsest (exact) probabilistic bisimulation, and can be used for predictive energy scheduling. The abstraction procedure allows for the quantification of the induced error in the form of (3.3).

The construction of finite abstractions for stochastic control systems is presented by Soudjani, Abate & Majumdar (2015, 2017). These studies investigate the problem of finite-horizon probabilistic invariance for dt-SCS by providing a closeness guarantee between two systems in the form of (3.3). The proposed approach is more general than (Lavaei, Soudjani & Zamani 2020*d*, Lavaei et al. 2018), since the provided framework does not require original systems to be $\delta$-ISS. On the other hand, the abstraction error in (Soudjani, Abate & Majumdar 2015, 2017) depends on the Lipschitz constants of the stochastic kernels associated with the system, and accordingly, it grows to infinity as the standard deviation of the noise goes to zero, which is not the case in (Lavaei, Soudjani & Zamani 2020*d*, Lavaei et al. 2018).

A method to generate finite Markovian abstractions for discrete-time linear stochastic systems is presented by Lahijanian et al. (2012). The proposed approach proceeds by approximating the transition probabilities from one partition set to another by calculating the probability from a single representative point in the first region. The work employs an adaptive refinement algorithm that takes advantage of the dynamics of the system to achieve a desired error value. The proposed approach is similar to that of (Soudjani & Abate 2013) with a closeness guarantee in the form of (3.3), however here the transition probabilities are averaged over partition sets.

The construction of finite abstractions for discrete-time stochastic systems is also pursued by Lavaei et al. (2018). Focusing on a specific class of linear dt-SCS, these results employ notions of stochastic simulation (or storage) functions by providing a probabilistic distance between the interconnection of stochastic control subsystems and that of their finite abstractions based on (3.6).

The construction of finite MDPs for stochastic systems that are not necessarily stabilizable is presented by Lavaei, Soudjani & Zamani (2020*c*). The proposed frameworks rely on a relation between a system and its finite abstraction employing a new notion called *finite-step* stochastic simulation. In comparison with the existing notions of simulation functions in which stability or stabilizability of each subsystem is required, a *finite-step* stochastic simulation function needs to decay only after some finite numbers of steps (rather than at each time step). This results in a less conservative approach in the sense that one can compositionally construct finite MDPs such that stabilizability of each subsystem is not necessarily required. The work

in (Lavaei, Soudjani & Zamani 2020*c*) provides a closeness guarantee between output trajectories of $\Sigma$ and $\widehat{\Sigma}$ as per (3.6).

The construction of finite abstractions for stochastic *switched* systems is presented by Lavaei, Soudjani & Zamani (2020*a*), Lavaei & Zamani (2021). The transition map switches between a finite set of modes and the switched system accepts multiple Lyapunov (or storage) functions with a dwell-time condition that puts a lower bound on the interval between two consecutive switching time instants. The dwell-time is deterministic and always met by the controller designed using the finite MDP. In particular, switching signals in those works are control inputs and the main goal is to synthesize them with a specific dwell-time, such that the output of original systems satisfies some high-level specifications, such as safety, reachability, etc. Those works utilize notions of stochastic simulation (or storage) functions and provide a closeness guarantee in the form of (3.6) but adapted to the switched setup. These works also show that under standard assumptions ensuring incremental input-to-state stability of switched systems similar to Definition 2.4 (*i.e.,* existence of common incremental Lyapunov (or storage) functions, or multiple incremental Lyapunov (or storage) functions with some dwell-time conditions), one can construct finite MDPs for nonlinear stochastic switched systems. These results also propose an approach to construct finite MDPs together with their corresponding stochastic simulation (or storage) functions for a particular class of nonlinear stochastic switched systems whose nonlinearity $\Upsilon$ satisfies either a slope restriction similar to (4.7), or an incremental quadratic inequality as

$$\begin{bmatrix} d_2 - d_1 \\ \Upsilon_p(k, d_2) - \Upsilon_p(k, d_1) \end{bmatrix}^T \bar{Q}_p \begin{bmatrix} d_2 - d_1 \\ \Upsilon_p(k, d_2) - \Upsilon_p(k, d_1) \end{bmatrix} \geq 0, \tag{5.5}$$

for all $k \in \mathbb{N}$, $d_1, d_2 \in \mathbb{R}$, for all switching modes $p \in P = \{1, \ldots, m\}$, and for all $\bar{Q}_p \in \bar{\mathcal{Q}}_p$, where $\bar{\mathcal{Q}}_p$ is the set of symmetric matrices referred to as "incremental multiplier" matrices. For this class of nonlinear systems, the aforementioned incremental stability property can be readily checked via matrix inequalities. The quadratic inequality in (5.5) is called "incremental," as the difference between $d_1, d_2$ and their functions $\Upsilon_p(k, d_1), \Upsilon_p(k, d_2)$ appears in the inequality.

Abstraction-based synthesis of general MDPs using approximate probabilistic relations is proposed by Lavaei, Soudjani & Zamani (2020*b*). The abstraction framework is based on the notion of $\delta$-lifted relations, which is similar to (Haesaert, Soudjani & Abate 2017), using which one can quantify the distance in probability between dt-SCS and that of their abstractions as a version of the closeness guarantee proposed in (3.3). The works focus on a class of stochastic nonlinear dynamical systems and construct their (in)finite abstractions using both model order reduction and state space discretization.

A general abstraction technique for verifying safety problems for probabilistic hybrid systems is proposed by Zhang et al. (2010). Safety verification of linear discrete-time stochastic systems over bounded and unbounded time horizons is studied by Lal & Prabhakar (2020). For bounded safety verification, the work reduces the problem to the satisfiability of a semidefinite programming problem, whereas for the unbounded safety verification, the paper proposes an abstraction procedure to reduce the safety problem to that of a finite graph, wherein, the nodes of the graph correspond to the regions of a partition of the state space. A counterexample-guided abstraction refinement algorithm for a subclass of probabilistic hybrid systems,

called polyhedral probabilistic hybrid systems, is proposed by Lal & Prabhakar (2019), where the continuous dynamics are specified using a polyhedral set within which the derivatives of the continuous executions lie.

An abstraction-based reachability analysis for finite stochastic hybrid systems is studied by Zhang et al. (2017). The work addresses the problem of computing the probability of reaching a desired set in a subclass of SHS, wherein the stochasticity arises from the randomness of the initial distribution of continuous states, and the probabilistic transitions in the underlying finite-state Markov chain. Hierarchical abstractions for reachability analysis of probabilistic hybrid systems are proposed by Lal & Prabhakar (2018), in which discrete and probabilistic dynamics are captured using finite-state MDPs, and the continuous dynamics are modeled by annotating the states of the MDP with differential equations and inclusions. An abstraction-based framework to check probabilistic specifications of MDPs using stochastic two-player game abstractions is proposed by Kattenbelt & Huth (2009). The work also proposes a four-valued PCTL semantics for the developed game abstractions. A counterexample-guided abstraction refinement technique for the automatic verification of probabilistic systems is proposed by Hermanns et al. (2008). A survey on various abstraction-based techniques of probabilistic systems is presented by Dehnert et al. (2012).

**Running example (continued).** We construct a finite MDP from the model in (2.4) according to Algorithm 1, with the state discretization parameter $\delta = 0.005$. By taking the initial states of the two models $\Sigma$ and $\widehat{\Sigma}$ to be equal to 20, and using the proposed bound in (3.6), one can guarantee that the distance between the outputs of $\Sigma$ and $\widehat{\Sigma}$ does not exceed $\varepsilon = 0.5$ over the time horizon $T_d = 100$, with a probability of at least 98%, *i.e.,*

$$\mathbb{P}\Big\{\|y(k) - \hat{y}(k)\| \leq 0.5,\ \forall k \in [0, 100]\Big\} \geq 0.98. \tag{5.6}$$

Let us now synthesize a controller, where $U = [0, 0.6]$, for $\Sigma$ via its finite abstraction $\widehat{\Sigma}$, such that for $\Sigma$ the temperature of the room remains in the safe region $[19, 21]$. This is attained by employing the software tool AMYTISS (Lavaei, Khaled, Soudjani & Zamani 2020). The synthesized policy of the room as a function of state is illustrated in Fig. 10. Closed-loop state trajectories describing the dynamics of the room temperature over the finite-time horizon $T_d = 100$, under 10 different noise realizations, are illustrated in Fig. 9. We remark that the synthesized concrete policy for this example is simply chosen as $\nu = \hat{\nu}$, which is a special case of the interface function discussed in Remark 2.6.

In order to better understand the provided probabilistic bound in (5.6), we also run Monte Carlo simulation of 10000 runs. One expects that the distance between the outputs of $\Sigma$ and $\widehat{\Sigma}$ is always less than or equal to 0.5 with a probability at least of 98%, and indeed this bound is easily matched in practice. Indeed, we expected the empirical outcomes to be tighter, due to the conservative nature of Lyapunov-like techniques (simulation functions) and associated error bounds.

We now further elaborate the running example and quantify other closeness bounds discussed in Section 3. One can compute the closeness bound $\lambda_1$ in (3.3) with $\delta = 0.005$, $T_d = 100$, $\sigma = 0.6$, $\mathscr{H} = \dfrac{2|a|}{\sigma\sqrt{2\pi}} = 0.39$ (cf. Remark 3.5) as

$$|\mathbb{P}(\Sigma_{\hat{\nu}} \vDash \varphi) - \mathbb{P}(\widehat{\Sigma}_{\hat{\nu}} \vDash \varphi)| \leq 0.19. \tag{5.7}$$

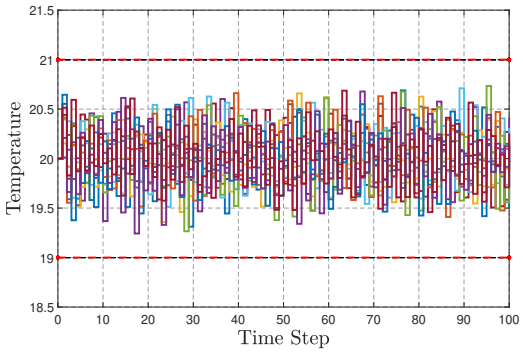


FIGURE 9. Closed-loop state trajectories with 10 different noise realizations for the finite time horizon $T_d = 100$.

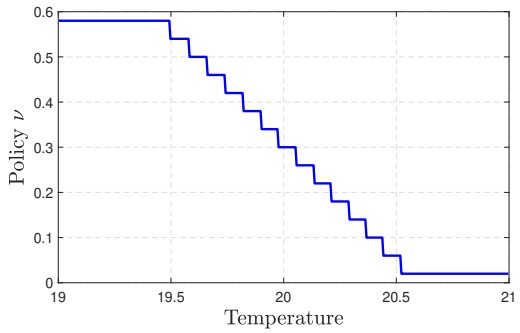


FIGURE 10. Synthesized policy as a function of state (room temperature).

According to Remark 3.5, one can quantify $\bar{\mathscr{H}} = \dfrac{2(|a| + |\gamma T_h|)}{\sigma\sqrt{2\pi}} = 17.02$, and accordingly, compute the closeness bound $\bar{\lambda}_1$ in (3.4) as 8.51. Similarly, one can readily compute the proposed closeness bound in (3.5) as 17.02. As can be observed, the obtained closeness error bounds from (3.4)-(3.5) are vacuous and even the closeness bound in (5.7) is more conservative than the one from (5.6). This issue is expected and the main reason is that closeness bounds (3.3)-(3.5) do not require the concrete model to be $\delta$-ISS (which is the case in (3.6)) but rather only the Lipschitz continuous. In addition, the closeness bounds (3.3)-(3.5) do not require the original and abstract systems to share the same source of stochasticity (which is the case in (3.6)), but this comes at the cost of providing more conservative closeness guarantees. One can readily improve the closeness bounds from (3.3)-(3.5) using an abstraction refinement approach. □

5.2. **Abstractions for Infinite-Horizon Specifications.** The construction of finite Markov chains for discrete-time stochastic models with continuous state spaces and their use to verify infinite-horizon properties (*e.g.,* safety and reachability specifications) are proposed by Tkachev & Abate (2011, 2012*a*,*b*, 2014). The proposed approaches employ notions of stochastic bisimulation functions and provide a lower bound for

infinite-time probabilistic invariance (cf. equations (3.7) and (3.8)) by decomposing this property into a finite-time reach-avoid together with an infinite-time invariance around absorbing sets (cf. Definition 5.2 below) over the state space of the model.

A quantitative abstraction-based controller synthesis for SHS is discussed by Tkachev et al. (2013) and later extended by Tkachev et al. (2017). The problem is reformulated as an optimization of a probabilistic reachability property over a product process (known as product automaton), which is obtained from the model of the specification and that of the system. The work develops a discretization procedure, which results in a standard synthesis problem over Markov decision processes with history-independent Markov policies, with errors of the form (3.3) and (3.4), respectively.

The satisfaction probability of infinite-horizon properties is theoretically investigated by Tkachev & Abate (2014). Extending to control-dependent models in (Tkachev et al. 2017), it is shown that the satisfaction probability depends on the existence of *absorbing sets*, as defined next.

**Definition 5.2.** *The set $\mathcal{A} \in \mathcal{B}(X)$ is called (weakly) absorbing if there exists a randomized selector $\bar{\mu}$ such that for all $x \in \mathcal{A}$, it holds that $\bar{\mu}(U \,|\, x) = 1$ and*

$$\int_U T_{\mathsf{x}}(\mathcal{A} \,|\, x, \nu)\bar{\mu}(\mathsf{d}\nu \,|\, x) = 1.$$

*We say that the set $\mathcal{A} \in \mathcal{B}(X)$ is simple if it does not have non-empty (weakly) absorbing subsets.*

Employing Definition 5.2, the following theorem is proposed by Tkachev et al. (2017).

**Theorem 5.3.** *The infinite-horizon safety probability for a continuous-space dt-SCS and a compact safe set $\mathbb{S}$ is equal to zero over the entire set if and only if the safe set $\mathbb{S}$ does not contain any absorbing sets.*

If the underlying system is a finite MDP, the simple absorbing sets in Definition 5.2 are bottom strongly connected components (BSCCs) of the MDP, and computing these BSCCs is straightforwardly done by graph search algorithms. Conversely, no computational method is proposed in the literature for finding absorbing sets of continuous systems. Motivated by Definition 5.2 and Theorem 5.3, we present the following open problem, which would allow to expand beyond the results in (Tkachev & Abate 2014, Tkachev et al. 2017).

**Open Problem 2.** *Given a dt-SCS with continuous-state space, compute its absorbing sets, or to compute over- and under-approximations of such absorbing sets within an a-priori precision.*

An alternative approach to handle infinite-horizon specifications is to employ interval MCs or interval MDPs, as discussed in the next subsection.

5.3. **Abstractions as Interval Markov Models.** The classical finite-state Markov models seen in the previous sections are not the only possible architecture for abstractions: uncertain Markov models can as well be employed for this task, and indeed they have been in particular employed to construct finite abstractions that are capable of satisfying infinite-time horizon properties in a more natural manner than standard Markov

models. Uncertain Markov models have been studied under different but related perspectives and semantics: (Junges 2020) provides an overview of existing approaches, and in particular focuses on models where the uncertainty is described parametrically, where probabilities are symbolic expressions rather than concrete values. (Junges 2020) discusses the parameter synthesis problem for the analysis of this class of Markov models. Alternatively, when the probabilities of transition between states belong to intervals, we can use interval Markov chains (IMC) and interval Markov decision processes (IMDP).

The definition of IMDP is similar to finite MDP as in Algorithm 1 with a tuple $\widehat{\Sigma}_{\mathrm{I}} = (\hat{X}, \hat{U}, \hat{T}_{\mathsf{x}_1}, \hat{T}_{\mathsf{x}_2}, \hat{Y}, \hat{h})$ where the exact transition probabilities are not known but are bounded above and below as $\hat{T}_{\mathsf{x}_1} \leq \hat{T}_{\mathsf{x}} \leq \hat{T}_{\mathsf{x}_2}$. An IMC is schematically depicted in Figure 11.

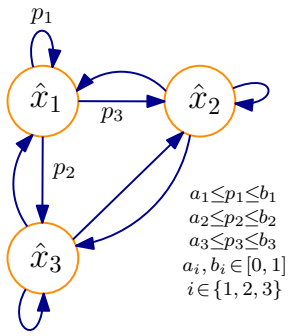


FIGURE 11. Example of an IMC.

It is worth mentioning that constructing IMCs/IMDPs can be more complicated compared to standard MCs/MDPs since one needs to provide both lower and upper bounds for probabilities of transitions among partition sets, by solving max-min optimization problems. However, one can mitigate the construction complexity if the system has a property called *mixed-monotonicity* and if the noise of the system has some nice properties.

**Definition 5.4.** *A function $f : X \to X$ is called mixed monotone if there exists a decomposition function $g : X \times X \to X$ satisfying (Smith 2008, Coogan & Arcak 2015)*

- $\forall x \in X \colon f(x) = g(x, x)$,
- $\forall x_1, x_2, z \in X \colon x_1 \leq x_2$ *implies* $g(x_1, z) \leq g(x_2, z)$,
- $\forall x, z_1, z_2 \in X \colon z_1 \leq z_2$ *implies* $g(x, z_2) \leq g(x, z_1)$.

Mixed monotonicity generalizes the notion of monotonicity in dynamical systems, which is recovered when $g(x, z) = f(x)$ for all $x, z$.

Consider the discrete-time stochastic system $x(k+1) = f(x(k)) + \varsigma(k)$, where $f(\cdot)$ is mixed monotone and entries of the noise $\varsigma(\cdot)$ are independent with unimodal distributions. Then, an IMC as abstraction of this system can be computed without the need for the optimization required in the computation of $\hat{T}_{\mathsf{x}_1}, \hat{T}_{\mathsf{x}_2}$ (Dutreix & Coogan 2020).

Specification-guided verification and abstraction refinement for mixed-monotone stochastic systems against omega-regular specifications are proposed by Dutreix & Coogan (2020). The article presents a procedure to compute a finite-state interval-valued Markov chain abstraction of discrete-time mixed-monotone stochastic systems subject to additive noise, given a rectangular partition of the state space. An algorithm is proposed for performing verification against omega-regular properties in IMCs that aims to compute bounds on the probability of satisfying a specification from any initial state of the IMC, in the form of (3.7). 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 original specification.

The results of (Dutreix & Coogan 2020) have been recently extended to the controller synthesis problem for discrete-time, continuous-state stochastic systems, under omega-regular specifications (Dutreix et al. 2022). The work presents a synthesis algorithm for optimizing the probability that a discrete-time stochastic switched system with a finite number of modes satisfies an omega-regular property. The approach relies on a finite-state abstraction of the underlying dynamics in the form of a bounded-parameter Markov decision process arising from a finite partition of the model's domain, with errors in the form of (3.7). Such Markovian abstractions allow for a range of probabilities of transitions between states for each selected action representing a mode of the original system. The proposed framework decomposes the synthesis into a qualitative problem, where the so-called greatest permanent winning or losing components of the product automaton are created.

**Remark 5.5.** *The results by Dutreix & Coogan (2020) and Dutreix et al. (2022) leverage the mixed-monotonicity property (cf. Definition 5.4) of the deterministic part of the map $f$ by assuming that (i) the stochasticity is additive, (ii) the distribution of the noise is unimodal, and (iii) noises of different states are independent from each other. This observation leads to the following open problem.*

**Open Problem 3.** *Provide a suitable definition of mixed-monotonicity for stochastic systems based on their stochastic kernels and investigate which classes of systems satisfy that property.*

An abstraction framework for mapping a discrete-time stochastic system to an IMC and mapping a switched dt-SCS to a bounded-parameter Markov decision process (BMDP) is proposed by Lahijanian et al. (2015). The work constructs model checking algorithms for IMCs and BMDPs against PCTL formulae to find sets of initial states that *definitely*, *possibly*, and *never* satisfy a given specification. It also develops an algorithm for BMDPs that synthesizes a policy maximizing the probability of satisfaction, and further proposes an adaptive refinement algorithm that exploits the dynamics of the system and the geometry of the partition to increase the precision of the solution. The work proposes a closeness guarantee in the form of (3.7).

Approximate abstractions of dt-SCS with interval MDPs are proposed by Zacchia Lun et al. (2018). The abstraction leverages the semantics of IMDPs and the standard notion of approximate probabilistic bisimulation. The resulting model presents a smaller one-step bisimulation error, in the form of equation (3.3) or (3.4), when compared to a Markov chain abstraction. The work outlines a method to perform probabilistic model checking, and shows that the computational complexity of the new method is comparable to that of standard abstractions based on approximate probabilistic bisimulations.

A constructive procedure for obtaining a finite abstraction of a discrete-time SHS is proposed by Abate et al. (2011), with errors as in equation (3.7) but usable over infinite horizons. Similar to the finite abstractions discussed above, the procedure consists of a partition of the state space of the system which depends on a controllable parameter. Given proper continuity assumptions on the model, the approximation errors introduced by the abstraction procedure are explicitly computed and it is shown that they can be tuned through the parameter of the partition. An analogous approach, with similar results, has been more recently pursued by Jaeger et al. (2020).

An efficient abstraction framework for formal analysis and control synthesis of a class of discrete-time SHS with linear dynamics is developed by Cauchi, Laurenti, Lahijanian, Abate, Kwiatkowska & Cardelli (2019). The work constructs IMDPs and focuses on temporal logic specifications over both finite- and infinite-time horizons. A strategy that maximizes the satisfaction probability of the given specification is synthesized over the IMDP and mapped to the underlying SHS. In contrast to existing formal approaches, which are by and large limited to finite-time properties and rely on conservative over-approximations, the article shows that the exact abstraction error can be computed as a solution of convex optimization problems and can be embedded into the IMDP abstraction. This is later used in the synthesis step over both bounded- and unbounded-time properties, mitigating the known state-space explosion problem but at the cost of lack of convergence guarantees. Much in the same line, IMDPs can be employed as abstract models for formal policy synthesis also for concrete dynamics described by deep, Bayesian neural networks (Laurenti et al. 2021).

The contribution in (Badings et al. 2022) presents a planning method for models with unknown disturbances, which computes a controller providing probabilistic guarantees on safely reaching a target. The continuous system is abstracted into an IMDP, adapting tools from the scenario approach (Campi & Garatti 2018) to compute probably approximately correct (PAC) assertions. The obtained IMDP is robust against uncertainty in the transition probabilities, and the tightness of the probability intervals can be controlled through the number of samples. Verification techniques are used to provide guarantees on the IMDP, and compute a controller for which these guarantees carry over to the concrete system.

**Open Problem 4.** *The discussed results in the setting of IMCs/IMDPs by and large provide a guarantee in the form of* (3.7)*. In particular, the satisfaction probability computed over the IMDPs gives a lower bound for the probability of satisfaction over the original system. Quantify instead the distance between the probability of satisfactions over the two systems in the form of equation* (3.3) *or* (3.4)*.*

## 6. Discretization-Free Verification and Synthesis

As discussed in the previous sections, discretization-free approaches can prevent the curse of dimensionality arising in the construction of finite abstractions. In this section we discuss discretization-free approaches based on (control) barrier certificates that have been proposed in recent years. Work in this domain follows theoretical (Prajna et al. 2007) and computational (Abate, Ahmed, Edwards, Giacobbe & Peruffo 2021) contributions, which however have been developed for continuous-time models. We first formally define control

barrier certificates in the context of this work, emphasizing that similar notions are also employed to study termination of related probabilistic programs (Roy et al. 2021).

**Definition 6.1.** *Consider a dt-SCS $\Sigma = (X, U, \varsigma, f, Y, h)$ with sets $X_0, X_u \subseteq X$ that are respectively initial and unsafe sets of the system. A function $\mathbb{B} : X \to \mathbb{R}_{\geq 0}$ is called a control barrier certificate (CBC) for $\Sigma$ if there exist constants $\eta, \beta \in \mathbb{R}_{\geq 0}$ with $\beta > \eta$ such that*

$$\mathbb{B}(x) \leq \eta, \qquad \forall x \in X_0, \tag{6.1}$$

$$\mathbb{B}(x) \geq \beta, \qquad \forall x \in X_u, \tag{6.2}$$

*and $\forall x \in X$, $\exists \nu \in U$, such that*

$$\mathbb{E}\Big[\mathbb{B}(f(x, \nu, \varsigma)) \;\Big|\; x, \nu\Big] \leq \max\Big\{\kappa \mathbb{B}(x), c\Big\}, \tag{6.3}$$

*for constants $0 < \kappa \leq 1$ and $c \in \mathbb{R}_{\geq 0}$.*

**Remark 6.2.** *Note that the existential quantifier for the condition in* (6.3) *implies the existence of a feedback controller for a model satisfying the conditions.*

Employing Definition 6.1, one can propose an upper bound on the probability that the dt-SCS in (2.2) reaches an unsafe region over a finite time horizon, as presented in the next theorem. Note that the requirement $\beta > \eta$ is needed in order to propose meaningful probabilistic bounds. Corollary 6.4 and the subsequent remark discuss the choice of the constant $c$ in the statement above.

**Theorem 6.3.** *Consider a dt-SCS $\Sigma = (X, U, \varsigma, f, Y, h)$ and a CBC $\mathbb{B}$ for $\Sigma$. Then the probability that the solution process of $\Sigma$ starts from any initial state $x(0) \in X_0$ and reaches $X_u$ under the policy $\nu(\cdot)$ (associated with the CBC $\mathbb{B}$) within the time interval $[0, T_d]$ is bounded by $\bar{\delta}$, namely*

$$\mathbb{P}\Big\{x(k) \in X_u \text{ for some } k \in [0, T_d] \;\Big|x(0) \in X_0\Big\} \leq \bar{\delta}, \tag{6.4}$$

*where if $0 < \kappa < 1$:*

$$\bar{\delta} := \begin{cases} 1 - (1 - \frac{\eta}{\beta})(1 - \frac{c}{\beta})^{T_d}, & \text{if } \beta \geq \frac{c}{\kappa - 1}, \\ (\frac{\eta}{\beta})\kappa^{T_d} + (\frac{c}{(1-\kappa)\beta})(1 - \kappa^{T_d}), & \text{if } \beta < \frac{c}{\kappa - 1}, \end{cases} \tag{6.5}$$

*whereas if $0 < \kappa \leq 1$:*

$$\bar{\delta} := \frac{\eta + cT_d}{\beta}. \tag{6.6}$$

*The upper bound proposed in* (6.5) *is less conservative than that of* (6.6) *in the sense that* (6.5) *yields a tighter probabilistic bound. On the other hand, the proposed bound in* (6.6) *is more general, since there may not exists a $\kappa$ strictly less than one satisfying condition* (6.3) *for many classes of models and dynamics.*

The results in Theorem 6.3 provide upper bounds on the probability that the solution process of $\Sigma$ reaches unsafe regions within a *finite* time horizon. One can generalize the proposed results to an *infinite* time horizon, provided that the constant $c = 0$, as stated in the following corollary.

**Corollary 6.4.** *Let $\Sigma = (X, U, \varsigma, f, Y, h)$ be a dt-SCS and suppose $\mathbb{B}$ is a CBC for $\Sigma$ with $c = 0$ in (6.3). Then the probability that the trajectory of $\Sigma$ starts from any initial state $x(0) \in X_0$ and reaches $X_u$ under the policy $\nu(\cdot)$ is bounded by*

$$\mathbb{P}\Big\{x(k) \in X_u \text{ for some } k \geq 0 \;\Big|\; x(0)\Big\} \leq \frac{\eta}{\beta}. \tag{6.7}$$

**Remark 6.5.** *Note that a CBC $\mathbb{B}$ satisfying condition (6.3) with $c = 0$ is a non-negative supermartingale (Kushner 1967, Chapter I). Although the supermartingale property on $\mathbb{B}$ allows one to provide probabilistic guarantees for infinite-time horizons via Corollary 6.4, it is restrictive in the sense that a supermartingale $\mathbb{B}$ may not exist (Steinhardt & Tedrake 2012, Jagtap et al. 2020). One may therefore employ a more general c-martingale type condition as in (6.3) that does not require such an assumption at the cost of providing probabilistic guarantees for finite time horizons.*

Note that the computation underlying CBC does not generate an abstract model, and accordingly it does not rely on any similarity relation and closeness error as presented in Definition 1.2. Instead, one can employ Definition 6.1 together with Theorem 6.4 and directly compute an upper bound on the probability that a dt-SCS reaches an unsafe region in a finite time horizon, much alike (3.7).

**Remark 6.6.** *Note that the verification problem is a special case of the synthesis one, in which the main goal is to verify that the property of interest is satisfied by means of some lower bound on the probability. The statements above can be accordingly tailored by changing the quantifier '$\exists$' in (6.3) to '$\forall$', and the results follow.*

6.1. **Computation of CBC and of Control Policies.** In this subsection, we discuss suitable methods to search for CBCs and to synthesize corresponding control policies. We study two different approaches based on (i) sum-of-squares (SOS) optimization and on (ii) counter-example guided inductive synthesis (CEGIS) (Jagtap et al. 2020).

6.1.1. *Sum-of-Squares Optimization Problems.* We reformulate conditions (6.1)-(6.3) as an SOS optimization problem (Parrilo 2003), where a CBC is restricted to be a non-negative polynomial that can be written as a sum of squares of different polynomials. To do so, the following assumption is required.

**Assumption 1.** *The dt-SCS $\Sigma$ has a continuous state set $X \subseteq \mathbb{R}^n$, and continuous input set $U \subseteq \mathbb{R}^m$. Moreover, the vector field $f : X \times U \times \mathcal{V}_\varsigma \to X$ is a polynomial function of the state $x$ and of the input $u$. Sets $X$ and $U$ are bounded semi-algebraic sets (i.e., they can be represented by the intersection of polynomial inequalities).*

Under Assumption 1, one can reformulate conditions (6.1)-(6.3) as an SOS optimization problem to search for a polynomial CBC $\mathbb{B}$ and a polynomial controller $\nu(\cdot)$ for the dt-SCS $\Sigma$. The following lemma provides the SOS formulation.

**Lemma 6.7.** *Suppose Assumption 1 holds and sets $X_0$, $X_u$, $X$, and $U$ can be defined as $X_0 = \{x \in \mathbb{R}^n \mid g_0(x) \geq 0\}$, $X_u = \{x \in \mathbb{R}^n \mid g_u(x) \geq 0\}$, $X = \{x \in \mathbb{R}^n \mid g(x) \geq 0\}$, $U = \{\nu \in \mathbb{R}^m \mid g_\nu(x) \geq 0\}$, where $g_0, g_u$, $g$ and $g_\nu$ are vectors of polynomials and inequalities are intended element-wise. Suppose for a given dt-SCS $\Sigma$, there exists a sum-of-squares polynomial $\mathbb{B}(x)$, constants $\eta, \beta, \bar{c} \in \mathbb{R}_{\geq 0}$, with $\beta > \eta$, $0 < \bar{\kappa} < 1$, vectors of sum-of-squares polynomials $\lambda_0(x), \lambda_u(x)$, $\lambda(x,\nu), \lambda_\nu(x,\nu)$, and polynomials $\lambda_{\nu_j}(x)$ corresponding to the $j^{th}$ input in $\nu = (\nu_1, \nu_2, ..., \nu_m) \in U \subseteq \mathbb{R}^m$ of appropriate dimensions, such that the following expressions are sum-of-squares polynomials:*

$$\begin{aligned}
&-\mathbb{B}(x) - \lambda_0^T(x) g_0(x) + \eta \\
&\mathbb{B}(x) - \lambda_u^T(x) g_u(x) - \beta \\
&-\mathbb{E}\Big[\mathbb{B}(f(x,\nu,\varsigma)) \mid x, \nu\Big] + \bar{\kappa}\mathbb{B}(x) + \bar{c} - \sum_{j=1}^{m} (\nu_j - \lambda_{\nu_j}(x)) - \lambda^T(x,\nu) g(x) - \lambda_\nu^T(x,\nu) g_\nu(x).
\end{aligned} \tag{6.8}$$

*Then $\mathbb{B}(x)$ is a CBC satisfying conditions* (6.1)-(6.3) *and $\nu = [\lambda_{\nu_1}(x); \ldots; \lambda_{\nu_m}(x)]$, is the corresponding controller of the dt-SCS $\Sigma$, where*

$$\kappa = 1 - (1-\pi)(1-\bar{\kappa}), \quad c = \frac{\bar{c}}{\pi(1-\bar{\kappa})},$$

*with $0 < \pi < 1$.*

For such computations, one can readily employ existing software tools available in the literature such as SOSTOOLS (Papachristodoulou et al. 2013), together with a semidefinite programming (SDP) solver (Sturm 1999, Yurtsever et al. 2021).

6.1.2. *Counter-Example Guided Inductive Synthesis.* One can find a CBC with a given parametric form, *e.g.,* a polynomial, by utilizing satisfiability modulo theories (SMT) solvers such as Z3 (De Moura & Bjørner 2008), dReal (Gao et al. 2012) or MathSat (Cimatti et al. 2013). The counter-example guided inductive synthesis (CEGIS) (Solar-Lezama et al. 2006) scheme can compute CBC for finite input sets, and it does not require any restrictions on underlying dynamics beyond what required by the SMT solver of choice. One can employ the following lemma and reformulate conditions (6.1)-(6.3) as a satisfiability modulo theory problem, as follows.

**Lemma 6.8.** *Consider the dt-SCS $\Sigma$. Suppose there exists a function $\mathbb{B}(x)$, constants $\eta, \beta, c \in \mathbb{R}_{\geq 0}$, and $0 < \kappa < 1$ such that*

$$\Theta(x) = \bigwedge_{x \in X_0} (\mathbb{B}(x) \leq \eta) \bigwedge_{x \in X_u} (\mathbb{B}(x) \geq \beta) \bigwedge_{x \in X} \bigvee_{\nu \in U} \Big( \mathbb{E}\Big[\mathbb{B}(f(x,\nu,\varsigma)) \big| x, \nu\Big] \leq \max\Big\{\kappa \mathbb{B}(x), c\Big\}\Big),$$

*where the index sets of conjunctions and disjunctions are possibly infinite. Then $\mathbb{B}(x)$ is a CBC satisfying conditions* (6.1)-(6.3).

Note that in the CEGIS approach, SMT solvers are employed to compute the CBC $\mathbb{B}(x)$ given a finite set $\overline{X} \subset X$ of data samples. If $\neg\Theta(x)$ has no feasible solution, this implies that $\mathbb{B}(x)$ is a true CBC. However, if $\neg\Theta(x)$ is feasible for some $\overline{x} \in X$, then $\overline{x}$ is a counter example. In this case, data samples should be updated to $\overline{X} = \overline{X} \cup \overline{x}$ and coefficients of the barrier should be recomputed iteratively until $\neg\Theta(x)$ becomes infeasible.

The control policy corresponding to the true CBC would be the sequence of inputs from the finite input set $U$ that renders $\Theta(x)$ feasible.

**Remark 6.9.** *The computational complexity in the construction of finite MDPs as in Algorithm 1 grows exponentially with the dimension of the state set. In contrast, in the case of sum-of-squares optimization, the computational complexity depends on both the degree of the polynomials and the number of state variables. It is shown that for fixed degree of the polynomials, the required computation grows polynomially with the dimension (Wongpiromsarn et al. 2015). Hence, we expect this technique to be more scalable than discretization-based approaches to study specific problems, such as safety analysis. The CEGIS approach (Jagtap et al. 2020) has a bottleneck that resides with the SMT solver, and it is difficult to provide any analysis on the computational complexity due to its iterative nature and lack of completeness (termination) guarantees.*

6.1.3. *Related Work on Barrier Certificates.* Within this line of work, the synthesis of invariants to study probabilistic safety of infinite-state models (probabilistic programs) is discussed by Chakarov & Sankaranarayanan (2013), much in line with results in Tkachev & Abate (2011, 2014). The proposed analysis employs concentration inequalities and martingales theory.

Finite-time safety verification of stochastic nonlinear systems using barrier certificates is proposed by Steinhardt & Tedrake (2012). The work considers the problem of bounding the probability of failure (defined as leaving a given bounded region of the state space) over a finite time horizon for continuous-time continuous -space stochastic nonlinear systems. The proposed approach searches for exponential barrier functions (*e.g.,* $\mathbb{B}(x,k) = e^{\frac{1}{2}x^T M(k)x} - 1$) that provide bounds using $c$-martingale type conditions as in (6.3), however in continuous time.

Probabilistic safety verification of systems using barrier certificates is proposed by Huang et al. (2017). The paper considers stochastic hybrid systems where the dynamics are represented as polynomial relations (equalities and inequalities) over the system variables and where random variables denote the initial discrete mode. The proposed approach guarantees the safety over an infinite time horizon. Control barrier certificates for a class of stochastic nonlinear systems against safety specifications are discussed by Liu et al. (2018). The proposed scheme provides probabilistic safety guarantees by reasoning over the stability properties of the model.

Temporal logic verification of stochastic systems via barrier certificates is proposed by Jagtap et al. (2018). The goal is to find a lower bound on the probability that a complex temporal logic property is satisfied by finite trajectories of the system (cf. (6.6)), in the spirit of (3.7). Considering properties expressed as safe LTL formulae, the proposed approach relies on decomposing the negation of the specification into a union of sequential reachabilities and then using barrier certificates to compute upper bounds for those reachability probabilities. The results of (Jagtap et al. 2018) are recently extended by Jagtap et al. (2020) to provide a formal synthesis framework for stochastic systems. The extended work distinguishes uncountable and finite input sets in the computation of control barrier certificates, using respectively SOS optimization and the CEGIS approach.

A controller synthesis framework for stochastic control systems based on control barrier functions is also provided by Clark (2019). The paper considers both complete information systems, in which the controller

has access to the full system information, as well as incomplete information systems where the state must be reconstructed from noisy measurements. In the complete information case, it formulates barrier functions that leads to sufficient conditions for safety with probability 1. However, in order to provide infinite-time horizon guarantees, this result requires that the control barrier functions exhibit supermartingale property, which presupposes stochastic stability and vanishing noise at the equilibrium point of the system. This approach is only applicable to systems with unbounded input sets and it does not provide any probabilistic guarantee 1 when the input set is bounded. In the incomplete information case, it formulates barrier functions that take an estimate from an extended Kalman filter (cf. Blom & Bloem (2004, 2007) for this notion in the context of SHS) as input, and derives bounds on the probability of safety as a function of the asymptotic error for the filter. The results in (Jahanshahi et al. 2020*a*) study formal synthesis of control policies for partially observed jump-diffusion systems (affected by both Poisson processes and Brownian motions) against complex logic specifications. Given a state estimator, the results in (Jahanshahi et al. 2020*a*) synthesize control policies providing (potentially maximizing) lower bounds on the probabilities that the trajectories of the partially observed jump-diffusion systems satisfy some complex specifications expressed by deterministic finite automata as in Definition 2.9.

A methodology for safety verification of non-stochastic systems using barrier certificates is proposed by Prajna & Rantzer (2005). Using the concepts of convex duality and density functions, the paper presents a converse statement for barrier certificates, showing that the existence of a barrier certificate is also necessary for safety. The results are then extended by Wisniewski & Sloth (2015) to more general classes of dynamical systems: in particular, (Wisniewski & Sloth 2015) proves converse barrier certificate theorems for a class of structurally stable dynamical systems. Dovetailing on these recent results, we present the following challenge, which is later generalized further.

**Open Problem 5.** *There is in general no a-priori guarantee on the existence of barrier certificates for a given SHS. In particular, Definition 6.1 provides a set of* sufficient *conditions for the existence of CBC. One interesting direction as a future work is to investigate* necessary and sufficient *conditions for the existence of control barrier certificates for SHS.*

**Open Problem 6.** *Develop computational techniques to construct CBC for general, nonlinear dt-SCS (not only polynomial-type models).*

**Running example (continued)**. The regions of interest in this example are considered as $X \in [1, 50]$, $X_0 \in [19.5, 20]$, and $X_u = [1, 17] \cup [23, 50]$. The main goal is to find a CBC for the system, for which a safety controller is synthesized maintaining the temperature of the room in a comfort zone $[17, 23]$.

We employ software tool SOSTOOLS (Papachristodoulou et al. 2013) and the SDP solver SeDuMi (Sturm 1999) to compute CBC as described in Definition 6.1. We compute CBC of an order 2 as $\mathbb{B}(T) = 0.86043T^2 - 33.78116T + 331.57433$ and the corresponding safety controller as $\nu(T) = -0.0120155T + 0.9$. Furthermore, the corresponding parameters in Definition 6.1 satisfying conditions (6.1)-(6.3) are quantified as $\eta = 0.13, \beta = 4.4, \kappa = 0.99$, and $c = 99 \times 10^{-4}$.

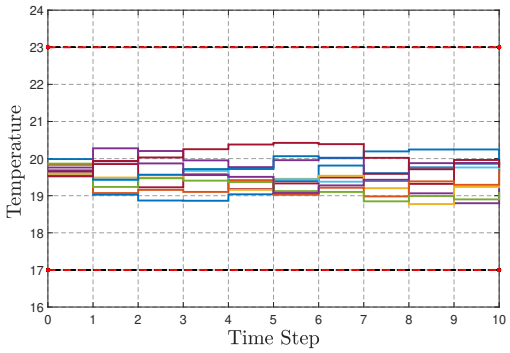


FIGURE 12. Closed-loop state trajectories with 10 different noise realizations for the finite-time horizon $T_d = 10$.

By employing Theorem 6.3, one can guarantee that the temperature of the room starting from the initial set $X_0 = [19.5, 20]$ remains in the safe set $X_u = [17, 23]$ during the time horizon $T_d = 10$ with a probability of at least 95%, *i.e.,*

$$\mathbb{P}\Big\{x(k) \notin X_u \text{ for all } k \in [0, T_d] \;\Big|\; a\Big\} \geq 0.95\,. \tag{6.9}$$

Closed-loop state trajectories with 10 noise realizations are illustrated in Figure 12. □

6.2. **Analysis of SHS with Optimization-based Methods.** For the sake of context and of completeness, we should mention literature dealing with alternative, discretization-free techniques for the analysis of SHS, which are mainly based on optimization approaches. Obviously the nature of these approaches is quite different than those presented so far, and indeed they tend to focus on other objectives, such as stability. However, in some instances, they might also accommodate some basic forms of specifications: for instance, safety might be able to be asserted for models that are shown to be stable. We should however remark that, in general, these alternative approaches are essentially different in nature than the techniques at the core of this survey.

Lyapunov-based conditions for stability and recurrence for a class of stochastic hybrid systems are presented by Teel (2013), where solutions are not necessarily unique, either due to nontrivial overlap of the flow and jump sets, a set-valued jump map, or a set-valued flow map. Different notions of stability for *stochastic* hybrid systems including Lyapunov, Lagrange, asymptotic stability, and recurrence analysis are overviewed by Teel et al. (2014). A moment closure technique for stochastic chemically reacting systems based on derivative-matching, which closes the moment equations by approximating higher-order moments as nonlinear functions of lower-order moments is investigated by Singh & Hespanha (2010*a*). A moment-based analysis for a class of SHS, so-called linear time-triggered SHS, is presented by Soltani & Singh (2017). The approach relies on embedding a Markov chain based on phase-type processes to model timing of events, and showing that the resulting system has closed moment dynamics.

A "shrinking-horizon" model predictive control (MPC) scheme for discrete-time linear systems with signal temporal logic (STL) specification constraints is proposed by Farahani et al. (2018). The control objective is to maximize a function under the restriction that a given STL specification is satisfied with high probability

against stochastic uncertainties. An MPC problem for a discrete-time linear system constrained to satisfy a co-safe LTL is studied by Gol et al. (2015). An overview of the main developments in the area of stochastic model predictive control (SMPC), together with potential perspectives for future research, can be found in (Mesbah 2016).

A Lagrangian technique to compute under-and-over approximations of target tube problem, a more generalized version of the finite-time horizon reach-avoid problem, for discrete-time nonlinear systems is proposed by Gleason et al. (2021). The proposed Lagrangian technique eliminates the necessity to grid the state, input, and disturbance spaces allowing for increased scalability and faster computation. A stochastic reachability problem, which maximizes the probability that the state remains within time-varying state constraints (*i.e.,* a "target tube"), despite bounded control inputs is studied by Vinod & Oishi (2021): the work proposes sufficient conditions under which the reach set is closed, compact and convex by providing an under-approximative interpolation technique for reach set. A linear programming approach for stochastic reach-avoid problems is proposed by Kariotoglou et al. (2017), where the objective is to synthesize a control policy to maximize the probability of reaching a target set at a given time, while staying in a safe set at all prior times. A class of stochastic reachability problems with state constraints from an optimal control perspective and set characterization for diffusions is proposed by Mohajerin Esfahani et al. (2016).

**Remark 6.10.** *Let us again comment on the comparison between abstraction-based techniques and SMPC, bearing in mind the discussed differences between the two classes of problems. Since abstraction-based approaches rely on discretizing state and input sets, they can readily handle any type of nonlinearity in models, and the spatial sets (characterizing specifications of interest) can be non-convex. In comparison, SMPC would be very challenging with nonlinear dynamics or non-convex constraints. Furthermore, whilst in principle SMPC may ensure a form of invariance via its recursive feasibility feature, it is not straightforward to enforce more complicated, high-level logical properties as optimization constraints in SMPC. In conclusion, SMPC is considered as an open and challenging problem (Mesbah 2016) that, in general, cannot be easily utilized to provide the formal guarantees on verification and controller synthesis for complex SHS which are the main focus of this survey paper.*

## 7. Temporal Logic Verification and Synthesis

In this section, we discuss how one can perform verification and synthesis for stochastic hybrid systems over interesting requirements, such as safety, reachability, or even more complex specifications encompassed by temporal logic or omega-regular languages. In presenting work at the interface between control theory and formal methods, we mainly focus on LTL and PCTL properties (or related expressions as automata) in this survey for the sake of better readability, thus leaving the survey of results on different temporal requirements to bibliographical pointers.

Let us start with basic specifications, expressed over DFAs: a quantitative, abstraction-based controller synthesis for SHS is proposed by Tkachev et al. (2013). The problem is first reformulated as an optimization of a probabilistic reachability property over a product process obtained from the model for the specification and the model of the system. The article develops a discretization procedure leading into standard dynamic

programming problems over finite MDPs with history-independent Markov policies. Errors are in the form of equation (3.4). A similar controller design scheme for stochastic hybrid systems is also provided by Kamgarpour et al. (2013). As a generalization, an optimal control synthesis approach defined over general discrete-time Markov decision processes is proposed by Tkachev et al. (2017) in which the probability of a given event is optimized: it is shown that the optimization over a wide class of LTL and $\omega$-regular properties can be reduced to the solution of one of two fundamental problems: reachability and repeated reachability.

A policy refinement scheme for dt-SCS via approximate similarity relations based on $\delta$-lifting is proposed by Haesaert, Soudjani & Abate (2017) by providing a closeness guarantee similar to (3.3). In particular, given safety properties over the concrete system, the work constructs an epsilon-perturbed specification over the abstract model whose probability of satisfaction gives a lower bound for the probability of satisfaction in the concrete domain with some quantified error bounds in the form of (3.3). The work is then generalized by Haesaert et al. (2018) to a larger class of temporal properties ((bounded) probabilistic reachability and other temporal logic specifications) and by Haesaert & Soudjani (2020) to synthesize policies for a robust satisfaction of specifications.

Policy synthesis with respect to co-safe linear temporal logic for stochastic control systems is proposed by Lavaei et al. (2019) in which it is discussed how synthesized policies for abstract systems can be refined back to original models while providing guarantees on the probability of satisfaction. All the results in (Haesaert, Soudjani & Abate 2017, Haesaert et al. 2018, Haesaert & Soudjani 2020, Lavaei et al. 2019) quantify a probabilistic distance between the original system and its epsilon-perturbed abstraction, as a version of closeness guarantee proposed in (3.3). We should highlight that, given the DFA $\mathcal{A}$ in Definition 2.9, the epsilon-perturbed specification in (Haesaert, Soudjani & Abate 2017, Haesaert et al. 2018, Haesaert & Soudjani 2020, Lavaei et al. 2019) corresponds to a new DFA $\hat{\mathcal{A}}_\varphi = (\bar{Q}_\ell, q_0, \bar{\Sigma}_{\mathsf{a}}, F_{\mathsf{a}}, \bar{\mathsf{t}})$ in which one absorbing location $q_{\mathsf{abs}}$ and one letter $\phi_\circ$ are added as $\bar{Q}_\ell := Q_\ell \cup \{q_{\mathsf{abs}}\}$ and $\bar{\Sigma}_{\mathsf{a}} := \Sigma_{\mathsf{a}} \cup \{\phi_\circ\}$. The initial and accept locations are the same with $\mathcal{A}_\phi$. The transition relation is defined, $\forall q \in \bar{Q}_\ell, \forall a \in \bar{\Sigma}_{\mathsf{a}}$, as

$$\bar{\mathsf{t}}(q,a) := \begin{cases} \mathsf{t}(q,a) & \text{if } q \in Q_\ell, a \in \Sigma_{\mathsf{a}}, \\ q_{\mathsf{abs}} & \text{if } a = \phi_\circ, q \in \bar{Q}_\ell, \\ q_{\mathsf{abs}} & \text{if } q = q_{\mathsf{abs}}, a \in \bar{\Sigma}_{\mathsf{a}}. \end{cases}$$

In other words, an absorbing state $q_{\mathsf{abs}}$ is added and all states will jump to this absorbing state with label $\phi_\circ$. As an example, the modified DFA of the reach-avoid specification in Figure 5 is plotted in Figure 13.

Temporal logic verification and synthesis of stochastic systems via control barrier certificates against a fragment of linear temporal logic, *i.e.,* safe LTL, over finite traces are presented by Jagtap et al. (2018, 2020). Those results got extended to $\omega$-regular specifications in (Anand et al. 2021).

Forward stochastic reachability analysis for uncontrolled linear systems with affine (bounded or unbounded) disturbance is presented by Vinod et al. (2017). The proposed method utilizes Fourier transforms to efficiently compute the forward stochastic reach probability measure (density) and the forward stochastic reach set. Underpinned by the same technique, an under-approximation of the stochastic reach-avoid probability for high-dimensional linear stochastic systems is presented by Vinod & Oishi (2017) while providing guarantees of

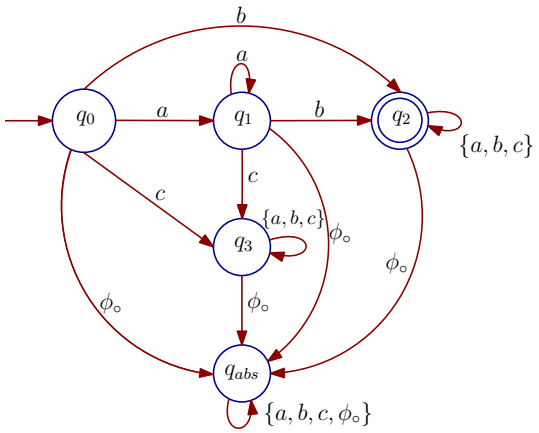


FIGURE 13. Modified DFA $\hat{\mathcal{A}}_\varphi$ of the specification $(a \cup b)$.

the type is equations (3.7) and (3.8). The proposed framework exploits fixed control sequences parameterized by the initial condition (an open-loop control policy) to generate the under-approximation. For Gaussian disturbances, the under-approximation can be obtained using existing efficient algorithms by solving a convex optimization problem. The work in (Vinod & Oishi 2018) proposes a scalable algorithm to construct a polytopic under-approximation of the terminal hitting time stochastic reach-avoid set, for the verification of high-dimensional linear stochastic systems with arbitrary stochastic disturbance. The existence of a polytopic under-approximation is proved by characterizing sufficient conditions under which the stochastic reach-avoid set and the proposed open-loop under-approximation are compact and convex.

A framework for analyzing probabilistic safety and reachability problems for discrete-time SHS, in scenarios where system dynamics are affected by competing agents, is proposed by Kamgarpour et al. (2011). The provided framework considers a zero-sum game formulation of the probabilistic reach-avoid problem, in which the control objective is to maximize the probability of reaching a desired subset of the hybrid state space, while avoiding an unsafe region, subject to the worst-case behavior of a rational adversary. The results are then extended by Ding, Kamgarpour, Summers, Abate, Lygeros & Tomlin (2013) to demonstrate how the proposed results can be specialized to address the safety problem, by computing the minimal probability that the system state reaches an unsafe subset of the state space.

Under-approximation of finite-time horizon, stochastic reach-avoid sets for discrete-time stochastic nonlinear systems is discussed by Gleason et al. (2017) via Lagrangian methods. The article utilizes the concept of target-tube reachability to define robust reach-avoid sets that are parameterized by the target set, safe set, and the set which the disturbance is drawn from. The proposed framework unifies two existing Lagrangian approaches to compute these sets, and establishes that there exists an optimal Markov control policy for the robust reach-avoid sets. The results characterize a subset of the disturbance space whose corresponding robust reach-avoid set for a given target and safe set is a guaranteed underapproximation of the stochastic reach-avoid level set of interest. Although the proposed method is conservative, it does not rely on a grid, implying scalability features that now hinge on geometrical computations.

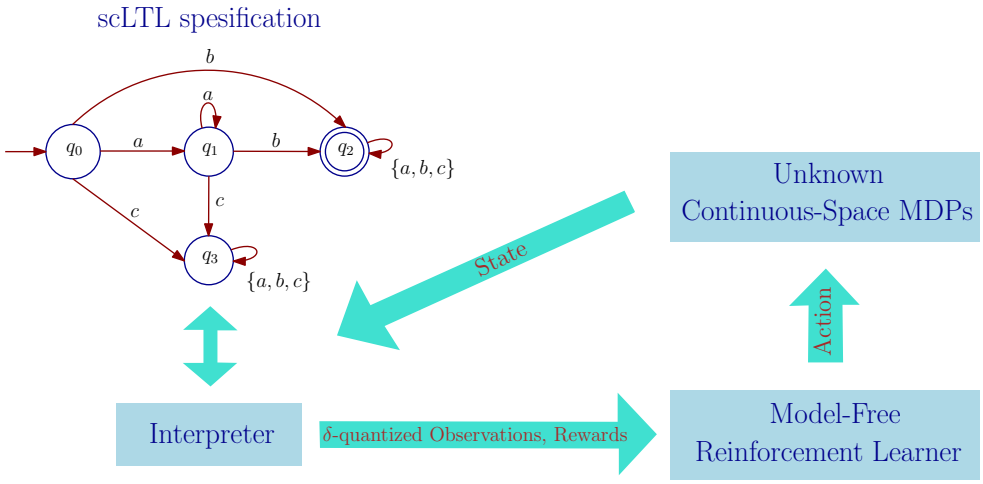


FIGURE 14. Model-free reinforcement learning is employed by a DFA corresponding to an scLTL objective. In particular, the $\delta$-quantized observation set of the dt-SCS $\Sigma$ is used by an *interpreter* process to compute a run of the DFA. When the run reaches a final state, the interpreter gives the reinforcement learner a positive reward and the training episode terminates. Any converging reinforcement learning algorithm over such $\delta$-quantized observation set is guaranteed to maximize the probability of satisfaction of the scLTL objective and converge to an optimal strategy over the unknown dt-SCS $\Sigma$.

(Lesser & Abate 2018) investigate multi-objective optimal control for dt-SHS with safety as a priority, by means of a *lexicographic* approach that priorities the safety as a constraint to be met prior to optimizing a given reward. The work by Haesaert et al. (2021) proposes an abstraction framework for computation of policies to satisfy multiple specifications with different priorities by encoding them in a multi-objective framework. A tutorial covering multi-objective probabilistic model checking, to analyze trade-offs between several different quantitative properties, is in (Forejt et al. 2011).

A computational framework for the automatic deployment of a robot with sensor, actuator noise and temporal logic specifications is proposed by Lahijanian et al. (2011). The work models the motion of the robot in the environment as a finite-state MDP and translates the motion specification to a formula of probabilistic computation tree logic. There are alternative approaches that deploy similar results on robotics applications modelled via stochastic systems (Lacerda et al. 2014).

We discuss briefly the formal synthesis over SHS via learning and data-driven approaches in Subsection 12.1 as an open research direction. Here, to conclude this section, we only present a few limited recent works with a focus on temporal logic verification and synthesis via learning approaches. A reinforcement learning framework for controller synthesis of finite MDPs with unknown transition probabilities against LTL objectives with a proof of convergence is proposed by Hasanbeig et al. (2019*a*), Hahn et al. (2019). A key feature of the proposed techniques is the compilation of $\omega$-regular properties into limit deterministic Büchi automata (LDBA), instead

of the Rabin automata that are standard with MDPs. If the dt-SCS is finite, theoretical guarantees are provided on the convergence of the RL algorithm to an optimal policy, maximizing the satisfaction probability.

A model-free reinforcement learning scheme to synthesize policies for unknown continuous-space dt-SCS is proposed by Lavaei, Somenzi, Soudjani, Trivedi & Zamani (2020). The proposed approach is schematically illustrated in Figure 14: the properties of interest for the system belong to so-called syntactically co-safe linear temporal logic formulae, and the synthesis requirement is to maximize the probability of satisfaction within a given bounded time horizon. The work provides control strategies maximizing the probability of satisfaction over unknown continuous-space dt-SCS while providing probabilistic closeness guarantees in the form of (3.3). Similarly based on the scheme in Figure 14, extensions to continuous spaces and $\omega$-regular properties are studied by Hasanbeig et al. (2019*b*), Kazemi & Soudjani (2020), as well as by Hasanbeig et al. (2020), Cai et al. (2021) by means of deep neural nets generalizers, but without providing optimality guarantees for the synthesized policies when refined over unknown dt-SCS.

**Open Problem 7.** *Provide (approximate) optimality guarantee for learning-based approaches that compute a controller for dt-SCS to satisfy any given LTL specification.*

## 8. Compositional Techniques

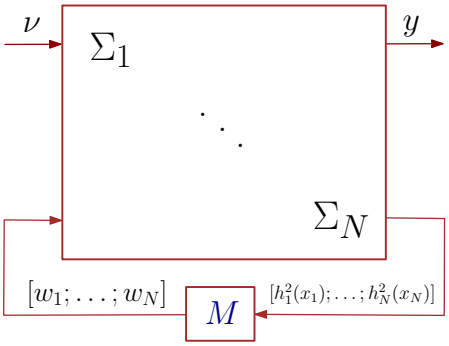


Figure 15. An interconnected dt-SCS with external input and output signals $\nu$ and $y$, respectively. Note that in so-called small-gain settings discussed later, the interconnection matrix $M$ is a permutation matrix that results in an element-wise interconnection constraint (*i.e.,* $\forall i,j \in \{1,\ldots,N\}, i \neq j\colon w_{ij} = h_{ji}^2(x_j)$).

It is of interest to extend the techniques introduced in previous sections to interconnected models, or to models with specially structured dynamics or coupling between variables. Moreover, the construction of (in)finite abstractions for large-scale stochastic hybrid systems in a monolithic manner suffers severely from the curse of dimensionality. To mitigate this issue, one promising solution is to consider a large-scale model as an interconnected system composed of several smaller subsystems. Compositional techniques are specifically suitable to tackle these problems, and are broadly studied in this section. We overview results on compositional

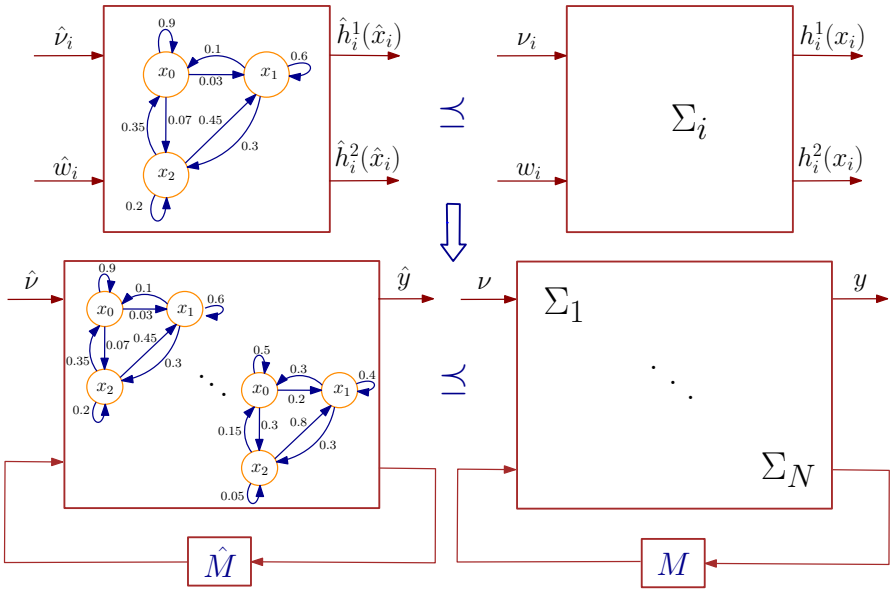


FIGURE 16. Representation of compositionality results. Here, we denote $\widehat{\Sigma} \preceq \Sigma$ if there exists an SSF $V$ from $\widehat{\Sigma}$ to $\Sigma$ (cf. Definition 2.5).

frameworks for the construction of (in)finite abstractions for interconnected systems using abstractions of smaller subsystems.

We first define stochastic control *subsystems* next. The term "internal" is employed for inputs and outputs of subsystems that are affecting each other in the interconnection: an internal output of a subsystem affects an internal input of another subsystem downstream. The term "external" instead is utilized to denote (exogenous) inputs and outputs that are not employed for the construction of the interconnection.

**Definition 8.1.** *A discrete-time stochastic control subsystem (dt-SCS) is described by the tuple*

$$\Sigma = (X, U, W, \varsigma, f, Y^1, Y^2, h^1, h^2), \tag{8.1}$$

*where*

- $X \subseteq \mathbb{R}^n$ *is a Borel space as the state space of the subsystem;*
- $U \subseteq \mathbb{R}^m$ *is a Borel space as the* external *input space of the subsystem;*
- $W \subseteq \mathbb{R}^{\bar{p}}$ *is a Borel space as the* internal *input space of the subsystem;*
- $\varsigma$ *is a sequence of i.i.d. random variables from a sample space* $\Omega$ *to the measurable space* $(\mathcal{V}_\varsigma, \mathcal{F}_\varsigma)$*;*

$$\varsigma := \{\varsigma(k) : (\Omega, \mathcal{F}_\Omega) \to (\mathcal{V}_\varsigma, \mathcal{F}_\varsigma),\ k \in \mathbb{N}\};$$

- $f : X \times U \times W \times \mathcal{V}_\varsigma \to X$ *is the transition map;*
- $Y^1 \subseteq \mathbb{R}^{q^1}$ *is a Borel space as the* external *output space of the subsystem;*
- $Y^2 \subseteq \mathbb{R}^{q^2}$ *is a Borel space as the* internal *output space of the subsystem;*
- $h^1 : X \to Y^1$ *is a measurable function as the* external *output map that takes a state* $x \in X$ *to its* external *output* $y^1 = h^1(x)$*;*

- $h^2 : X \to Y^2$ *is a measurable function as the* internal *output map that takes a state* $x \in X$ *to its* internal *output* $y^2 = h^2(x)$.

Properties of the interconnected system are specified over external outputs, as in Definition 8.1, and the synthesis objective is to control external inputs in order to satisfy desired properties over external outputs; whereas internal signals are utilized for the sake of interconnections amongst subsystems. We are now well equipped to define an interconnected dt-SCS.

**Definition 8.2.** *Consider* $N \in \mathbb{N}_{\geq 1}$ *stochastic control subsystems* $\Sigma_i = (X_i, U_i, W_i, \varsigma_i, f_i, Y_i^1, Y_i^2, h_i^1, h_i^2)$, $\forall i \in \{1, \ldots, N\}$, *and a matrix* $M$ *of appropriate dimensions, defining the coupling between these subsystems. The interconnection of* $\Sigma_i$, $i \in \{1, \ldots, N\}$, *is the dt-SCS* $\Sigma = (X, U, \varsigma, f, Y, h)$, *denoted by* $\mathcal{I}(\Sigma_1, \ldots, \Sigma_N)$, *such that* $X := \prod_{i=1}^N X_i$, $U := \prod_{i=1}^N U_i$, $f := \prod_{i=1}^N f_i$, $Y := \prod_{i=1}^N Y_i^1$, *and* $h = \prod_{i=1}^N h_i^1$, *subjected to the following interconnection constraint:*

$$[w_1; \ldots; w_N] = M[h_1^2(x_1); \ldots; h_N^2(x_N)].$$

An interconnected dt-SCS based on Definition 8.2 is schematically depicted in Fig. 15.

In this section, we discuss two different compositional approaches based on small-gain and dissipativity conditions. Small-gain and dissipativity techniques have been traditionally employed in the context of stability analysis for networks of interconnected systems (Dashkovskiy et al. 2010, Arcak et al. 2016). Suppose $\Sigma$ is an interconnected dt-SCS with $N$ stable subsystems $\Sigma_1, \ldots, \Sigma_N$. Under some small-gain or dissipativity conditions, one can ensure that the composed network $\Sigma = \mathcal{I}(\Sigma_1, \ldots, \Sigma_N)$ is also stable. A similar idea can be utilized here in the setting of similarity relations, using which one can establish a formal relation between an interconnected system $\Sigma$ and its abstraction $\widehat{\Sigma}$, based on relations between subsystems and their corresponding abstractions. We present the semantics of compositionality techniques in the following.

**Semantics of compositionality techniques.** Consider an interconnected dt-SCS $\Sigma = \mathcal{I}(\Sigma_1, \ldots, \Sigma_N)$, and assume proper relations between subsystems $\Sigma_i$ and their corresponding abstractions $\widehat{\Sigma}_i$ in the sense of Definition 2.5. Under some compositionality conditions, one can construct an overall relation between the two interconnected systems $\widehat{\Sigma} = \widehat{\mathcal{I}}(\widehat{\Sigma}_1, \ldots, \widehat{\Sigma}_N)$[4] and $\Sigma$, based on the relations between the subsystems and their abstractions.

Compositionality conditions based on dissipativity approaches are in the form of LMI that can be readily checked via semidefinite programming (SDP) solvers such as SeDuMi (Sturm 1999). For small-gain reasoning, we distinguish the corresponding compositionality conditions based on so-called sum-type and max-type small-gain approaches. In particular, in sum-type small-gain approach, the second condition of the stochastic simulation function (SSF) is in the form of (2.7), and the overall SSF is a weighted sum of SSF of subsystems. Accordingly, one deals with a spectral radius of some matrix that needs to be strictly less than one as the compositionality condition. In contrast, in the max-type small-gain approach, the upper bound in (2.7) is in the max form and the overall SSF is based on the maximum of SSF of subsystems. We refer the interested

[4]Interconnection topology in the abstract domain can be constructed similar to the interconnection topology of the concrete domain (see *e.g.,* (Lavaei, Soudjani & Zamani 2020*d*, Section VI)).

readers for more details on the compositionality conditions in (sum and max-type) small-gain and dissipativity approaches respectively to (Lavaei et al. 2017, Lavaei, Soudjani & Zamani 2020*d*) and (Lavaei et al. 2019). It will be further discussed in this section that the closeness guarantees are of the type in (3.6).

Compositionality results have been schematically depicted in Fig. 16. As illustrated, if there exists a local stochastic simulation function between each original subsystem and its corresponding finite MDP, one can construct an overall stochastic simulation function between the original interconnected system and its interconnected finite abstraction provided that some compositionality conditions are satisfied.

**Remark 8.3.** *Note that the proposed compositionality results based on sum-type small-gain approaches (Lavaei et al. 2017) require linear growth on gains of subsystems (cf. Lavaei et al. (2017, Assumption 1)) and provide an additive overall error (i.e., the error of the interconnected abstraction is linear combination of errors of abstractions of subsystems). In contrast, the* max*-type small-gain approaches (Lavaei, Soudjani & Zamani 2020*d*) are more general, since they do not require any linearity assumption on gains of subsystems and the overall error is the maximum error of abstractions of subsystems. Both errors provide closeness guarantees of the type in* (3.6)*. On the other hand, checking the compositionality condition in the sum-type small-gain is much easier than the* max*-type one, since it is based on the spectral radius of some matrix that can be easily checked.*

Compositional techniques based on infinite and finite abstractions have been schematically illustrated in Fig. 17.

A compositional reasoning methodology for the design of *finite* systems with stochastic and/or non-deterministic aspects is proposed by Delahaye et al. (2011). The work focuses on models of assume/guarantee contracts for stochastic systems, in which the contract allows to distinguish hypotheses made on a system (the guarantees) from those made on its environment (the assumptions). An automated technique for assume-guarantee style checking of strong simulation between a system and a specification, both expressed as non-deterministic *finite* labeled probabilistic transition systems is presented by Komuravelli et al. (2012).

Compositional construction of finite abstractions for stochastic control systems is presented by Soudjani, Abate & Majumdar (2015, 2017). These results investigate the finite-horizon probabilistic invariance for dt-SCS and provide a closeness guarantee between two systems in the form of (3.3). The compositional framework is based on finite *dynamic Bayesian networks (DBNs)* and the results exploit the structure of the underlying Markov process to compute the abstraction separately for each dimension and discuss how factor graphs and the sum-product algorithm for DBNs can be utilized to solve the finite-horizon probabilistic invariance problem.

Compositional construction of *infinite* abstractions for interconnected dt-SCS is proposed by Lavaei et al. (2017) via sum-type small-gain conditions. The abstraction framework is based on notions of SSF in Definition 2.5, using which one can quantify the probabilistic distance between original interconnected stochastic control systems and their abstractions based on the closeness type in (3.6). A compositional scheme for constructing *infinite* abstractions based on dissipativity approaches is presented by Lavaei et al. (2019). The proposed scheme employs the interconnection matrix and joint dissipativity-type properties of subsystems and their abstractions described by a notion of stochastic storage functions (Lavaei et al. 2019, Definition 3.1).

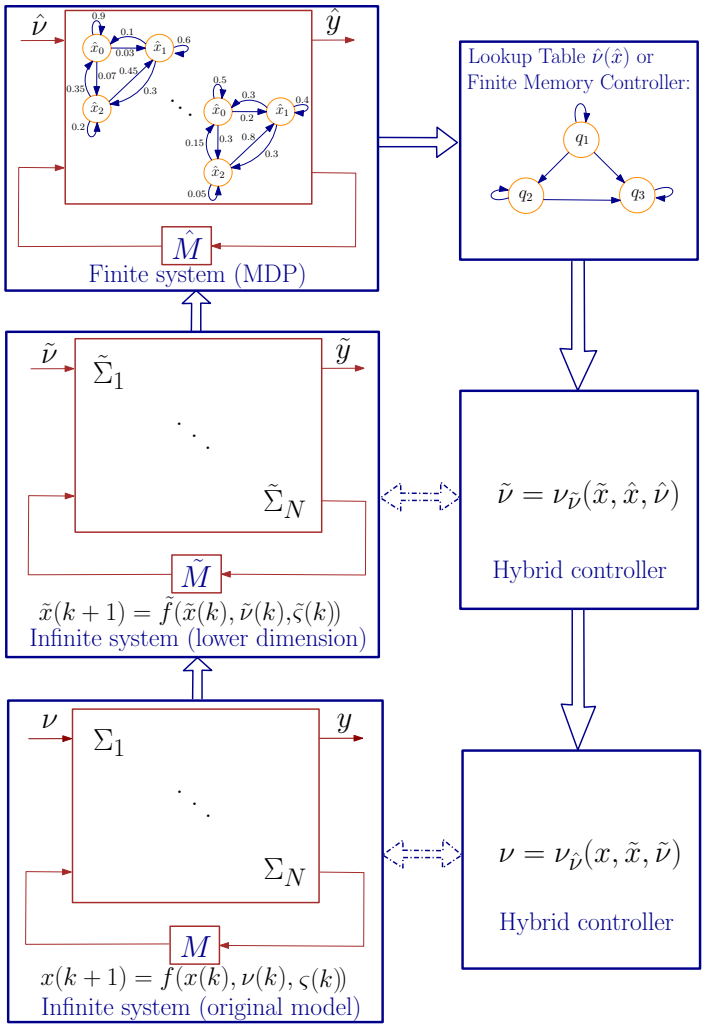


FIGURE 17. Compositional techniques based on (in)finite abstractions.

Compositional construction of both infinite and finite abstractions via max-type small-gain conditions is discussed by Lavaei, Soudjani & Zamani (2020*d*). The proposed overall error is computed based on maximum errors of subsystems. The article by Lavaei, Soudjani & Zamani (2020*d*) employs a variant of notions of SSF in Definition 2.5 and provides the probabilistic distance between the interconnection of stochastic control subsystems and that of their (in)finite abstractions based on (3.6). The proposed framework also leverages the $\delta$-ISS property of original systems as in Definition 2.4 and provides an approach to construct finite MDPs of the concrete models (or their reduced-order versions).

Compositional construction of finite abstractions for stochastic control systems is also presented by Lavaei et al. (2018) but using dissipativity conditions. This work provides a closeness in the form of (3.6) and proposes an

approach to construct finite MDPs for the general setting of discrete-time nonlinear SCS satisfying a passivity-like property, whereby one can construct finite MDPs by selecting a suitable discretization of the input and state sets. Moreover, for linear dt-SCS, the aforementioned property boils down to a matrix inequality.

Compositional construction of finite MDPs for networks of not necessarily stabilizable stochastic systems is presented by Lavaei, Soudjani & Zamani (2020*c*) via *relaxed* dissipativity approaches. The proposed framework relies on a relation between each subsystem and its finite abstraction employing a new notion of simulation functions, called *finite-step* stochastic simulation functions (Lavaei, Soudjani & Zamani 2020*c*, Definition A.4). In comparison with the existing notions of simulation functions in which stability or stabilizability of each subsystem is required, a *finite-step* stochastic simulation function needs to decay only after some finite numbers of steps instead of at each time step. This relaxation results in a *less conservative* version of small-gain or dissipativity conditions.

Compositional construction of finite abstractions for networks of classes of stochastic hybrid systems, namely, stochastic *switched* systems, is presented by Lavaei, Soudjani & Zamani (2020*a*), Lavaei & Zamani (2021) via respectively small-gain and dissipativity approaches. These contributions utilize notions of stochastic simulation (or storage) functions and provide a closeness guarantee in the form of (3.6), however adapted to switched models.

Compositional abstraction-based synthesis of dt-SCS using approximate probabilistic relations is proposed by Lavaei, Soudjani & Zamani (2020*b*). The abstraction framework is based on the notion of $\delta$-lifted relations, using which one can quantify the distance in probability between the interconnected dt-SCS and that of their abstractions as a version of closeness guarantee proposed in (3.3). Those results provide some matrix (in)equality conditions for simultaneous existence of relations incorporating the structure of the network. It is shown that the unified compositional scheme is less conservative than the two-step consecutive procedure that independently constructs infinite and finite abstractions (*e.g.,* Lavaei et al. (2018, 2019)).

Results on compositional multi-objective synthesis for finite probabilistic models are proposed by Kwiatkowska et al. (2013), and for continuous-space models by Haesaert et al. (2021) but with in a monolithic manner. However, those approaches are not, to the best of our knowledge, applicable or in general computationally tractable for large-scale SHS. Therefore, we raise the following open problem.

**Open Problem 8.** *Develop a compositional, multi-objective synthesis framework for continuous-space SHS, and display its applicability to classes of real-life large-scale problems.*

Although the proposed compositional frameworks for constructing finite abstractions can mitigate the state-space explosion problem, the curse of dimensionality may still arise at the level of single subsystems. As discussed in Section 6, an alternative direction is to employ *control barrier functions* as a discretization-free technique for controller synthesis of complex stochastic systems. However, searching for control barrier certificates for large-scale systems can be also computationally expensive. Consequently, developing compositional techniques for constructing control barrier functions is a promising solution to alleviate this complexity. Compositional construction of control barrier certificates for large-scale stochastic control systems is presented

by Anand et al. (2022) and Anand et al. (2021). The proposed compositional methodologies are based on a notion of *control sub-barrier certificates*, enabling one to construct control barrier certificates of interconnected systems by leveraging some max-type small-gain or dissipativity-type compositionality conditions, respectively. Compositional construction of control barrier certificates for discrete-time stochastic switched systems accepting multiple barrier certificates with some dwell-time conditions is also proposed by Nejati et al. (2020*a*).

The results in (Jahanshahi et al. 2022) also propose a compositional framework for the synthesis of safety controllers for networks of partially-observed discrete-time stochastic control systems (a.k.a., continuous-space POMDPs). The proposed framework is based on a notion of so-called local control barrier functions computed for subsystems in two different ways. In the first scheme, no prior knowledge of estimation accuracy is needed. The second framework utilizes a probability bound on the estimation accuracy using a notion of so called stochastic simulation functions. In both proposed schemes, sufficient small-gain type conditions are derived in order to compositionally construct control barrier functions for interconnected POMDPs using local barrier functions computed for subsystems. Leveraging compositionality results, the constructed control barrier functions enable computing lower bounds on the probabilities that the interconnected POMDPs avoid certain unsafe regions in finite-time horizons.

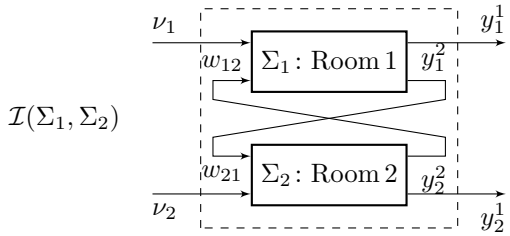


FIGURE 18. Interconnection of two rooms $\Sigma_1$ and $\Sigma_2$.

**Running example (continued)**. We present the dynamics of (2.4) based on a network of two rooms, as illustrated in Figure 18. The evolution of the temperatures $T_i$ can be described by (2.4) in which $A$ is a matrix with diagonal elements $\bar{a}_{ii} = (1 - 2\sigma - \theta - \gamma\nu_i(k))$, $i \in \{1, 2\}$, and off-diagonal elements $a_{1,2} = a_{2,1} = \sigma$. Parameter $\sigma = 0.1$ is the conduction factor between the rooms. Furthermore, $T(k) = [T_1(k); T_2(k)]$, $\nu(k) = [\nu_1(k); \nu_2(k)]$, $\varsigma(k) = [\varsigma_1(k); \varsigma_2(k)]$, $T_E = [T_{e_1}; T_{e_2}]$, and $R = 0.3\mathbb{I}_2$. By considering the individual rooms $\Sigma_i$ as

$$\Sigma_i : \begin{cases} T_i(k+1) = \bar{a}_{ii} T_i(k) + \gamma T_h \nu_i(k) + D_i w_i(k) + \theta T_{ei} + 0.3\varsigma_i(k), \\ y_i(k) = T_i(k), \end{cases} \tag{8.2}$$

one can readily verify that $\Sigma = \mathcal{I}(\Sigma_1, \Sigma_2)$ where $D_i = \sigma$, and $w_i(k) = y_{i-1}^2(k)$ for any $i \in \{1, 2\}$ (with $y_0^2 = y_2^2$). One can also establish a quadratic stochastic simulation function between $\Sigma_i$ and $\widehat{\Sigma}_i$ in the form of $S_i(T_i, \hat{T}_i) = (T_i - \hat{T}_i)^2$ satisfying (Lavaei, Soudjani & Zamani 2020*d*, conditions (III.1), (III.2)) with $\alpha_i(s) = s^2$, $\kappa_i(s) = 0.99s$, $\rho_{\text{int}i}(s) = 0.97s^2$, $\rho_{\text{ext}i}(s) = 0$, $\forall s \in \mathbb{R}_{\geq 0}$, and $\psi_i = 6.06\,\delta_i^2$, for any $i \in \{1, 2\}$.

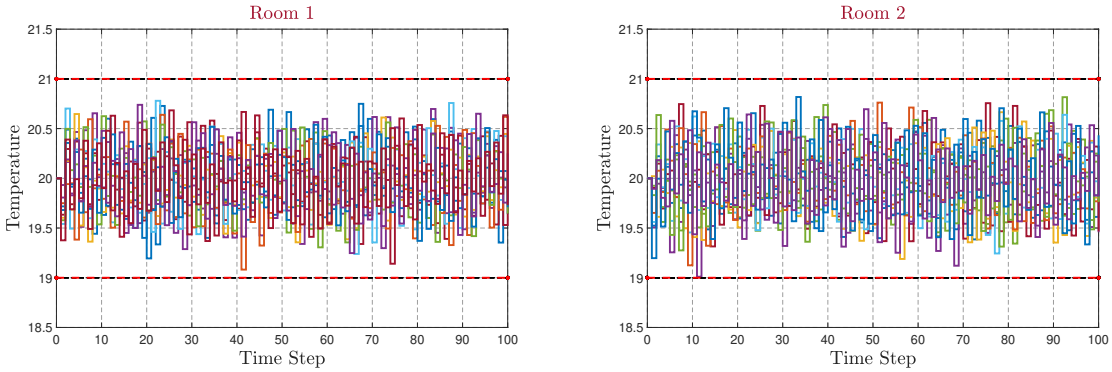


FIGURE 19. Closed-loop state trajectories of Room 1 (top) and Room 2 (bottom) with 10 different noise realizations for the finite-time horizon $T_d = 100$.

Now we need to check the small-gain condition for the interconnected system (2.4). The small-gain condition (Lavaei, Soudjani & Zamani 2020*d*, equation (V.2)) is readily satisfied if

$$\kappa_{12}.\kappa_{21} < 1, \tag{8.3}$$

where $\kappa_{ij}$ for any $i, j \in \{1, 2\}, i \neq j$, is defined as $\kappa_{ij}(s) = \rho_{\text{int}i}(\alpha_j^{-1}(s))$. Since $\kappa_{12} = \kappa_{21} = 0.97$, the small-gain condition (8.3) is simply satisfied. Hence, $V(T, \hat{T}) = \max_i (T_i - \hat{T}_i)^2$ for any $i \in \{1, 2\}$ is a stochastic simulation function from the interconnected system $\widehat{\Sigma}$ to $\Sigma$.[5]

By taking the *state* discretization parameter $\delta = 0.005$, the initial states of the interconnected systems $\Sigma$ and $\widehat{\Sigma}$ as $20\mathbb{1}_2$, and using (Lavaei, Soudjani & Zamani 2020*d*, ineqality (III.3)), one can guarantee that the distance between outputs of $\Sigma$ and $\widehat{\Sigma}$ will not exceed $\varepsilon = 0.5$ during the time horizon $T_d = 100$ with a probability of at least 98%, in the form of (3.6), *i.e.*,

$$\mathbb{P}\Big\{\|y(k) - \hat{y}(k)\| \leq 0.5,\ \forall k \in [0, 100]\Big\} \geq 0.98\,. \tag{8.4}$$

Let us now synthesize a controller for $\Sigma$ via its finite abstraction $\widehat{\Sigma}$ such that the controller maintains the temperature of any room in the comfort zone $[19, 21]$. We design a local controller for the abstract subsystem $\widehat{\Sigma}_i$, and then refine it back to the subsystem $\Sigma_i$ using an *interface* map. Consequently, the overall controller for the interconnected system $\Sigma$ would be a vector such that each of its components is the controller for subsystems $\Sigma_i$. We employ the software tool AMYTISS (Lavaei, Khaled, Soudjani & Zamani 2020) to synthesize controllers for $\Sigma_i$. Closed-loop state trajectories of two rooms with 10 different noise realizations are illustrated in Figure 19. The simulations show that none of 10 trajectories violates the specification, which is in accordance with the theoretical guarantee (8.4). As discussed in Section 5, if one employs our designed controllers and run Monte Carlo simulations of the closed-loop model, the distance between outputs of $\Sigma$ and $\widehat{\Sigma}$ will likely be empirically closer than 0.5 with the same probability as in (8.4). This is as expected, in view of the conservative nature of formal guarantees provided using Lyapunov-like techniques (simulation functions).

[5]We should highlight that condition (8.3) is similar to what proposed by Zames (1966) in the context of stability verification of feedback interconnection of two linear systems.

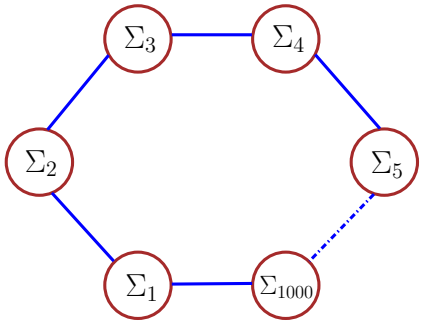


FIGURE 20. A circular interconnection for a network of 1000 rooms.

We now increase the number of rooms to $n = 1000$ and interconnect them in a circular fashion, as depicted in Figure 20. In this case, $A$ in (2.4) is a matrix with diagonal elements $\bar{a}_{ii} = (1-2\sigma-\theta-\gamma\nu_i(k))$, $i \in \{1,\dots,n\}$, off-diagonal elements $\bar{a}_{i,i+1} = \bar{a}_{i+1,i} = \bar{a}_{1,n} = \bar{a}_{n,1} = \sigma$, $i \in \{1,\dots,n-1\}$, and all other elements are identically zero. Moreover, $\sigma$ is a conduction factor between pairs of room $i \pm 1$ and $i$, $T(k) = [T_1(k);\dots;T_n(k)]$, $\nu(k) = [\nu_1(k);\dots;\nu_n(k)]$, $\varsigma(k) = [\varsigma_1(k);\dots;\varsigma_n(k)]$, $T_E = [T_{e1};\dots;T_{en}]$, $R = 0.3\mathbb{I}_n$. Considering the individual rooms $\Sigma_i$ as (8.2), one can readily verify that $\Sigma = \mathcal{I}(\Sigma_1,\dots,\Sigma_n)$ where $D_i = [\sigma;\sigma]^T$, and $w_i(k) = [y^2_{i-1}(k); y^2_{i+1}(k)]$ (with $y^2_0 = y^2_n$ and $y^2_{n+1} = y^2_1$).

We set the state discretization parameter $\delta = 0.005$, and initial states of the interconnected systems $\Sigma$ and $\widehat{\Sigma}$ as $20\mathbb{1}_{1000}$. Using the proposed bound by Lavaei, Soudjani & Zamani (2020*d*, ineqality (III.3)), one can guarantee that the distance between outputs of $\Sigma$ and $\widehat{\Sigma}$ will not exceed $\varepsilon = 0.5$ during the time horizon $T_d = 100$ with the probability at least 98%, *i.e.,*

$$\mathbb{P}\Big\{\|y(k) - \hat{y}(k)\| \leq 0.5,\ \forall k \in [0,100]\Big\} \geq 0.98. \tag{8.5}$$

We employ AMYTISS (Lavaei, Khaled, Soudjani & Zamani 2020) and synthesize a controller compositionally for $\Sigma$ via the abstraction $\widehat{\Sigma}$ such that the controller maintains the temperature of any room in the comfort zone $[19, 21]$. Closed-loop state trajectories of a representative room with 10 different noise realizations are illustrated in Figure 21 for a finite-time horizon $T_d = 100$. □

## 9. CONTINUOUS-TIME STOCHASTIC HYBRID SYSTEMS

Foundations of continuous-time SHS can be traced back to the work on piecewise-deterministic Markov models by Davis (1993), which was extended to diffusion processes by Hu et al. (2000). An early survey of work can be found in (Pola et al. 2003). In this survey, we present selected results for the continuous-time setting, categorized according to the different topics discussed in the previous sections.

**Notations.** We assume that for continuous-time processes, the triple $(\Omega, \mathcal{F}_\Omega, \mathbb{P}_\Omega)$ denotes a probability space endowed with a filtration $\mathbb{F} = (\mathcal{F}_s)_{s\geq 0}$ satisfying the standard conditions of completeness and right-continuity (Oksendal 2013). In addition, we denote by $(\mathbb{W}_s)_{s\geq 0}$ a $\mathsf{b}$-dimensional $\mathbb{F}$-Brownian motion. We now define continuous-time stochastic control systems which are studied in this section.

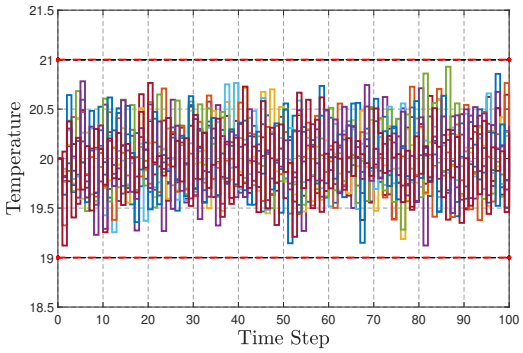


FIGURE 21. Closed-loop state trajectories of a representative room with 10 different noise realizations, for the network of 1000 rooms.

**Definition 9.1.** *A continuous-time stochastic control system (ct-SCS) is characterized by the tuple*

$$\Sigma = (X, U, \mathcal{U}, f, \sigma, Y, h), \tag{9.1}$$

*where:*

- $X \subseteq \mathbb{R}^n$ *is the state space of the system;*
- $U \subseteq \mathbb{R}^m$ *is the input space of the system;*
- $\mathcal{U}$ *is a subset of the sets of all* $\mathbb{F}$*-progressively measurable processes taking values in* $\mathbb{R}^m$*;*
- $f : X \times U \to X$ *is the drift term which is globally Lipschitz continuous: there exist constants* $\mathscr{L}_x, \mathscr{L}_\nu \in \mathbb{R}_{\geq 0}$ *such that* $\|f(x,\nu) - f(x',\nu')\| \leq \mathscr{L}_x\|x - x'\| + \mathscr{L}_\nu\|\nu - \nu'\|$ *for all* $x, x' \in X$*, and for all* $\nu, \nu' \in U$*;*
- $\sigma : \mathbb{R}^n \to \mathbb{R}^{n\times b}$ *is the diffusion term which is globally Lipschitz continuous with the Lipschitz constant* $\mathscr{L}_\sigma$*;*
- $Y \subseteq \mathbb{R}^q$ *is the output space of the system;*
- $h : X \to Y$ *is the output map.*

A continuous-time stochastic control system $\Sigma$ satisfies

$$\Sigma : \begin{cases} \mathsf{d}x(t) = f(x(t), \nu(t))\,\mathsf{d}t + \sigma(x(t))\,\mathsf{d}\mathbb{W}_t, \\ y(t) = h(x(t)), \end{cases} \tag{9.2}$$

$\mathbb{P}$-almost surely ($\mathbb{P}$-a.s.) for any $\nu \in \mathcal{U}$, where stochastic processes $x : \Omega \times \mathbb{R}_{\geq 0} \to X$ and $y : \Omega \times \mathbb{R}_{\geq 0} \to Y$ are called the *solution process* and the *output trajectory* of $\Sigma$, respectively. We also employ $x_{a\nu}(t)$ to denote the value of the solution process at time $t \in \mathbb{R}_{\geq 0}$ under an input trajectory $\nu$ from an initial condition $x_{a\nu}(0) = a$ $\mathbb{P}$-a.s., where $a$ is a random variable that is $\mathcal{F}_0$-measurable. We also denote by $y_{a\nu}$ the *output trajectory* corresponding to the *solution process* $x_{a\nu}$.

**Stochastic Similarity Relations.** Here, we first define the notion of stochastic simulation functions (SSF) for *continuous-time* stochastic control systems as a counterpart of Definition 2.5.

**Definition 9.2.** *Consider two ct-SCS* $\Sigma = (X, U, \mathcal{U}, f, \sigma, Y, h)$ *and* $\widehat{\Sigma} = (\hat{X}, \hat{U}, \hat{\mathcal{U}}, \hat{f}, \hat{\sigma}, \hat{Y}, \hat{h})$. *A twice-differentiable function* $V : X \times \hat{X} \to \mathbb{R}_{\geq 0}$ *is called a* stochastic simulation function *(SSF) from* $\widehat{\Sigma}$ *to* $\Sigma$ *if*

- $\exists \alpha \in \mathcal{K}_\infty$ *such that*

$$\forall x \in X, \forall \hat{x} \in \hat{X}, \quad \alpha(\|h(x) - \hat{h}(\hat{x})\|^{\bar{q}}) \leq V(x, \hat{x}),$$

- $\forall x \in X, \forall \hat{x} \in \hat{X}, \forall \hat{\nu} \in \hat{U}$, $\exists \nu \in U$ *such that*

$$\mathcal{L}V(x, \hat{x}) \leq -\kappa(V(x, \hat{x})) + \rho_{\text{ext}}(\|\hat{\nu}\|^{\bar{q}}) + \psi,$$

*for some* $\kappa \in \mathcal{K}_\infty$, $\rho_{\text{ext}} \in \mathcal{K}_\infty \cup \{0\}$, *and* $\psi \in \mathbb{R}_{\geq 0}$, *where* $\bar{q} \in \mathbb{N}_{\geq 1}$ *denoting the moment of a random variable and* $\mathcal{L}V$ *is the infinitesimal generator of the stochastic process acting on the function* $V$ *(Oksendal 2013), defined as*

$$\mathcal{L}V(x, \hat{x}) = \partial_x V f(x, \nu) + \partial_{\hat{x}} V \hat{f}(\hat{x}, \hat{\nu}) + \frac{1}{2}\mathsf{Tr}(\sigma(x)\sigma(x)^T \partial_{x,x} V) + \frac{1}{2}\mathsf{Tr}(\hat{\sigma}(\hat{x})\hat{\sigma}(\hat{x})^T \partial_{\hat{x},\hat{x}} V).$$

In the next theorem, we present a result on the closeness in *expectation (moment)* of the difference between output trajectories of original continuous-time systems $\Sigma$ and their corresponding abstractions $\widehat{\Sigma}$, as proposed by Zamani, Mohajerin Esfahani, Majumdar, Abate & Lygeros (2014). This relates to *condition (iii)* in Definition 1.2, which is the last of the four closeness results that have been introduced in this article.

**Theorem 9.3.** *Let* $\Sigma$ *be a continuous-time SCS and* $\widehat{\Sigma}$ *be its abstraction. Suppose there exists a stochastic simulation function* $V : X \times \hat{X} \to \mathbb{R}_{\geq 0}$ *from* $\widehat{\Sigma}$ *to* $\Sigma$ *as in Definition 9.2. For any input trajectory* $\hat{\nu}(\cdot) \in \hat{\mathbb{U}}$ *that preserves the Markov property for the closed-loop* $\widehat{\Sigma}$, *and for any random variables* $a$ *and* $\hat{a}$ *as the initial states of* $\Sigma$ *and* $\widehat{\Sigma}$, *respectively, one can construct an input trajectory* $\nu(\cdot) \in \mathbb{U}$ *for* $\Sigma$ *through an interface function associated with* $V$ *(cf. Def. 2.5), such that:*

$$\mathbb{E}\Big[\|y_{a\nu}(t) - \hat{y}_{\hat{a}\hat{\nu}}(t)\|^{\bar{q}}\Big] \leq \lambda_4, \quad \forall t \in \mathbb{R}_{>0}, \tag{9.3}$$

*where,*

$$\lambda_4 := \beta(\mathbb{E}[V(a, \hat{a})], t) + \rho_{\text{ext}}(\mathbb{E}[\|\hat{\nu}\|_\infty^{\bar{q}}]) + c,$$

*with* $\beta \in \mathcal{KL}, \rho_{\text{ext}} \in \mathcal{K}_\infty, c \in \mathbb{R}_{>0}$, *and* $\bar{q} \in \mathbb{N}_{\geq 1}$ *denoting the moment of a random variable.*

**Remark 9.4.** *Note that one can leverage the bound in* (9.3) *together with the Markov inequality (Oksendal 2013) to provide a lower bound on the probability of satisfaction of logic specifications for which satisfiability is only concerned at single time instances (e.g., reachability). The new bound is similar to* (3.6)*, but the supremum appears outside of the probability operator (Zamani, Mohajerin Esfahani, Majumdar, Abate & Lygeros 2014). As such, note that this reasoning is not useful for providing probabilistic bounds on satisfaction of general logic specifications, such as those involving the always (*$\Box$*) or until (*$\mathsf{U}$*) operators, since it does not capture the joint distributions across different time instants. Then, one can conclude that the bound in* (9.3) *implies* (3.6) *and accordingly* (3.3) *and* (3.7)*.*

The characterization and computation of probabilistic bisimulations of SCS are discussed by Abate (2009). This work proposes sufficient conditions for the existence of a stochastic simulation function based on the use of contractivity analysis (a variant of incremental stability) for probabilistic systems. The results are then extended by Abate (2010) to probabilistic simulations between two SCS that are additionally endowed with switching and resetting behaviors. A notion of stochastic simulation function for continuous-time stochastic systems is discussed by Zamani, Rungger & Mohajerin Esfahani (2017, Definition 3.2).

**Infinite abstractions.** The construction of infinite abstractions for a class of continuous-time SHS was initially proposed by Julius & Pappas (2009). The approximation framework is based on stochastic simulation functions and the work provides a closeness guarantee in the form of continuous-time counterpart of (3.6) but for infinite-time horizons (*i.e.,* $0 \leq k < \infty$). For the class of jump linear stochastic systems and linear stochastic hybrid automata, the article shows that the computation of stochastic simulation functions can be cast as a linear matrix inequality (LMI) problem. A method for verifying continuous-time SHS using the Mori-Zwanzig model reduction is proposed by Wang et al. (2016), where properties are specified as Metric Interval Temporal Logic (MITL) formulas. MITL is an extension of LTL that deals with models with dense time. By partitioning the state space of the continuous-time SHS and computing the optimal transition rates between partitions, the work provides a procedure to both reduce a continuous-time SHS to a continuous-time Markov chain (CTMC), and the associated MITL formulae defined on the continuous-time SHS to MITL specifications on the CTMC.

**Finite abstractions.** Construction of symbolic models (*i.e.,* finite abstractions) for incrementally stable stochastic control and switched systems is proposed by Zamani, Mohajerin Esfahani, Majumdar, Abate & Lygeros (2014) and Zamani et al. (2015), respectively. The underlying switched systems in (Zamani et al. 2015) have a probabilistic evolution over a continuous domain and control-dependent discrete dynamics over a finite set of modes. Both papers constructively derive approximately equivalent (bisimilar) symbolic models of stochastic systems. The result in (Zamani et al. 2015) provides two different symbolic abstraction techniques: one requires state space discretization, but the other one does not require any space discretization which can be potentially more efficient than the first one especially when dealing with higher dimensional stochastic switched systems. Construction of symbolic models for *randomly switched* stochastic system is studied by Zamani & Abate (2014*a*). The proposed framework is based on approximate bisimilar relations and leverages some incremental stability assumption over randomly switched stochastic systems to establish the relation between original systems and their corresponding finite symbolic models. All those aforementioned results provide a closeness bound between the original system and its bisimilar finite abstraction based on the one proposed in (9.3).

**Control barrier certificates.** Discretization-free approaches based on barrier certificates for safety verification of continuous-time SHS are initially proposed by Prajna et al. (2007). The article leverages the supermartingale property and quantifies an upper bound on the probability that a trajectory of the system ever reaches a given unsafe set (over an infinite time horizon) as proposed in (6.7). For polynomial-type systems, barrier certificates can be constructed using convex optimization, which is computationally tractable.

Verification and control for finite-time safety of stochastic systems via barrier functions are discussed by Santoyo et al. (2019). The proposed certificate condition includes a state-dependent bound on the infinitesimal generator, allowing for tighter probability bounds. Moreover, for stochastic systems where the drift dynamics are affine in the control input, the paper proposes a method for synthesizing a polynomial state-feedback controller that achieves a specified safety probability.

Control barrier functions for stochastic systems under process and measurement noise have been proposed by Clark (2021). The article first considers the case where the system state is known at each time step, and presents a construction that guarantees almost sure safety. It then extends the approach to models with incomplete state information, where the state must be estimated: it is shown that the proposed certificates ensure safety with probability 1 when the state estimate is within a given bound of the true state, which can be achieved using an Extended Kalman Filter[6] when the system is linear or the process and measurement noises are sufficiently small.

Synthesis for stochastic systems with partial state information via control barrier functions is discussed by Jahanshahi et al. (2020*b*). Given an estimator with a probabilistic guarantee on its accuracy, the paper proposes an approach to compute a controller providing a lower bound on the probability that the trajectories of the stochastic control system remain safe over a finite time horizon (similar to (6.6)). This work does not require a supermartingale property on the control barrier functions, and in particular it does not require any stability assumption on the model. The results of this article are generalized by Jahanshahi et al. (2020*a*) in which no a-priori knowledge about the estimation accuracy is needed. Besides, the class of properties is extended to those expressed by nondeterministic finite automata (NFA), and the dynamics are also generalized to partially-observed jump-diffusion systems.

**Compositional techniques.** Compositional construction of infinite abstractions (in particular, reduced-order models) for a class of SHS is proposed by Zamani, Rungger & Mohajerin Esfahani (2017) using sum-type small-gain conditions. The class of systems includes both jump linear stochastic systems and linear stochastic hybrid automata. The work employs stochastic simulation functions to quantify an error between the interconnection of stochastic hybrid subsystems and that of their approximations, in the form of (9.3). It also focuses on a specific class of SHS, namely jump linear stochastic systems, and proposes a constructive scheme to determine approximations together with their corresponding stochastic simulation functions.

Compositional construction of finite abstractions for stochastic control systems is presented by Mallik et al. (2017). The proposed framework is based on a notion of (approximate) disturbance bisimulation relation, which results in a closeness guarantee in the form of (9.3). Given any SCS satisfying a stochastic version of a $\delta$-ISS property and a positive error bound, the article shows how to construct a finite-state transition system (whenever existing) which is disturbance-bisimilar to the given stochastic control system.

Compositional construction of finite MDPs for networks of continuous-time stochastic systems via max small-gain conditions is proposed by Nejati et al. (2021). The proposed framework leverages stochastic simulation functions to relate continuous-time stochastic systems with their discrete-time (in)finite counterparts. In order

[6]This and other particle filters have been developed for SHS by Blom & Bloem (2004, 2007).

to propose the construction procedure for finite abstractions, the paper first introduces infinite abstractions as time-discretized versions of original continuous-time stochastic hybrid systems (as a middle step) since the finite abstractions are constructed from the discrete-time counterparts. The work quantifies the distance in probability between original continuous-time stochastic hybrid systems and their discrete-time (finite or infinite) abstractions at sampling times in the form of (3.6). It also constructs finite abstractions together with their corresponding stochastic simulation functions for a particular class of *nonlinear* SHS. Although the original models in (Zamani et al. 2015, Zamani & Abate 2014*a*, Mallik et al. 2017) are stochastic, the constructed abstractions are finite, non-stochastic labeled transition systems. However, the finite abstractions in (Nejati et al. 2021) are finite MDPs and potentially less conservative for the sake of controller synthesis and satisfaction probabilities.

Compositional construction of infinite abstractions for networks of stochastic hybrid systems under randomly switched topologies are proposed by Awan & Zamani (2018). The proposed framework leverages the interconnection topology, switching randomly between $\mathcal{P}$ different interconnection topologies (it is modelled by a Markov chain), and the joint dissipativity-type properties of subsystems and their abstractions. The abstraction itself is a stochastic hybrid system (possibly with a lower dimension) and can be used as a substitute of the original system in the controller design process. The work provides a closeness guarantee based on different moments similar to (9.3).

Compositional construction of infinite abstractions of interconnected stochastic hybrid systems via dissipativity theory is discussed by Awan & Zamani (2019). The proposed results leverage a notion of stochastic simulation function in which the supply rate has its own dynamics. The stochastic noises and jumps in the concrete subsystem and its abstraction do not need to be the same. For a class of nonlinear stochastic hybrid subsystems with an incremental quadratic inequality on the nonlinearity, a set of matrix (in)equalities is established to facilitate the construction of their abstractions together with the corresponding stochastic storage functions. The article quantifies a formal error between the output behaviors of the original system and the ones of its infinite abstractions in the form of (9.3).

Compositional construction of control barrier functions for networks of continuous-time stochastic systems is presented by Nejati et al. (2020*b*). The proposed scheme is based on notions of *pseudo-barrier functions* (similar to Definition 6.1 but computed over subsystems), using which one can synthesize state-feedback controllers for interconnected systems enforcing safety specifications over a finite time horizon. This work leverages sum-type small-gain conditions to compositionally construct control barrier functions for interconnected systems based on the corresponding pseudo-barrier functions computed for subsystems. Then, using the constructed control barrier functions, it quantifies upper bounds on the probability that an interconnected system reaches certain *unsafe* regions in a finite time horizon, similar to Theorem 6.3 but in the continuous-time setting. The work also employs a systematic technique based on the SOS optimization program to search for pseudo-barrier functions of subsystems, while synthesizing safe controllers.

**Temporal logic verification and synthesis.** Early work on SHS is by Gao et al. (2006) and Koutsoukos & Riley (2006), which focused on probabilistic reachability analysis. Stochastic reachability analysis of hybrid

systems is also studied by Bujorianu & Lygeros (2003), Bujorianu (2012), and this line of work continues to these days (Wisniewski et al. 2020, Cosentino et al. 2021).

A reachability analysis problem for an aircraft conflict prediction, modelled as a stochastic hybrid system, is studied by Prandini & Hu (2008) in which a switching diffusion is presented to predict the future positions of an aircraft given a flight plan. The work proposes a numerical algorithm for estimating the probability that the aircraft either enters an unsafe region or closely approaches to another aircraft.

A probabilistic approach for control of continuous-time linear stochastic systems subject to LTL formulae over a set of linear predicates in the state of the system is presented by Lahijanian et al. (2009). The article defines a polyhedral partition of the state space and a finite collection of controllers, represented as symbols, and constructs a finite MDP. By utilizing an algorithm resembling LTL model checking, it determines a run satisfying the formula in a corresponding Kripke structure. A sequence of control actions in the MDP is determined to maximize the probability of following the run.

Measurability and safety verification of a class of SHS are discussed by Fränzle et al. (2011) in which the continuous-time behaviour is given by differential equations, but discrete jumps are chosen by probability distributions. In this work, non-determinism is also supported, and it is exploited in an abstraction and evaluation method that establishes safe upper bounds on reachability probabilities.

An optimal control problem for continuous-time stochastic systems subject to objectives specified in a fragment of metric interval temporal logic specifications, a temporal logic with real-time constraints, is presented by Fu & Topcu (2015). The work proposes a numerical method for computing an optimal policy with which the given specification is satisfied with maximal probability in the discrete approximation of the underlying stochastic system. It is shown that the policy obtained in the discrete approximation converges to the optimal one for satisfying the specification in the continuous or dense-time semantics, as the discretization becomes finer in both state and time.

Motion planning for continuous-time stochastic processes via dynamic programming is discussed by Mohajerin Esfahani et al. (2015). The work studies stochastic motion planning problems which involve a controlled process, with possibly discontinuous sample paths, visiting certain subsets of the state-space while avoiding others in a sequential fashion. A weak dynamic programming principle (DPP) is proposed that characterizes the set of initial states which admit a control enabling the process to execute the desired maneuver with probability no less than some pre-specified value. The proposed DPP comprises auxiliary value functions defined in terms of discontinuous payoff functions.

Reachability analysis for continuous-time stochastic hybrid systems with no resets is proposed by Laurenti et al. (2017) in which continuous dynamics described by linear stochastic differential equations. For this class of models, the article studies reachability (and dually, safety) properties on an abstraction defined in terms of a discrete-time and finite-space Markov chain, with provable error bounds. The paper provides a characterization of the uniform convergence of the time discretization of stochastic processes with respect to safety properties, and this allows to provide a complete and sound numerical procedure for reachability and safety computation over the stochastic systems.

An approach for the automated synthesis of safe and robust proportional-integral-derivative (PID) controllers for SHS is proposed by Shmarov et al. (2017). The work considers hybrid systems with nonlinear dynamics (Lipschitz continuous ordinary differential equations) and random parameters, and synthesizes PID controllers such that the resulting closed-loop systems satisfy safety and performance constraints given as probabilistic bounded-reachability properties. The proposed technique leverages SMT solvers over reals and nonlinear differential equations to provide formal guarantees that the synthesized controllers satisfy such properties. These controllers are also robust by design since they minimize the probability of reaching an unsafe set in the presence of random disturbances.

Automated synthesis of controllers with formal safety guarantees for nonlinear control systems with noisy output measurements, and stochastic disturbances is presented by Shmarov et al. (2020). The proposed method derives controllers such that the corresponding closed-loop system, modeled as a sampled-data stochastic control system, satisfies a safety specification with probability above a given threshold. If the obtained probability is not above the threshold, the approach expands the search space for candidates by increasing the controller degree.

A theoretical and computational synthesis framework for safety properties of continuous-time continuous-space switched diffusions is proposed by Laurenti et al. (2020). The work provides an appropriate discrete abstraction in the form of an uncertain Markov model that captures all possible behaviors of the system. This is achieved through a discretization of both time and space domains, each introducing an error such that the errors are formally characterized and represented as uncertain transition probabilities in the abstraction model. It also provides a robust strategy that optimizes a safety property over the abstraction. This strategy is computed by considering only the feasible transition probability distributions, preventing the explosion of the error term and resulting in achievable bounds for the safety probability. This robust strategy is mapped to a switching strategy for the system consisting of switched diffusions with the guarantee that safety probability bounds also hold for this system in the form of (3.7).

**Stability and optimal control.** In this survey we do not delve into the broad issues of stability analysis and optimal control synthesis for SHS, which have been widely investigated over the past two decades. For stability, we point the reader to the survey in (Teel et al. 2014), whereas for optimal control we refer the interested reader to the books (Blom et al. 2006, Cassandras & Lygeros 2006) (which covers seminal work (Davis 1993, Arapostathis et al. 1993)) and to the work of (Pakniyat & Caines 2016). It is worth mentioning that formal verification and synthesis have a tight connection to *optimal control* by expressing logic specifications as a part of constraints in the optimization problem (Lesser & Abate 2018, Haesaert et al. 2021). Finally, let us mention the qualitative connection with stochastic MPC (Mesbah 2016), discussed in more detail elsewhere in this survey.

## 10. Simulations and Statistical Model Checking of SHS

In this section, we study simulation-based analysis of stochastic hybrid systems, and also encompass work on statistical model checking (SMC). Early work deals with the development of filtering algorithms that are

applicable to SHS (Blom & Bloem 2004, 2007). Similarly, sequential Monte Carlo simulations and the use of Petri Nets are discussed by Blom et al. (2007), and employed in collision risk estimation of free flight operations. The proposed results are applicable to rare-event estimation over complex models. Safety risk analysis of an air traffic management operation over a SHS is similarly discussed by Blom et al. (2013). (Bouissou et al. 2014) puts forward techniques for efficient Monte Carlo simulation of stochastic hybrid systems.

So far, we have widely discussed formal verification and synthesis of SHS in which a closeness guarantee between original SHS and their abstractions is formally provided in four different forms, as elaborated in Definition 1.2. For simulation-based analysis instead, new guarantees are provided. Suppose $(\hat{x}_i)_{i=1}^{\bar{N}}$ are $\bar{N}$ i.i.d. sampled data from a set $\Omega$. Simulation-based guarantees are presented in two-layer probabilities (Calafiore & Campi 2006), as follows:

$$\mathbb{P}\Big\{(\hat{x}_i)_{i=1}^{\bar{N}} \in \Omega \colon \mathbb{P}\big\{\bar{y}_{a\nu} \not\models \varphi\big\} \leq \bar{\varepsilon}\Big\} \geq \bar{\beta}, \tag{10.1}$$

where $\varphi$ is the property of interest, $\bar{y}_{a\nu}$ is any given random output trace, and $\bar{\varepsilon}, \bar{\beta} \in [0,1]$ are respectively a threshold and confidence level. As a comparison with the studied approaches in the previous sections, the formal guarantee there comes with only one layer probability similar to (6.4). If the confidence level $\bar{\beta}$ is increased to one, the chance constrained problem (10.1) can be understood to be similar to (6.4).

A statistical model checking (SMC) algorithm to verify stochastic properties with unbounded until is presented by Sen et al. (2005). The algorithm is based on Monte Carlo simulation of the model and hypothesis testing of the samples, as opposed to sequential hypothesis testing. Statistical model checking for synthesizing policies on stochastic models including finite MDPs is presented by Henriques et al. (2012). The proposed framework develops an algorithm that resolves nondeterminism probabilistically, and then uses multiple rounds of sampling and reinforcement learning to provably improve resolutions of nondeterminism with respect to satisfying a bounded linear temporal logic (BLTL) property. The proposed algorithm thus reduces an MDP to a fully probabilistic Markov chain on which SMC may be applied to give an approximate estimation of the probability of the BLTL property.

A numerically rigorous Monte Carlo approach for computing probabilistic reachability in hybrid systems subject to random and nondeterministic parameters is proposed by Shmarov & Zuliani (2016). Instead of standard simulation, the work employs $\delta$-complete SMT procedures, which enables formal reasoning for nonlinear systems up to a user-definable numerical precision. Monte Carlo approaches for probability estimation assume that sampling is possible for the real system at hand, however when using $\delta$-complete simulations, one instead samples from an over-approximation of the random quantities at hand. The article introduces a Monte Carlo-SMT approach for computing probabilistic reachability confidence intervals that are both statistically and numerically rigorous. A survey on statistical model checking is provided by Agha & Palmskog (2018), which covers SMC algorithms, techniques, and tools, while emphasizing limitations and tradeoffs between precision and scalability.

A multilevel Monte Carlo method for statistical model checking of continuous-time stochastic hybrid systems is proposed by Soudjani, Majumdar & Nagapetyan (2017). The provided approach relies on a sequence of discrete-time stochastic processes whose executions approximate and converge weakly to that of the original continuous-time SHS with respect to the satisfaction of a property of interest. With focus on bounded-horizon

reachability, the paper casts the model checking problem as the computation of the distribution of an exit time, which is in turn formulated as the expectation of an indicator function. This latter computation involves estimating discontinuous functionals, which reduces the bound on the convergence rate of the Monte Carlo algorithm. The work then proposes a smoothing step with tuneable precision and formally quantifies the error in the mean-square sense, which is composed of smoothing error, bias, and variance.

A confidence bound for statistical model checking of probabilistic hybrid systems is proposed by Ellen et al. (2012). The work presents an approximation algorithm based on confidence intervals obtained from sampling which allow for an explicit trade-off between accuracy and computational effort. Although the algorithm gives only approximate results in terms of confidence intervals, it is still guaranteed to converge to the exact solution. A Bayesian approach to statistical model-checking of discrete-time Markov chains with respect to continuous stochastic logic specifications is presented by Lal et al. (2020). Related to this approach, (Molyneux & Abate 2020) formally integrates the use of sequential Monte Carlo techniques for approximate Bayesian inference with (Bayesian) statistical model checking.

In conclusion, statistical model checking approaches appear to be suitable for verification goals, whereas they have shown to be less efficient when synthesis is in order. The latter objective could be considered as an open problem for future research. In addition, most of the proposed SMC results are suitable for *finite-time* horizons. More precisely, in the setting of SMC approaches in *infinite-time* horizons, the proposed results require some strong assumptions that are not in general satisfiable by SHS. More emphasis on infinite-horizon properties via SMC can be cast as another future research direction.

## 11. Software Tools

In this section, we discuss software tools for verification and synthesis, as well as simulation, of stochastic hybrid systems. This is a growing and fast-pacing area, thus we focus on existing tools at the time of writing, particularly those that have participated in the *ARCH Initiative* (discussed below): we emphasize their architectures and relate them to the underlying theory, and conclude presenting a relevant open-science initiative in this area.

11.1. **The Modest Toolset.** Modest Toolset (Hartmanns & Hermanns 2014) performs modelling and analysis for hybrid, real-time, distributed and stochastic systems. At its core are models of networks of stochastic hybrid automata (SHA), which combine nondeterministic choices, continuous system dynamics, stochastic decisions and timing, and real-time behaviour, including nondeterministic delays. The Modest Toolset is a modular framework, supporting as input the high-level Modest modelling language and providing a variety of analysis backends for various special cases of SHA. Many existing automata-based formalisms are special cases of SHA.

11.2. **SReach.** SReach (Wang et al. 2015) solves probabilistic, bounded-time reachability problems for two classes of models: (i) nonlinear hybrid automata with parametric uncertainty, and (ii) probabilistic hybrid automata with additional randomness on both transition probabilities and variable resets. Standard approaches

to reachability analysis for linear hybrid systems require numerical solutions of large optimization problems, which become practically infeasible for systems involving both nonlinear dynamics and stochasticity. SReach instead encodes stochasticity by using a set of random variables, and combines $\delta$-complete decision procedures and statistical tests to solve $\delta$-reachability problems. Compared to standard simulation-based methods, SReach supports non-deterministic branching and allows one to increase the coverage of performed simulations.

11.3. **ProbReach.** ProbReach (Shmarov & Zuliani 2015) is a statistical model checking tool that studies bounded-time reachability and other quantitative properties. It handles SHS with random continuous quantities encompassing model parameters or initial conditions that are chosen within an initial set and which remain unchanged throughout the system evolution. For continuous dynamics, ProbReach can analyze any Lipschitz-continuous differential equations with stochastic parameters. Given an SHS with random continuous quantities and an arbitrarily small $\epsilon > 0$, ProbReach returns an interval of size not larger than $\epsilon$ containing the exact bounded-reachability probability. This result is guaranteed to be numerically sound, *e.g.,* free from floating-point inaccuracies. The introduction of discrete random parameters to the system will not affect the guarantees provided by ProbReach, however if the model features only discrete random parameters, then these guarantees do not hold: this happens because probability distributions over discrete random parameters are not continuous, hence an arbitrary precision cannot be provided any longer. Introducing nondeterministic continuous parameters affects the guarantees the tool provides, as well: this happens because nondeterministic parameters do not have any probability measure. In this case, ProbReach computes an enclosure that is guaranteed to contain all the possible reachability probabilities. In general, such an enclosure may have size larger than $\epsilon > 0$. ProbReach employs a validated integration procedure to obtain a partition over the random continuous quantities in such a way that the guarantees described above hold. This partition is then used to enclose the probabilistic outcome by computing under- and over-approximations.

11.4. **SReachTools.** SReachTools (Vinod et al. 2019) is an open-source Matlab toolbox for performing stochastic reachability of linear, potentially time-varying, discrete-time systems that are perturbed by a stochastic disturbance. More precisely, this tool addresses the problem of stochastic reachability of a target tube, which also encompasses terminal-time (hitting) problems, reach-avoid problems, and related viability problems (not discussed in this survey). The stochastic reachability of a target tube problem maximizes the likelihood that the state of a stochastic system will remain within a collection of time-dependent target sets for a given time horizon, while respecting system dynamics and utilizing inputs within a bounded control domain. SReachTools implements several algorithms based on convex optimization, computational geometry, and Fourier transforms, to efficiently compute over- and under-approximations of stochastic reach sets. SReachTools can be employed to perform probabilistic verification of closed-loop systems, and can also perform controller synthesis via open-loop or affine state-feedback controllers.

11.5. **HYPEG.** A statistical simulator for hybrid Petri nets with general transitions, called HYPEG, is presented by Pilch et al. (2017). It combines discrete and continuous components with a possibly large number of random variables, whose stochastic behavior follows arbitrary probability distributions. HYPEG employs

time-bounded discrete-event simulation and well-known statistical model checking techniques to verify properties, including time-bounded reachability.

11.6. **Mascot-SDS.** Mascot-SDS (Majumdar et al. 2020) is an open-source tool for synthesizing controllers with formal correctness guarantees for discrete-time dynamical systems in the presence of stochastic perturbations. Mascot-SDS is written in C++, and is an extension of Mascot (Hsu et al. 2018). The tool supports *infinite-horizon* control specifications for stochastic dynamical systems and computes over- and under-approximations of the set of states that satisfy a given specification with probability one. The current version of the tool is developed for "always eventually" specifications, namely for specifications dealing with "infinitely often" ($\omega$-regular) requirements.

11.7. **Level-Set Toolbox.** The Level-Set Toolbox (Mitchell 2007) is a software package for solving time-dependent Hamilton-Jacobi partial differential equations (PDEs) in the Matlab programming environment. Level set methods are often used for simulation of dynamic implicit surfaces in graphics, fluid and combustion simulations, image processing, and computer vision. Hamilton-Jacobi and related PDEs arise in fields such as control, robotics, differential games, dynamic programming, mesh generation, stochastic differential equations, financial mathematics, and verification. All source code for the toolbox is provided as plain text in the Matlab m-file programming language. The toolbox is designed to allow quick and easy experimentation with level set methods, although it is not by itself a level set tutorial and so should be used in combination with the existing literature. The Level-Set Toolbox has been in particular used for the analysis of stochastic models in (Park et al. 2014, Sprinkle et al. 2005, Ding 2012, Choi et al. 2022).

11.8. **FAUST $^2$.** FAUST$^2$ (Soudjani, Gevaerts & Abate 2015) generates formal abstractions for continuous-space discrete-time Markov processes defined over uncountable (continuous) state spaces, and performs verification and synthesis for safety and reachability specifications. The abstract model is formally put in a relationship with the concrete model via a user-defined maximum threshold on the approximation error introduced by the abstraction procedure. FAUST$^2$ allows exporting the abstract model to well-known probabilistic model checkers, such as PRISM (Kwiatkowska et al. 2002) or Storm (Dehnert et al. 2017). Alternatively, it can handle internally the computation of PCTL properties (*e.g.,* safety or reachability) over the abstract model. It also allows refining the outcomes of the verification procedures over the concrete model in view of the quantified and tuneable error, which depends on the concrete dynamics and on the given PCTL formula.

11.9. **StocHy.** StocHy (Cauchi, Degiorgio & Abate 2019) performs quantitative analysis of discrete-time stochastic hybrid systems. The tool allows to (i) simulate the SHS evolution over a given time horizon; and to automatically construct finite abstractions of the SHS. Abstractions are then employed for (ii) formal verification or (iii) control synthesis satisfying safety and reachability specifications. The tool is implemented in C++ and employs manipulations based on vector calculus, using sparse matrices, the symbolic construction of probabilistic kernels, and multi-threading. StocHy allows for modular modelling, and has separate simulation, verification and synthesis engines which are implemented as independent libraries. This allows for libraries to be readily used and for extensions to be easily built.

11.10. **AMYTISS.** AMYTISS (Lavaei, Khaled, Soudjani & Zamani 2020) is developed in C++/OpenCL for designing correct-by-construction controllers of large-scale discrete-time stochastic control systems. AMYTISS natively supports both additive and multiplicative noises with different distributions including normal, uniform, exponential, and beta. This software tool provides scalable parallel algorithms that allow to (i) construct finite MDPs from discrete-time stochastic control systems, and (ii) synthesize controllers satisfying complex logic properties including safety, reachability, and reach-avoid specifications. AMYTISS employs high-performance computing platforms and cloud-computing services to alleviate the effects of the state-explosion problem. This tool improves performances over computation time and memory usage by parallel execution over different heterogeneous computing platforms including CPUs, GPUs and hardware accelerators (*e.g.,* FPGAs). AMYTISS significantly reduces the memory usage by setting a probability threshold $\gamma \in [0, 1]$ to control how many partition elements around the mean of the system should be stored. Such an approximation allows controlling the sparsity of the columns of $\hat{T}_{\mathsf{x}}$ (transition probability matrix of constructed finite MDP). AMYTISS also proposes another technique that further reduces the required memory for computing $\hat{T}_{\mathsf{x}}$, named *on-the-fly abstraction* (OFA). In OFA, computing and storing the probability transition matrix $\hat{T}_{\mathsf{x}}$ are skipped. Instead the required entries of $\hat{T}_{\mathsf{x}}$ on-the-fly are computed as they are needed for the synthesis part via the standard dynamic programming. This reduces the required memory for $\hat{T}_{\mathsf{x}}$ but at the cost of repeated computation of their entries in each time step from 1 to a finite-time horizon $T_d$. AMYTISS has been successfully applied to some large-scale applications including autonomous vehicles.

11.11. **The ARCH Initiative.** The ARCH competition aims at providing an updated point of reference on the current state of the art in the area of models for hybrid systems, together with the currently available tools and frameworks for performing formal verification and optimal policy synthesis. The initiative further provides a set of benchmarks aiming to push forward the development of current and future tools. To provide a fair and comprehensive comparison of results, which also allows tools designed for multi-core architectures to highlight their capabilities, the competition is performed via a centralized execution of the benchmarks. To establish further trustworthiness of the results, and to bolster related *open science* initiatives, the code describing the benchmarks together with the code used to compute the results are also published in a public server. The tools compete based on different aspects including implementation languages, class of models, platforms, algorithms, specifications, type of stochasticity, type of distributions, type of disturbances, etc. Presentation and discussion of outcomes of yearly benchmarking competitions on tools for formal verification and policy synthesis of stochastic models (and in particular SHS) are provided by Abate et al. (2018, 2019, 2020), Abate, Blom, Bouissou, Cauchi, Chraibi, Delicaris, Haesaert, Hartmanns, Khaled, Lavaei, Ma, Mallik, Niehage, Remke, Schupp, Shmarov, Soudjani, Thorpe, Turcuman & Zuliani (2021). We refer the interested readers to (Abate et al. 2020, Table 2) for more details on recent results of competitions.

## 12. Directions for Open Research

In this section, we present and discuss a few open topics that can be taken up as future research initiatives.

12.1. **Formal Analysis of SHS via Learning and Data-Driven Approaches.** We discuss a few results on formal synthesis of SHS via learning and data-driven approaches, which is still considered as an open direction. A deterministic policy gradient algorithm for reinforcement learning with continuous actions is presented by Silver et al. (2014). The framework introduces an off-policy actor-critic algorithm that learns a deterministic target policy from an exploratory behaviour policy. It shows that the deterministic policy gradient can be estimated much more efficiently than the usual stochastic policy gradient especially in high-dimensional action spaces. However, the results do not provide any quantitative guarantee on the optimality of synthesized policies for original MDPs. Policy synthesis over MDPs via sampling approaches (such as Reinforcement Learning) has been pursued by (Brázdil et al. 2014, Jaeger et al. 2020).

A model-free reinforcement learning framework of $\omega$-regular objectives for finite Markov decision processes is proposed by Hahn et al. (2019), Hasanbeig et al. (2019*a*). The $\omega$-regular properties are compiled into limit-deterministic Büchi automata (LDBA) instead of the traditional Rabin automata; this choice sidesteps difficulties that have marred previous proposals. (Hahn et al. 2019) present a constructive reduction from the almost-sure satisfaction of $\omega$-regular objectives to an almost-sure reachability problem, and learn how to control an unknown model so that the chance of satisfying the objective is maximized. (Hasanbeig et al. 2019*a*) exploit the structure of the LDBA and shapes a synchronous reward function on-the-fly, so that an RL algorithm can synthesize a policy resulting in traces that maximize the probability of satisfying the linear temporal property. The approach by Bozkurt et al. (2020) proposes a reward scheme associated to the given specification that requires two discounting factors in the reinforcement learning algorithm and provides a condition on these discounting to guarantee convergence of the learned policy to the optimal policy. These three approaches can be applied with off-the-shelf reinforcement learning algorithms to compute optimal strategies from the sample paths of the finite MDP. Extensions to continuous-space (and -actions) models are investigated by Lavaei, Somenzi, Soudjani, Trivedi & Zamani (2020), Kazemi & Soudjani (2020) and by Hasanbeig et al. (2019*b*, 2020), Cai et al. (2021), as surveyed in Section 7. The contribution in (Hammond et al. 2021) extends this setup to multi-agent cooperative games, and (Skalse et al. 2022) to a multi-objective RL framework. Early tool support for these methods has recently come to light (Hahn et al. 2021, Hasanbeig et al. 2022).

A data-driven verification approach under signal temporal logic constraints is proposed by Salamati et al. (2020, 2021). As the dynamics are parameterized and partially unknown, the framework collects data from the system and employs Bayesian inference techniques to associate a confidence value to the satisfaction of the property. The results combine both data-driven and model-based techniques in order to have a two-layer probabilistic reasoning over the behavior of the system: one layer is related to the stochastic noise inside the system and the next layer is related to the noisy data collected from the system. Approximate algorithms are also provided for computing the confidence for linear dynamical systems.

A data-driven technique for satisfying temporal properties on unknown stochastic processes with continuous spaces is presented by Kazemi & Soudjani (2020). The proposed framework is based on reinforcement learning that is used to compute sub-optimal policies that are finite-memory and deterministic. The work addresses properties expressed by LTL and uses their automaton representation to give a path-dependent reward function maximized via the RL algorithm. It also develops theoretical foundations characterizing the convergence of

the learned policy to the optimal one in the continuous space. To improve the performance of the learning on the constructed sparse reward function, the paper proposes a learning procedure based on a sequence of labelling functions obtained from the positive normal form of the LTL specification. This procedure is utilized to guide the RL algorithm towards the optimal policy. It is shown that the proposed approach can provide guaranteed lower bounds for the optimal satisfaction probability.

12.2. **Formal Analysis of Partially-Observed SHS.** With a few mentioned exceptions, most of the surveyed work on automated verification and synthesis of SHS assumes complete state information. However, in many real applications we do not have access to full information. There have been a limited work on formal synthesis of partially-observed SHS. An early formulation is put forward by Ding, Abate & Tomlin (2013), which characterizes the safety problem measure-theoretically and develops an application in air traffic management. Reachability analysis of partially observable discrete-time SHS is proposed by Lesser & Oishi (2014). A dynamic programming recursion is also developed for the solution of the equivalent perfect information problem, proving that the recursion is valid, an optimal solution exists, and results in the same solution as to the original problem.

A finite-state approximation for safety verification and control of partially observable SHS is presented by Lesser & Oishi (2015*b*, 2016). The papers solve a dynamic program over the finite state approximation to generate a lower bound to the viability probability, using a point-based method that generates samples of the information state. The proposed approach produces approximate probabilistic viable sets and synthesizes a controller to satisfy safety specifications. It also provides error bounds and convergence results, assuming additive Gaussian noise in the continuous-state dynamics and observations. Computing probabilistic viable sets for partially observable systems using truncated Gaussians and adaptive gridding is presented by Lesser & Oishi (2015*a*).

Verification of uncertain POMDPs using barrier certificates is discussed by Ahmadi et al. (2018). A class of POMDPs is considered with uncertain transition and/or observation probabilities in which the uncertainty takes the form of probability intervals. Given an uncertain POMDP representation of the system, the main goal is to propose a method for checking whether the system will satisfy an optimal performance, while not violating a safety requirement. A policy synthesis in multi-agent POMDPs via discrete-time barrier functions to enforce safety is proposed by Ahmadi et al. (2019). The method is implemented online by a sequence of one-step greedy algorithms as a standalone safe controller or as a safety-filter given a nominal planning policy. Verification of partial-information probabilistic systems using counterexample-guided refinements is studied by Giro & Rabe (2012).

A perception-aware point-based value iteration for POMDPs is presented by Ghasemi & Topcu (2019). The approach avoids combinatorial expansion over the action space from the integration of planning and perception decisions, through a greedy strategy for observation selection that minimizes an information-theoretic measure of the state uncertainty. The article develops a point-based value iteration algorithm that incorporates this greedy strategy to pick perception actions for each sampled belief point in each iteration. A sequential decision making process using POMDPs is studied by Wu et al. (2019). The work aims to find strategies that actively

interact with the system, and observe its reactions so that the true model is determined efficiently and with high confidence.

Synthesis of stochastic systems with partial state information via control barrier functions is proposed by Jahanshahi et al. (2020*b*,*a*, 2022), as surveyed in Section 9.

**Open Problem 9.** *Formal synthesis of POMDPs (even with finite set of states) is a hard problem and the available methods are not scalable. Developing scalable algorithmic techniques for POMDPs to make the synthesis problem more tractable is a potential future research direction.*

12.3. **Secure-by-Construction Controller Synthesis.** Security-related attacks are increasingly becoming pervasive in safety-critical applications, such as autonomous vehicles, implantable and wearable medical devices, smart systems and infrastructures. While most of the well-known attacks—such as vehicle hacking, pacemaker and Implantable Cardioverter Defibrillator (ICD) attacks (Halperin et al. 2008, Raghunathan & Jha 2011)—exploit unencrypted wireless communication, such attacks can be readily guarded against by following well-established cryptographic protocols. On the other hand, security vulnerabilities related to information leaks via side-channels may be impossible to mitigate without requiring a non-trivial modification to control software, as the side-channels are products of the interaction of the embedded control software with its physical environment (which clearly represents a "hybrid" feature). Furthermore, the presence of wide variety of physical variables (such as temperature, electro-magnetic emissions, velocity and so on) in these control systems expose corresponding attack surfaces to the intruder and render those systems even more vulnerable than traditional digital software/hardware systems. The source of stochasticity in those systems is either due to the noisy environment or the measurement noise of intruders' observations. Hence, SHS are a good modeling framework for studying those security vulnerabilities. We refer the interested readers to the recent vision paper in (Liu et al. 2022) explaining in detail different security notions for hybrid systems, some interesting examples, and initial verification and controller synthesis approaches suitable for them.

While the controller synthesis approach for SHS has been heavily investigated for safety requirements as discussed in details in the previous sections, the secrecy requirements in SHS are often verified in a post facto manner after the design of controllers. Hence, if the system leaks information beyond an acceptable range, the controller needs to be redesigned incurring very high verification and validation costs. The secure-by-construction synthesis approach advocates a paradigm shift in the development of safe and secure SHS by proposing a controller design scheme which generalizes existing correct-by-construction synthesis methods by considering security properties, simultaneously to the safety ones discussed elsewhere in this survey, during the design phase.

**Open Problem 10.** *Correct-by-construction controller synthesis approaches for SHS provide embedded control software from high-level safety requirements in an automated and formal manner. Proposing secure-by-construction controller synthesis schemes which generalize the correct-by-construction ones by integrating security requirements with the safety ones in the controller synthesis phase is a potential future research direction.*

12.4. **(Mix)-monotonicity of SHS.** As discussed in Subsection 5.3, the construction of IMC/IMDP can be much more complex than standard abstractions based on MC/MDP, since one needs to compute lower and upper bounds for the probabilities of transition between states, rather than computing just a single number as in standard MCs/MDPs. Under some assumptions (*e.g.,* additive stochasticity, unimodal noise distribution, independent noises affecting different states), contributions in (Dutreix & Coogan 2018, 2020, Dutreix et al. 2022) utilize the mix-monotonicity property of the deterministic part of the map $f$ and propose an approach to compute those lower and upper bounds in an efficient way. Further research in this direction is deemed worthy of attention.

12.5. **Compositional Construction of IMC/IMDP.** Since constructing IMC/IMDP is more complex than standard abstractions, as discussed in Subsection 5.3, a promising approach to mitigate the related computational complexity is to develop compositional techniques: these might allow constructing IMC/IMDP of high-dimensional systems based on IMC/IMDP of smaller subsystems.

12.6. **Compositional Controller Synthesis for SHS.** In this survey, we mainly discussed different compositional approaches for the construction of (in)finite abstractions for networks of stochastic systems. Potential future work concerns the investigation of compositional controller synthesis for stochastic hybrid systems. In particular, given a specification over the interconnected system, it is of interest to find a formal relation between the satisfaction probabilities provided by local controllers for individual subsystems, as well as the optimal satisfaction probability for the specification on the monolithic (overall) system.

12.7. **Extensions and Development of Software Tools.** Developing efficient software tools based on theoretical and algorithmic results is essential for the practical use of automated verification and synthesis of SHS. Most of the tools discussed in Section 11 are developed for discrete-time models. The software tools Level-Set Toolbox (Mitchell 2007) and Uppaal (David et al. 2015) allow for the analysis of continuous-time models. Although software tools (Rungger & Zamani 2016, Khaled & Zamani 2019, 2021) are developed for formal controller synthesis of non-stochastic control systems, they can be readily utilized for the proposed results for incrementally stable SHS in (Zamani, Mohajerin Esfahani, Majumdar, Abate & Lygeros 2014, Zamani, Tkachev & Abate 2014, Zamani et al. 2015, Mallik et al. 2017). A future direction is to develop more general and scalable software tools for *continuous-time stochastic hybrid systems*. Moreover, developing software tools to handle *infinite-horizon specifications* is another unmet extension. Finally, there is no tool at the moment that handles the construction of finite MDPs *compositionally*: further developing software tools for compositional purposes is therefore of interest.

A comprehensive and up-to-date discussion about different software tools together with their potential directions of extension can be found in (Abate et al. 2020) and available at https://bit.ly/3nGechr.

## 13. Closing Discussion

In this article, we have provided the first survey of work on automated formal verification and control synthesis of stochastic hybrid systems (SHS). We have focused on most recent and sharpest results, and for the

sake of a clear and streamlined presentation we have presented selected analysis methods, applications, and results in detail, and instead briefly overviewed alternative approaches. We have distinguished approaches as discretization-based and -free, and have investigated four different closeness guarantees between a concrete SHS model and its abstractions. We have discussed different problems including stochastic similarity relations, infinite and finite abstractions, the use of control barrier certificates, temporal logic verification and synthesis, compositional techniques, continuous-time stochastic models, data-drive approaches, and finally overviewed existing software tools that implement the discussed approaches. Throughout this survey, we have also added the discussion of a few open problems.

We hope that this survey article provides an introduction to the foundations of SHS, towards an easier understanding of many challenges and existing solutions related to formal verification and control synthesis of these models, together with the associated software tools.

## References

Abate, A. (2009), A contractivity approach for probabilistic bisimulations of diffusion processes, *in* 'Proceedings of the 48th IEEE Conference of Decision and Control', pp. 2230–2235.

Abate, A. (2010), Probabilistic bisimulations of switching and resetting diffusions, *in* 'Proceedings of the 49th IEEE Conference of Decision and Control', pp. 5918–5923.

Abate, A. (2013), 'Approximation metrics based on probabilistic bisimulations for general state-space Markov processes: a survey', *Electronic Notes in Theoretical Computer Science* **297**, 3–25.

Abate, A., Ahmed, D., Edwards, A., Giacobbe, M. & Peruffo, A. (2021), FOSSIL: A software tool for the formal synthesis of lyapunov functions and barrier certificates using neural networks, *in* 'Proceedings of HSCC', pp. 1–11.

Abate, A., Blom, H., Bouissou, M., Cauchi, N., Chraibi, H., Delicaris, J., Haesaert, S., Hartmanns, A., Khaled, M., Lavaei, A., Ma, H., Mallik, K., Niehage, M., Remke, A., Schupp, S., Shmarov, F., Soudjani, S., Thorpe, A., Turcuman, V. & Zuliani, P. (2021), Arch-comp21 category report: Stochastic models, *in* '8th International Workshop on Applied Verification of Continuous and Hybrid Systems', pp. 55–89.

Abate, A., Blom, H., Cauchi, N., Degiorgio, K., Fraenzle, M., Hahn, E. M., Haesaert, S., Ma, H., Oishi, M., Pilch, C., Remke, A., Salamati, M., Soudjani, S., van Huijgevoort, B. & Vinod, A. (2019), ARCH-COMP19 category report: Stochastic modelling, *in* '6th International Workshop on Applied Verification of Continuous and Hybrid Systems', Vol. 61 of *EPiC Series in Computing*, pp. 62–102.

Abate, A., Blom, H., Cauchi, N., Delicaris, J., Hartmanns, A., Khaled, M., Lavaei, A., Pilch, C., Remke, A., Schupp, S., Shmarov, F., Soudjani, S., Vinod, A., Wooding, B., Zamani, M. & Zuliani, P. (2020), ARCH-COMP20 category report: Stochastic models, *in* '7th International Workshop on Applied Verification of Continuous and Hybrid Systems (ARCH20)', Vol. 74 of *EPiC Series in Computing*, pp. 76–106.

Abate, A., Blom, H., Cauchi, N., Haesaert, S., Hartmanns, A., Lesser, K., Oishi, M., Sivaramakrishnan, V., Soudjani, S., Vasile, C.-I. & Vinod, A. P. (2018), ARCH-COMP18 category report: Stochastic modelling, *in* '5th International Workshop on Applied Verification of Continuous and Hybrid Systems', Vol. 54 of *EPiC Series in Computing*, pp. 71–103.

Abate, A., D'Innocenzo, A. & Benedetto, M. D. (2011), 'Approximate abstractions of stochastic hybrid systems', *IEEE Transactions on Automatic Control* **56**(11), 2688–2694.

Abate, A., Katoen, J., Lygeros, J. & Prandini, M. (2010), 'Approximate model checking of stochastic hybrid systems', *European Journal of Control* **16**(6), 624–641.

Abate, A., Kwiatkowska, M., Norman, G. & Parker, D. (2014), Probabilistic model checking of labelled Markov processes via finite approximate bisimulations, *in* 'Horizons of the Mind. A Tribute to Prakash Panangaden', Springer, pp. 40–58.

Abate, A., Prandini, M., Lygeros, J. & Sastry, S. (2008), 'Probabilistic reachability and safety for controlled discrete-time stochastic hybrid systems', *Automatica* **44**(11), 2724–2734.

Agha, G. & Palmskog, K. (2018), 'A survey of statistical model checking', *ACM Transactions on Modeling and Computer Simulation (TOMACS)* **28**(1), 1–39.

Ahmadi, M., Cubuktepe, M., Jansen, N. & Topcu, U. (2018), Verification of uncertain POMDPs using barrier certificates, *in* 'Proceedings of the Annual Allerton Conference on Communication, Control, and Computing', pp. 115–122.

Ahmadi, M., Singletary, A., Burdick, J. W. & Ames, A. D. (2019), Safe policy synthesis in multi-agent POMDPs via discrete-time barrier functions, *in* 'Proceedings of the 58th Conference on Decision and Control (CDC)', pp. 4797–4803.

Amin, S., Abate, A., Prandini, M., Lygeros, J. & Sastry, S. (2006), Reachability analysis for controlled discrete time stochastic hybrid systems, *in* 'Proceedings of HSCC06, LNCS 3927', Springer Verlag, pp. 49–63.

Anand, M., Lavaei, A. & Zamani, M. (2021), 'Compositional synthesis of control barrier certificates for networks of stochastic systems against omega-regular specifications', *arXiv:2103.02226* .

Anand, M., Lavaei, A. & Zamani, M. (2022), 'From small-gain theory to compositional construction of barrier certificates for large-scale stochastic systems', *IEEE Transactions on Automatic Control* .

Antoulas, A. C. (2005), *Approximation of large-scale dynamical systems*, SIAM.

Arapostathis, A., Borkar, V. S., Fernández-Gaucherand, E., Ghosh, M. K. & Marcus, S. I. (1993), 'Discrete-time controlled Markov processes with average cost criterion: A survey', *SIAM J. Control Optim.* **31**(2), 282–344.

Arcak, M., Meissen, C. & Packard, A. (2016), *Networks of dissipative systems: compositional certification of stability, performance, and safety*, Springer.

Awan, A. U. & Zamani, M. (2018), Compositional abstractions of networks of stochastic hybrid systems under randomly switched topologies, *in* 'Proceedings of the American Control Conference (ACC)', pp. 1586–1591.

Awan, A. U. & Zamani, M. (2019), 'From dissipativity theory to compositional abstractions of interconnected stochastic hybrid systems', *IEEE Transactions on Control of Network Systems* **7**(1), 433–445.

Badings, T., Abate, A., Jansen, N., Parker, D., Poonawala, H. & Stoelinga, M. (2022), Sampling-based robust control of autonomous systems with non-gaussian noise, *in* 'Proceedings of AAAI'.

Baier, C. & Katoen, J.-P. (2008), *Principles of model checking*, MIT Press.

Belta, C., Yordanov, B. & Gol, E. A. (2017), *Formal Methods for Discrete-Time Dynamical Systems*, Vol. 89 of *Studies in Systems, Decision and Control*, Springer.

Bertsekas, D. P. & Shreve, S. E. (1996), *Stochastic optimal control: The discrete-time case*, Athena Scientific.

Bian, G. & Abate, A. (2017), On the relationship between bisimulation and trace equivalence in an approximate probabilistic context, *in* 'Proceedings of the International Conference on Foundations of Software Science and Computation Structures', pp. 321–337.

Blom, H. A., Krystul, J., Bakker, G., Klompstra, M. B. & Obbink, B. K. (2007), 'Free flight collision risk estimation by sequential MC simulation', pp. 249–281.

Blom, H. A., Lygeros, J., Everdij, M., Loizou, S. & Kyriakopoulos, K. (2006), *Stochastic hybrid systems: theory and safety critical applications*, Vol. 337, Springer.

Blom, H. A., Stroeve, S. H. & Bosse, T. (2013), Modelling of potential hazards in agent-based safety risk analysis, *in* 'Proceedings of the 10th USA/Europe Air Traffic Management Research and Development Seminar'.

Blom, H. & Bloem, E. (2004), Particle filtering for stochastic hybrid systems, *in* '43rd IEEE Conference on Decision and Control (CDC)', Vol. 3, pp. 3221–3226.

Blom, H. & Bloem, E. (2007), 'Exact bayesian and particle filtering of stochastic hybrid systems', *IEEE Transactions on Aerospace and Electronic Systems* **43**(1), 55–70.

Bouissou, M., Elmqvist, H., Otter, M. & Benveniste, A. (2014), Efficient Monte Carlo simulation of stochastic hybrid systems, *in* 'Proceedings of the 10th International Modelica Conference'.

Bozkurt, A., Wang, Y., Zavlanos, M. M. & Pajic, M. (2020), 'Control synthesis from linear temporal logic specifications using model-free reinforcement learning', *2020 IEEE International Conference on Robotics and Automation (ICRA)* pp. 10349–10355.

Brázdil, T., Chatterjee, K., Chmelík, M., Forejt, V., Křetínský, J., Kwiatkowska, M., Parker, D. & Ujma, M. (2014), Verification of Markov decision processes using learning algorithms, *in* F. Cassez & J. Raskin, eds, 'ATVA', Springer Verlag, pp. 98–114.

Bujorianu, L. M. (2012), *Stochastic reachability analysis of hybrid systems*, Springer Science & Business Media.

Bujorianu, M. L. & Lygeros, J. (2003), Reachability questions in piecewise deterministic Markov processes, *in* 'Proceedings of the 6th International Workshop on Hybrid Systems: Computation and Control', Vol. 2623 of *Lecture Notes in Computer Science*, Springer, pp. 126–140.

Cai, M., Hasanbeig, M., Xiao, S., Abate, A. & Kan, Z. (2021), 'Modular deep reinforcement learning for continuous motion planning with temporal logic', *IEEE Robotics and Automation Letters* **6**(4), 7973–7980.

Calafiore, G. C. & Campi, M. C. (2006), 'The scenario approach to robust control design', *IEEE Transactions on Automatic Control* **51**(5), 742–753.

Campi, M. C. & Garatti, S. (2018), *Introduction to the scenario approach*, SIAM.

Cassandras, C. G. & Lygeros, J. (2006), *Stochastic hybrid systems*, CRC Press.

Cauchi, N. & Abate, A. (2018), Benchmarks for cyber-physical systems: A modular model library for building automation systems, *in* 'Proceedings of ADHS', pp. 49–54.

Cauchi, N., Degiorgio, K. & Abate, A. (2019), StocHy: Automated verification and synthesis of stochastic processes, *in* 'Proceedings of TACAS'19', Lecture Notes in Computer Science, Springer, pp. 247–264.

Cauchi, N., Laurenti, L., Lahijanian, M., Abate, A., Kwiatkowska, M. & Cardelli, L. (2019), Efficiency through uncertainty: Scalable formal synthesis for stochastic hybrid systems, *in* 'Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control', pp. 240–251.

Chakarov, A. & Sankaranarayanan, S. (2013), Probabilistic program analysis with martingales, *in* 'Proceedings of the International Conference on Computer Aided Verification', pp. 511–526.

Cheng, X., Kawano, Y. & Scherpen, J. M. (2017), 'Reduction of second-order network systems with structure preservation', *IEEE Transactions on Automatic Control* **62**(10), 5026–5038.

Choi, J. J., Agrawal, A., Sreenath, K., Tomlin, C. J. & Bansal, S. (2022), 'Computation of regions of attraction for hybrid limit cycles using reachability: An application to walking robots', *IEEE Robotics and Automation Letters* **7**(2), 4504–4511.

Ciesinski, F. & Größer, M. (2004), On probabilistic computation tree logic, *in* 'Validation of Stochastic Systems', Springer, pp. 147–188.

Cimatti, A., Griggio, A., Schaafsma, B. J. & Sebastiani, R. (2013), The MathSAT5 SMT solver, *in* 'Tools and Algorithms for the Construction and Analysis of Systems', Lecture Notes in Computer Science, pp. 93–107.

Clark, A. (2019), Control barrier functions for complete and incomplete information stochastic systems, *in* 'Proceedings of the American Control Conference (ACC)', pp. 2928–2935.

Clark, A. (2021), 'Control barrier functions for stochastic systems', *Automatica* **130**.

Coogan, S. & Arcak, M. (2015), Efficient finite abstraction of mixed monotone systems, *in* 'Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control', pp. 58–67.

Cosentino, F., Oberhauser, H. & Abate, A. (2021), 'Grid-free computation of probabilistic safety with malliavin calculus', *arXiv preprint: 2104.14691* .

Dashkovskiy, S. N., Rüffer, B. S. & Wirth, F. R. (2010), 'Small gain theorems for large scale systems and construction of ISS Lyapunov functions', *SIAM Journal on Control and Optimization* **48**(6), 4089–4118.

David, A., Jensen, P. G., Larsen, K. G., Mikucionis, M. & Taankvist, J. H. (2015), Uppaal stratego, *in* 'TACAS', Springer, pp. 206–211.

Davis, M. H. A. (1993), *Markov models and optimization*, Vol. 49 of *Monographs on Statistics and Applied Probability*, Chapman & Hall, London.

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

Dehnert, C., Gebler, D., Volpato, M. & Jansen, D. N. (2012), On abstraction of probabilistic systems, *in* 'International Autumn School on Rigorous Dependability Analysis Using Model Checking Techniques for Stochastic Systems', Springer, pp. 87–116.

Dehnert, C., Junges, S., Katoen, J. & Volk, M. (2017), A storm is coming: A modern probabilistic model checker, *in* 'Proceedings of the 29th International Conference on Computer Aided Verification (CAV)', Vol. 10427 of *Lecture Notes in Computer Science*, Springer, pp. 592–600.

Delahaye, B., Caillaud, B. & Legay, A. (2011), 'Probabilistic contracts: a compositional reasoning methodology for the design of systems with stochastic and/or non-deterministic aspects', *Formal Methods in System Design* **38**(1), 1–32.

Desharnais, J., Gupta, V., Jagadeesan, R. & Panangaden, P. (2004), 'Metrics for labelled Markov processes', *Theoretical computer science* **318**(3), 323–354.

Desharnais, J., Laviolette, F. & Tracol, M. (2008), Approximate analysis of probabilistic processes: Logic, simulation and games, *in* 'Proceedings of the 5th international conference on quantitative evaluation of

system', pp. 264–273.

Ding, J. (2012), *Methods for reachability-based hybrid controller design*, Ph.D. Dissertation, University of California, Berkeley.

Ding, J., Abate, A. & Tomlin, C. (2013), Optimal control of partially observable discrete time stochastic hybrid systems for safety specifications, *in* 'Proceedings of the 2013 American Control Conference', pp. 6231–6236.

Ding, J., Kamgarpour, M., Summers, S., Abate, A., Lygeros, J. & Tomlin, C. (2013), 'A stochastic games framework for verification and control of discrete time stochastic hybrid systems', *Automatica* **49**(9), 2665–2674.

D'Innocenzo, A., Abate, A. & Katoen, J. (2012), Robust PCTL model checking, *in* 'Proceedings of the 15th ACM international conference on Hybrid Systems: Computation and Control', pp. 275–286.

Dutreix, M. & Coogan, S. (2018), Efficient verification for stochastic mixed monotone systems, *in* 'Proceedings of the 9th International Conference on Cyber-Physical Systems (ICCPS)', pp. 150–161.

Dutreix, M. & Coogan, S. (2020), 'Specification-guided verification and abstraction refinement of mixed monotone stochastic systems', *IEEE Transactions on Automatic Control* **66**(7), 2975–2990.

Dutreix, M., Huh, J. & Coogan, S. (2022), 'Abstraction-based synthesis for stochastic systems with omega-regular objectives', *Nonlinear Analysis: Hybrid Systems* **45**.

Ellen, C., Gerwinn, S. & Fränzle, M. (2012), Confidence bounds for statistical model checking of probabilistic hybrid systems, *in* 'International Conference on Formal Modeling and Analysis of Timed Systems', Springer, pp. 123–138.

Farahani, S. S., Majumdar, R., Prabhu, V. S. & Soudjani, S. (2018), 'Shrinking horizon model predictive control with signal temporal logic constraints under stochastic disturbances', *IEEE Transactions on Automatic Control* **64**(8), 3324–3331.

Fehnker, A. & Ivančić, F. (2004), Benchmarks for hybrid systems verification, *in* 'International Workshop on Hybrid Systems: Computation and Control', Springer, pp. 326–341.

Forejt, V., Kwiatkowska, M., Norman, G. & Parker, D. (2011), Automated verification techniques for probabilistic systems, *in* 'International school on formal methods for the design of computer, communication and software systems', Springer, pp. 53–113.

Fränzle, M., Hahn, E. M., Hermanns, H., Wolovick, N. & Zhang, L. (2011), Measurability and safety verification for stochastic hybrid systems, *in* 'Proceedings of the 14th international conference on hybrid systems: computation and control', pp. 43–52.

Fu, J. & Topcu, U. (2015), Computational methods for stochastic control with metric interval temporal logic specifications, *in* 'Proceedings of the 54th IEEE Conference on Decision and Control (CDC)', pp. 7440–7447.

Gao, S., Avigad, J. & Clarke, E. M. (2012), $\delta$-complete decision procedures for satisfiability over the reals, *in* 'Automated Reasoning', Lecture Notes in Computer Science, pp. 286–300.

Gao, Y., Lygeros, J. & Quincampoix, M. (2006), The reachability problem for uncertain hybrid systems revisited: A viability theory perspective, *in* 'Proceedings of the 9th International Workshop on Hybrid Systems: Computation and Control', Vol. 3927 of *Lecture Notes in Computer Science*, Springer, pp. 242–256.

Ghasemi, M. & Topcu, U. (2019), Perception-aware point-based value iteration for partially observable Markov decision processes, *in* 'Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)', pp. 2371–2377.

Giro, S. & Rabe, M. N. (2012), Verification of partial-information probabilistic systems using counterexample-guided refinements, *in* 'International Symposium on Automated Technology for Verification and Analysis', Springer, pp. 333–348.

Gleason, J. D., Vinod, A. P. & Oishi, M. M. (2017), Underapproximation of reach-avoid sets for discrete-time stochastic systems via Lagrangian methods, *in* 'Proceedings of the 56th Conference on Decision and Control', pp. 4283–4290.

Gleason, J. D., Vinod, A. P. & Oishi, M. M. (2021), 'Lagrangian approximations for stochastic reachability of a target tube', *Automatica* **128**, 109546.

Gol, E. A., Lazar, M. & Belta, C. (2015), 'Temporal logic model predictive control', *Automatica* **56**, 78–85.

Haesaert, S., Cauchi, N. & Abate, A. (2017), 'Certified policy synthesis for general Markov decision processes: An application in building automation systems', *Performance Evaluation* **117**, 75–103.

Haesaert, S., Nilsson, P. & Soudjani, S. (2021), 'Formal multi-objective synthesis of continuous-state MDPs', *IEEE Control Systems Letters* **5**(5), 1765–1770.

Haesaert, S. & Soudjani, S. (2020), 'Robust dynamic programming for temporal logic control of stochastic systems', *IEEE Transactions on Automatic Control* **66**(6), 2496–2511.

Haesaert, S., Soudjani, S. & Abate, A. (2017), 'Verification of general Markov decision processes by approximate similarity relations and policy refinement', *SIAM Journal on Control and Optimization* **55**(4), 2333–2367.

Haesaert, S., Soudjani, S. & Abate, A. (2018), 'Temporal logic control of general Markov decision processes by approximate policy refinement', *IFAC-PapersOnLine* **51**(16), 73–78.

Hahn, E. M., Perez, M., Schewe, S., Somenzi, F., Trivedi, A. & Wojtczak, D. (2019), Omega-regular objectives in model-free reinforcement learning, *in* 'Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems', pp. 395–412.

Hahn, E., Perez, M., Schewe, S., Somenzi, F., Trivedi, A. & Wojtczak, D. (2021), 'Mungojerrie: Reinforcement learning of linear-time objectives', *arXiv:2106.09161* .

Hall, P. & Heyde, C. C. (2014), *Martingale limit theory and its application*, Academic Press.

Halperin, D., Heydt-Benjamin, T. S., Ransford, B., Clark, S. S., Defend, B., Morgan, W., Fu, K., Kohno, T. & Maisel, W. H. (2008), Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses, *in* '2008 IEEE Symposium on Security and Privacy', pp. 129–142.

Hammond, L., Abate, A., Gutierrez, J. & Wooldridge, M. (2021), Multi-agent reinforcement learning with temporal logic specifications, *in* 'Proceedings of AAAMAS', pp. 583–592.

Hartfiel, D. J. (2006), *Markov set-chains*, Springer.

Hartmanns, A. & Hermanns, H. (2014), The modest toolset: An integrated environment for quantitative modelling and verification, *in* 'Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems', pp. 593–598.

Hasanbeig, M., Abate, A. & Kroening, D. (2019*a*), 'Certified reinforcement learning with logic guidance', *arXiv:1902.00778* .

Hasanbeig, M., Abate, A. & Kroening, D. (2019*b*), Logically-Constrained Neural Fitted Q-Iteration, *in* 'Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)', pp. 2012–2014.

Hasanbeig, M., Kroening, D. & Abate, A. (2020), Deep reinforcement learning with temporal logics, *in* 'Proceedings of FORMATS22', Springer LNCS 12288, pp. 1–22.

Hasanbeig, M., Kroening, D. & Abate, A. (2022), LCRL: Certified policy synthesis via logically-constrained reinforcement learning, *in* 'Proceedings of QEST22'.

Henriques, D., Martins, J. G., Zuliani, P., Platzer, A. & Clarke, E. M. (2012), Statistical model checking for Markov decision processes, *in* 'Proceedings of the 9th international conference on quantitative evaluation of systems', pp. 84–93.

Hermanns, H., Wachter, B. & Zhang, L. (2008), Probabilistic CEGAR, *in* 'International Conference on Computer Aided Verification', Springer, pp. 162–175.

Hernández-Lerma, O. & Lasserre, J. B. (1996), *Discrete-time Markov control processes*, Appl. Math. 30, Springe, New York.

Hespanha, J. P. & Singh, A. (2005), 'Stochastic models for chemically reacting systems using polynomial stochastic hybrid systems', *International Journal of Robust and Nonlinear Control: IFAC-Affiliated Journal* **15**(15), 669–689.

Hsu, K., Majumdar, R., Mallik, K. & Schmuck, A.-K. (2018), Multi-layered abstraction-based controller synthesis for continuous-time systems, *in* 'Proceedings of the 21st International Conference on Hybrid Systems: Computation and Control', pp. 120–129.

Hu, J., Lygeros, J. & Sastry, S. (2000), Towars a theory of stochastic hybrid systems, *in* 'Proceedings of the 3rd International Workshop on Hybrid Systems: Computation and Control', Vol. 1790 of *Lecture Notes in Computer Science*, Springer, pp. 160–173.

Hu, J., Wu, W.-C. & Sastry, S. (2004), Modeling subtilin production in bacillus subtilis using stochastic hybrid systems, *in* 'International Workshop on Hybrid Systems: Computation and Control', Springer, pp. 417–431.

Huang, C., Chen, X., Lin, W., Yang, Z. & Li, X. (2017), 'Probabilistic safety verification of stochastic hybrid systems using barrier certificates', *ACM Transactions on Embedded Computing Systems (TECS)* **16**(5s), 186.

Ionescu, T. C. & Astolfi, A. (2015), 'Nonlinear moment matching-based model order reduction', *IEEE Transactions on Automatic Control* **61**(10), 2837–2847.

Jaeger, M., Bacci, G., Bacci, G., Larsen, K. & Jensen, P. (2020), Approximating euclidean by imprecise markov decision processes, *in* T. Margaria & B. Steffen, eds, 'Proceedings of ISoLA20', Vol. 12476 of *LNCS*, Springer, pp. 275–289.

Jagtap, P., Soudjani, S. & Zamani, M. (2018), Temporal logic verification of stochastic systems using barrier certificates, *in* 'Proceedings of the International Symposium on Automated Technology for Verification and Analysis', pp. 177–193.

Jagtap, P., Soudjani, S. & Zamani, M. (2020), 'Formal synthesis of stochastic systems via control barrier certificates', *IEEE Transactions on Automatic Control* **66**(7), 3097–3110.

Jahanshahi, N., Jagtap, P. & Zamani, M. (2020*a*), 'Synthesis of partially observed jump-diffusion systems via control barrier functions', *IEEE Control Systems Letters* **5**(1), 253–258.

Jahanshahi, N., Jagtap, P. & Zamani, M. (2020*b*), 'Synthesis of stochastic systems with partial information via control barrier functions', *IFAC-PapersOnLine* **53**(2), 2441–2446.

Jahanshahi, N., Lavaei, A. & Zamani, M. (2022), 'Compositional construction of safety controllers for networks of continuous-space pomdps', *IEEE Transactions on Control of Network Systems* .

Julius, A. A. & Pappas, G. J. (2009), 'Approximations of stochastic hybrid systems', *IEEE Transactions on Automatic Control* **54**(6), 1193–1203.

Junges, S. (2020), Parameter Synthesis in Markov Models, PhD thesis, RWTH Aachen University, Germany.

Kallenberg, O. (1997), *Foundations of modern probability*, Springer-Verlag, New York.

Kamgarpour, M., Ding, J., Summers, S., Abate, A., Lygeros, J. & Tomlin, C. (2011), Discrete time stochastic hybrid dynamical games: Verification & controller synthesis, *in* 'Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference', pp. 6122–6127.

Kamgarpour, M., Summers, S. & Lygeros, J. (2013), Control design for specifications on stochastic hybrid systems, *in* 'Proceedings of the 16th international conference on hybrid systems: computation and control', pp. 303–312.

Kariotoglou, N., Kamgarpour, M., Summers, T. H. & Lygeros, J. (2017), 'The linear programming approach to reach-avoid problems for Markov decision processes', *Journal of Artificial Intelligence Research* **60**, 263–285.

Kattenbelt, M. & Huth, M. (2009), Verification and refutation of probabilistic specifications via games, *in* 'IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science', Schloss Dagstuhl-Leibniz-Zentrum für Informatik.

Kazemi, M. & Soudjani, S. (2020), Formal policy synthesis for continuous-space systems via reinforcement learning, *in* 'International Conference on Integrated Formal Methods', pp. 3–21.

Khaled, M. & Zamani, M. (2019), pFaces: An acceleration ecosystem for symbolic control., *in* 'International Conference on Hybrid Systems: Computation and Control', pp. 252–257.

Khaled, M. & Zamani, M. (2021), OmegaThreads: Symbolic controller design for omega-regular objectives, *in* 'The 24th ACM International Conference on Hybrid Systems: Computation and Control (HSCC)', pp. 1–7.

Komuravelli, A., Păsăreanu, C. S. & Clarke, E. M. (2012), Assume-guarantee abstraction refinement for probabilistic systems, *in* 'International Conference on Computer Aided Verification', Springer, pp. 310–326.

Koutsoukos, X. D. & Riley, D. (2006), Computational methods for reachability analysis of stochastic hybrid systems, *in* 'Proceedings of the 9th International Workshop on Hybrid Systems: Computation and Control', Vol. 3927 of *Lecture Notes in Computer Science*, Springer, pp. 377–391.

Kupferman, O. & Vardi, M. Y. (2001), 'Model checking of safety properties', *Formal Methods in System Design* **19**(3), 291–314.

Kushner, H. J. (1967), *Stochastic Stability and Control*, Mathematics in Science and Engineering, Elsevier Science.

Kwiatkowska, M., Norman, G. & Parker, D. (2002), PRISM: Probabilistic symbolic model checker, *in* 'Proceedings of the International Conference on Modelling Techniques and Tools for Computer Performance Evaluation', pp. 200–204.

Kwiatkowska, M., Norman, G. & Parker, D. (2011), Prism 4.0: Verification of probabilistic real-time systems, *in* 'International conference on computer aided verification', Springer, pp. 585–591.

Kwiatkowska, M., Norman, G., Parker, D. & Qu, H. (2013), 'Compositional probabilistic verification through multi-objective model checking', *Information and Computation* **232**, 38–65.

Lacerda, B., Parker, D. & Hawes, N. (2014), Optimal and dynamic planning for markov decision processes with co-safe ltl specifications, *in* '2014 IEEE/RSJ International Conference on Intelligent Robots and Systems', pp. 1511–1516.

Lahijanian, M., Andersson, S. B. & Belta, C. (2009), A probabilistic approach for control of a stochastic system from LTL specifications, *in* 'Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 28th Chinese Control Conference', pp. 2236–2241.

Lahijanian, M., Andersson, S. B. & Belta, C. (2011), 'Temporal logic motion planning and control with probabilistic satisfaction guarantees', *IEEE Transactions on Robotics* **28**(2), 396–409.

Lahijanian, M., Andersson, S. B. & Belta, C. (2012), Approximate Markovian abstractions for linear stochastic systems, *in* 'Proceedings of the 51st IEEE Conference on Decision and Control (CDC)', pp. 5966–5971.

Lahijanian, M., Andersson, S. B. & Belta, C. (2015), 'Formal verification and synthesis for discrete-time stochastic systems', *IEEE Transactions on Automatic Control* **60**(8), 2031–2045.

Lal, R., Duan, W. & Prabhakar, P. (2020), Bayesian statistical model checking for continuous stochastic logic, *in* '18th ACM-IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE)', pp. 1–11.

Lal, R. & Prabhakar, P. (2018), Hierarchical abstractions for reachability analysis of probabilistic hybrid systems, *in* '2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton)', IEEE, pp. 848–855.

Lal, R. & Prabhakar, P. (2019), 'Counterexample guided abstraction refinement for polyhedral probabilistic hybrid systems', *ACM Transactions on Embedded Computing Systems (TECS)* **18**(5s), 1–23.

Lal, R. & Prabhakar, P. (2020), Safety analysis of linear discrete-time stochastic systems: Work-in-progress, *in* '2020 International Conference on Embedded Software (EMSOFT)', pp. 34–36.

Larsen, K. G. & Skou, A. (1991), 'Bisimulation through probabilistic testing', *Information and computation* **94**(1), 1–28.

Laurenti, L., Abate, A., Bortolussi, L., Cardelli, L., Ceska, M. & Kwiatkowska, M. (2017), Reachability computation for switching diffusions: Finite abstractions with certifiable and tuneable precision, *in* 'Proceedings of the 20th ACM International Conference on Hybrid Systems: Computation and Control', pp. 55–64.

Laurenti, L., Kwiatkovska, M., Patane, A., Wickert, M. & Abate, A. (2021), Strategy synthesis for probabilistic reach-avoid for learned bayesian neural network models, *in* 'Proceedings of UAI21 - PMLR 161', p. 1713–1723.

Laurenti, L., Lahijanian, M., Abate, A., Cardelli, L. & Kwiatkowska, M. (2020), 'Formal and efficient synthesis for continuous-time linear stochastic hybrid processes', *IEEE Transactions on Automatic Control* **66**(1), 17–32.

Lavaei, A., Khaled, M., Soudjani, S. & Zamani, M. (2020), AMYTISS: Parallelized automated controller synthesis for large-scale stochastic systems, *in* 'Proceedings of the 32nd International Conference on Computer-Aided Verification (CAV), Lecture Notes in Computer Science 12225', pp. 461–474.

Lavaei, A., Somenzi, F., Soudjani, S., Trivedi, A. & Zamani, M. (2020), Formal controller synthesis for continuous-space MDPs via model-free reinforcement learning, *in* 'Proceedings of the 11th ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS)', pp. 98–107.

Lavaei, A., Soudjani, S., Majumdar, R. & Zamani, M. (2017), Compositional abstractions of interconnected discrete-time stochastic control systems, *in* 'Proceedings of the 56th IEEE Conference on Decision and Control', pp. 3551–3556.

Lavaei, A., Soudjani, S. & Zamani, M. (2018), From dissipativity theory to compositional construction of finite Markov decision processes, *in* 'Proceedings of the 21st ACM International Conference on Hybrid Systems: Computation and Control', pp. 21–30.

Lavaei, A., Soudjani, S. & Zamani, M. (2019), 'Compositional construction of infinite abstractions for networks of stochastic control systems', *Automatica* **107**, 125–137.

Lavaei, A., Soudjani, S. & Zamani, M. (2020*a*), 'Compositional abstraction-based synthesis for networks of stochastic switched systems', *Automatica* **114**.

Lavaei, A., Soudjani, S. & Zamani, M. (2020*b*), 'Compositional abstraction-based synthesis of general MDPs via approximate probabilistic relations', *Nonlinear Analysis: Hybrid Systems* **39**.

Lavaei, A., Soudjani, S. & Zamani, M. (2020*c*), 'Compositional abstraction of large-scale stochastic systems: A relaxed dissipativity approach', *Nonlinear Analysis: Hybrid Systems* **36**.

Lavaei, A., Soudjani, S. & Zamani, M. (2020*d*), 'Compositional (in)finite abstractions for large-scale interconnected stochastic systems', *IEEE Transactions on Automatic Control* **65**(12), 5280–5295.

Lavaei, A. & Zamani, M. (2021), 'From dissipativity theory to compositional synthesis of large-scale stochastic switched systems', *IEEE Transactions on Automatic Control* .

Lesser, K. & Abate, A. (2018), 'Multi-objective optimal control with safety as a priority', *IEEE Transactions on Control Systems Technology* **26**(3), 1015–1027.

Lesser, K. & Oishi, M. (2014), 'Reachability for partially observable discrete time stochastic hybrid systems', *Automatica* **50**(8), 1989–1998.

Lesser, K. & Oishi, M. (2015*a*), Computing probabilistic viable sets for partially observable systems using truncated gaussians and adaptive gridding, *in* 'Proceedings of the American Control Conference (ACC)', pp. 1505–1512.

Lesser, K. & Oishi, M. (2015*b*), Finite state approximation for verification of partially observable stochastic hybrid systems, *in* 'Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control', pp. 159–168.

Lesser, K. & Oishi, M. (2016), 'Approximate safety verification and control of partially observable stochastic hybrid systems', *IEEE Transactions on Automatic Control* **62**(1), 81–96.

Liu, S., Trivedi, A., Yin, X. & Zamani, M. (2022), 'Secure-by-construction synthesis of cyber-physical systems', *Annual Reviews in Control* .

Liu, Y.-J., Lu, S., Tong, S., Chen, X., Chen, C. P. & Li, D.-J. (2018), 'Adaptive control-based barrier Lyapunov functions for a class of stochastic nonlinear systems with full state constraints', *Automatica* **87**, 83–93.

Majumdar, R., Mallik, K. & Soudjani, S. (2020), Symbolic controller synthesis for büchi specifications on stochastic systems, *in* 'Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control', pp. 1–11.

Maler, O. & Nickovic, D. (2004), Monitoring temporal properties of continuous signals, *in* 'Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems', Springer, pp. 152–166.

Maler, O., Nickovic, D. & Pnueli, A. (2005), Real time temporal logic: Past, present, future, *in* 'International Conference on Formal Modeling and Analysis of Timed Systems', Springer, pp. 2–16.

Maler, O., Nickovic, D. & Pnueli, A. (2008), Checking temporal properties of discrete, timed and continuous behaviors, *in* 'Pillars of computer science', Springer, pp. 475–505.

Maler, O. & Pnueli, A. (1995), Timing analysis of asynchronous circuits using timed automata, *in* 'Advanced Research Working Conference on Correct Hardware Design and Verification Methods', Springer, pp. 189–205.

Mallik, K., Soudjani, S., Schmuck, A.-K. & Majumdar, R. (2017), Compositional construction of finite state abstractions for stochastic control systems, *in* 'Proceedings of the 56th IEEE International Conference on Decision and Control (CDC)', pp. 550–557.

Mesbah, A. (2016), 'Stochastic model predictive control: An overview and perspectives for future research', *IEEE Control Systems Magazine* **36**(6), 30–44.

Meyer, P. J., Girard, A. & Witrant, E. (2017), 'Compositional abstraction and safety synthesis using overlapping symbolic models', *IEEE Transactions on Automatic Control* **63**(6), 1835–1841.

Meyn, S. P. & Tweedie, R. L. (1993), *Markov chains and stochastic stability*, Comm. Control Engrg., Springer, London.

Mitchell, I. M. (2007), 'A toolbox of level set methods', *UBC Department of Computer Science Technical Report TR-2007-11* .

Mohajerin Esfahani, P., Chatterjee, D. & Lygeros, J. (2015), 'Motion planning for continuous-time stochastic processes: A dynamic programming approach', *IEEE Transactions on Automatic Control* **61**(8), 2155–2170.

Mohajerin Esfahani, P., Chatterjee, D. & Lygeros, J. (2016), 'The stochastic reach-avoid problem and set characterization for diffusions', *Automatica* **70**, 43–56.

Molyneux, G. & Abate, A. (2020), ABC(SMC)$^2$: Simultaneous inference and formal verification, *in* 'Proceedings of CMSB, LNCS 12314', pp. 255–279.

Nejati, A., Soudjani, S. & Zamani, M. (2020*a*), 'Compositional construction of control barrier certificates for large-scale stochastic switched systems', *IEEE Control Systems Letters* **4**(4), 845–850.

Nejati, A., Soudjani, S. & Zamani, M. (2020*b*), 'Compositional construction of control barrier functions for networks of continuous-time stochastic systems', *IFAC-PapersOnLine* **53**(2), 1856–1861.

Nejati, A., Soudjani, S. & Zamani, M. (2021), 'Compositional abstraction-based synthesis for continuous-time stochastic hybrid systems', *European Journal of Control* **57**, 82–94.

Oksendal, B. (2013), *Stochastic differential equations: an introduction with applications*, Springer Science & Business Media.

Pakniyat, A. & Caines, P. E. (2016), On the stochastic minimum principle for hybrid systems, *in* '55th IEEE Conference on Decision and Control', pp. 1139–1144.

Panangaden, P. (2009), *Labelled Markov Processes*, Imperial College Press.

Papachristodoulou, A., Anderson, J., Valmorbida, G., Prajna, S., Seiler, P. & Parrilo, P. (2013), 'SOSTOOLS version 3.00 sum of squares optimization toolbox for MATLAB', *arXiv:1310.4716* .

Park, J., Kurt, A. & Özgüner, Ü. (2014), 'Hybrid systems modeling and reachability-based controller design methods for vehicular automation', *Unmanned Systems* **2**(02), 101–119.

Parrilo, P. A. (2003), 'Semidefinite programming relaxations for semialgebraic problems', *Mathematical Programming* **96**(2), 293–320.

Pilch, C., Edenfeld, F. & Remke, A. (2017), Hypeg: Statistical model checking for hybrid petri nets: Tool paper, *in* 'Proceedings of the 11th EAI International Conference on Performance Evaluation Methodologies and Tools', pp. 186–191.

Pnueli, A. (1977), The temporal logic of programs, *in* 'Proceedings of the 18th Annual Symposium on Foundations of Computer Science', pp. 46–57.

Pola, G., Bujorianu, M. L., Lygeros, J. & Benedetto, M. D. D. (2003), Stochastic hybrid models: An overview, *in* 'IFAC Conference on Analysis and Design of Hybrid Systems', Vol. 36-6, pp. 45–50.

Pola, G. & Pola, G. (2006), Optimal dynamic asset allocation: A stochastic invariance approach, *in* '45th IEEE Conference on Decision and Control', pp. 2589–2594.

Prajna, S., Jadbabaie, A. & Pappas, G. J. (2007), 'A framework for worst-case and stochastic safety verification using barrier certificates', *IEEE Transactions on Automatic Control* **52**(8), 1415–1428.

Prajna, S. & Rantzer, A. (2005), 'On the necessity of barrier certificates', *IFAC Proceedings Volumes* **38**(1), 526–531.

Prandini, M. & Hu, J. (2008), Application of reachability analysis for stochastic hybrid systems to aircraft conflict prediction, *in* '2008 47th IEEE Conference on Decision and Control', pp. 4036–4041.

Raghunathan, A. & Jha, N. K. (2011), Hijacking an insulin pump: Security attacks and defenses for a diabetes therapy system, *in* '2011 IEEE 13th International Conference on e-Health Networking, Applications and Services', pp. 150–156.

Ramponi, F., Chatterjee, D., Summers, S. & Lygeros, J. (2010), On the connections between PCTL and dynamic programming, *in* 'Proceedings of the 13th ACM International Conference on Hybrid Systems: Computation and Control', pp. 253–262.

Roy, D., Giacobbe, M. & Abate, A. (2021), Learning probabilistic termination proofs, *in* 'Proceedings of CAV21, LNCS 12760', p. 3–26.

Rungger, M. & Zamani, M. (2016), SCOTS: A tool for the synthesis of symbolic controllers, *in* 'Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control', pp. 99–104.

Salamati, A., Soudjani, S. & Zamani, M. (2020), 'Data-driven verification under signal temporal logic constraints', *IFAC-PapersOnLine* **53**(2), 69–74.

Salamati, A., Soudjani, S. & Zamani, M. (2021), 'Data-driven verification of stochastic linear systems with signal temporal logic constraints', *Automatica* **131**.

Santoyo, C., Dutreix, M. & Coogan, S. (2019), Verification and control for finite-time safety of stochastic systems via barrier functions, *in* 'Proceedings of the IEEE Conference on Control Technology and Applications', pp. 712–717.

Segala, R. & Lynch, N. (1995), 'Probabilistic simulations for probabilistic processes', *Nordic Journal of Computing* **2**(2), 250–273.

Sen, K., Viswanathan, M. & Agha, G. (2005), On statistical model checking of stochastic systems, *in* 'International Conference on Computer Aided Verification', pp. 266–280.

Shmarov, F., Paoletti, N., Bartocci, E., Lin, S., Smolka, S. A. & Zuliani, P. (2017), 'Automated synthesis of safe and robust PID controllers for stochastic hybrid systems', *Proceedings of the Haifa Verification Conference* pp. 131–146.

Shmarov, F., Soudjani, S., Paoletti, N., Bartocci, E., Lin, S., Smolka, S. A. & Zuliani, P. (2020), 'Automated synthesis of safe digital controllers for sampled-data stochastic nonlinear systems', *IEEE Access* **8**, 180825–180843.

Shmarov, F. & Zuliani, P. (2015), ProbReach: Verified probabilistic delta-reachability for stochastic hybrid systems, *in* 'Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control', pp. 134–139.

Shmarov, F. & Zuliani, P. (2016), Probabilistic hybrid systems verification via SMT and Monte Carlo techniques, *in* 'Proceedings of the Haifa Verification Conference', pp. 152–168.

Silver, D. et al. (2014), Deterministic policy gradient algorithms, *in* 'Proceedings of the 31st International Conference on International Conference on Machine Learning', pp. 387–395.

Singh, A. & Hespanha, J. P. (2010*a*), 'Approximate moment dynamics for chemically reacting systems', *IEEE Transactions on Automatic Control* **56**(2), 414–418.

Singh, A. & Hespanha, J. P. (2010*b*), 'Stochastic hybrid systems for studying biochemical processes', *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **368**(1930), 4995–5011.

Skalse, J., Hammond, L., Griffin, C. & Abate, A. (2022), Lexicographic multi-objective reinforcement learning, *in* 'Proceedings of IJCAI-ECAI22', pp. 3430–3436.

Smith, H. L. (2008), 'Global stability for mixed monotone systems', *Journal of Difference Equations and Applications* **14**(10-11), 1159–1164.

Solar-Lezama, A., Tancau, L., Bodik, R., Seshia, S. & Saraswat, V. (2006), 'Combinatorial sketching for finite programs', *ACM Sigplan Notices* **41**(11), 404–415.

Soltani, M. & Singh, A. (2017), 'Moment-based analysis of stochastic hybrid systems with renewal transitions', *Automatica* **84**, 62–69.

Soudjani, S. (2014), Formal Abstractions for Automated Verification and Synthesis of Stochastic Systems, PhD thesis, Technische Universiteit Delft, The Netherlands.

Soudjani, S. & Abate, A. (2012*a*), Higher-order approximations for verification of stochastic hybrid systems, *in* 'Automated Technology for Verification and Analysis', Vol. 7561 of *Lecture Notes in Computer Science*, Springer, pp. 416–434.

Soudjani, S. & Abate, A. (2012*b*), Probabilistic invariance of mixed deterministic-stochastic dynamical systems, *in* 'ACM Proceedings of the 15th International Conference on Hybrid Systems: Computation and Control', pp. 207–216.

Soudjani, S. & Abate, A. (2013), 'Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes', *SIAM Journal on Applied Dynamical Systems* **12**(2), 921–956.

Soudjani, S. & Abate, A. (2014*a*), Precise approximations of the probability distribution of a Markov process in time: an application to probabilistic invariance, *in* 'Proceedings of TACAS14, LNCS 8413', Springer Verlag, pp. 547–561.

Soudjani, S. & Abate, A. (2014*b*), 'Probabilistic reach-avoid computation for partially-degenerate stochastic processes', *IEEE Transactions on Automatic Control* **59**(2), 528–534.

Soudjani, S. & Abate, A. (2015), 'Quantitative approximation of the probability distribution of a Markov process by formal abstractions', *Logical Methods in Computer Science* **11**(3).

Soudjani, S., Abate, A. & Majumdar, R. (2015), Dynamic Bayesian networks as formal abstractions of structured stochastic processes, *in* 'Proceedings of the 26th International Conference on Concurrency Theory', pp. 1–14.

Soudjani, S., Abate, A. & Majumdar, R. (2017), 'Dynamic Bayesian networks for formal verification of structured stochastic processes', *Acta Informatica* **54**(2), 217–242.

Soudjani, S., Gerwinn, S., Ellen, C., Fränzle, M. & Abate, A. (2014), Formal synthesis and validation of inhomogeneous thermostatically controlled loads, *in* 'Proceedings of the International Conference on Quantitative Evaluation of Systems', pp. 57–73.

Soudjani, S., Gevaerts, C. & Abate, A. (2015), FAUST$^2$: Formal abstractions of uncountable-state stochastic processes, *in* 'Proceedings of TACAS'15', Vol. 9035 of *Lecture Notes in Computer Science*, Springer, pp. 272–286.

Soudjani, S., Majumdar, R. & Nagapetyan, T. (2017), Multilevel Monte Carlo method for statistical model checking of hybrid systems, *in* 'International Conference on Quantitative Evaluation of Systems', pp. 351–367.

Sprinkle, J., Miller, R., Shakernia, O. & Sastry, S. (2005), Using the hybrid systems interchange format to input design models to verification & validation tools, *in* 'IEEE Aerospace Conference', pp. 1–6.

Steinhardt, J. & Tedrake, R. (2012), 'Finite-time regional verification of stochastic non-linear systems', *The International Journal of Robotics Research* **31**(7), 901–923.

Sturm, J. F. (1999), 'Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones', *Optimization methods and software* **11**(1-4), 625–653.

Summers, S. & Lygeros, J. (2010), 'Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem', *Autom.* **46**(12), 1951–1961.

Tabuada, P. (2009), *Verification and control of hybrid systems: A symbolic approach*, Springer Science & Business Media.

Teel, A. R. (2013), 'Lyapunov conditions certifying stability and recurrence for a class of stochastic hybrid systems', *Annual Reviews in Control* **37**(1), 1–24.

Teel, A. R., Subbaraman, A. & Sferlazza, A. (2014), 'Stability analysis for stochastic hybrid systems: A survey', *Automatica* **50**(10), 2435–2456.

Tkachev, I. & Abate, A. (2011), On infinite-horizon probabilistic properties and stochastic bisimulation functions, *in* 'Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)', pp. 526–531.

Tkachev, I. & Abate, A. (2012*a*), Regularization of Bellman equations for infinite-horizon probabilistic properties, *in* 'Proceedings of the 15th ACM international conference on Hybrid Systems: computation and control', pp. 227–236.

Tkachev, I. & Abate, A. (2012*b*), Stability and attractivity of absorbing sets for discrete-time Markov processes, *in* 'Proceedings of the 51st IEEE Conference on Decision and Control', pp. 7652–7657.

Tkachev, I. & Abate, A. (2014), 'Characterization and computation of infinite horizon specifications over Markov processes', *Theoretical Computer Science* **515**, 1–18.

Tkachev, I., Mereacre, A., Katoen, J.-P. & Abate, A. (2013), Quantitative automata-based controller synthesis for non-autonomous stochastic hybrid systems, *in* 'Proceedings of the 16th ACM International Conference on Hybrid Systems: Computation and Control', pp. 293–302.

Tkachev, I., Mereacre, A., Katoen, J.-P. & Abate, A. (2017), 'Quantitative model-checking of controlled discrete-time Markov processes', *Information and Computation* **253**, 1 – 35.

van Breugel, F., Tang, Q., Mardare, R., Larsen, K. G., Bacci, G. & Bacci, G. (2021), 'Computing probabilistic bisimilarity distances for probabilistic automata', *Logical Methods in Computer Science* **17**.

van Schuppen, J. H. (1989), Stochastic realization problems, *in* 'Three decades of mathematical system theory', Springer, pp. 480–523.

Vardi, M. Y. (1985), Automatic verification of probabilistic concurrent finite state programs, *in* '26th Annual Symposium on Foundations of Computer Science (SFCS)', pp. 327–338.

Vargas-García, C. A. & Singh, A. (2018), Elucidating cell size control mechanisms with stochastic hybrid systems, *in* '2018 IEEE Conference on Decision and Control (CDC)', pp. 4366–4371.

Vinod, A. & Oishi, M. M. (2018), Scalable underapproximative verification of stochastic LTI systems using convexity and compactness, *in* 'Proceedings of the 21st International Conference on Hybrid Systems: Computation and Control', pp. 1–10.

Vinod, A. P., Gleason, J. D. & Oishi, M. M. (2019), SReachTools: A MATLAB stochastic reachability toolbox, *in* 'Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control', pp. 33–38.

Vinod, A. P., HomChaudhuri, B. & Oishi, M. M. (2017), Forward stochastic reachability analysis for uncontrolled linear systems using fourier transforms, *in* 'Proceedings of the 20th ACM International Conference on Hybrid Systems: Computation and Control', pp. 35–44.

Vinod, A. P. & Oishi, M. M. (2017), 'Scalable underapproximation for the stochastic reach-avoid problem for high-dimensional LTI systems using fourier transforms', *IEEE control systems letters* **1**(2), 316–321.

Vinod, A. P. & Oishi, M. M. (2021), 'Stochastic reachability of a target tube: Theory and computation', *Automatica* **125**, 109458.

Wang, Q., Zuliani, P., Kong, S., Gao, S. & Clarke, E. M. (2015), SReach: A probabilistic bounded delta-reachability analyzer for stochastic hybrid systems, *in* 'Proceedings of the International Conference on Computational Methods in Systems Biology', pp. 15–27.

Wang, Y., Roohi, N., West, M., Viswanathan, M. & Dullerud, G. E. (2016), Verifying continuous-time stochastic hybrid systems via Mori-Zwanzig model reduction, *in* '2016 IEEE 55th conference on decision and control (CDC)', pp. 3012–3017.

Wisniewski, R., Bujorianu, M. L. & Sloth, C. (2020), 'p-safe analysis of stochastic hybrid processes', *IEEE Transactions on Automatic Control* **65**(12), 5220–5235.

Wisniewski, R. & Sloth, C. (2015), 'Converse barrier certificate theorems', *IEEE Transactions on Automatic Control* **61**(5), 1356–1361.

Wongpiromsarn, T., Topcu, U. & Lamperski, A. (2015), 'Automata theory meets barrier certificates: Temporal logic verification of nonlinear systems', *IEEE Transactions on Automatic Control* **61**(11), 3344–3355.

Wu, B., Ahmadi, M., Bharadwaj, S. & Topcu, U. (2019), Cost-bounded active classification using partially observable Markov decision processes, *in* 'Proceedings of the American Control Conference (ACC)', pp. 1216–1223.

Yu, L., Cheng, X., Scherpen, J. M. & Gort, E. (2019), $H_2$ sub-optimal model reduction for second-order network systems, *in* '2019 IEEE 58th Conference on Decision and Control (CDC)', pp. 5062–5067.

Yu, L., Cheng, X., Scherpen, J. & Xiong, J. (2022), '$H_2$ model reduction for diffusively coupled second-order networks by convex-optimization', *Automatica* **137**.

Yurtsever, A., Tropp, J. A., Fercoq, O., Udell, M. & Cevher, V. (2021), 'Scalable semidefinite programming', *SIAM Journal on Mathematics of Data Science* **3**(1), 171–200.

Zacchia Lun, Y., Wheatley, J., D'Innocenzo, A. & Abate, A. (2018), 'Approximate abstractions of Markov chains with interval decision processes', *Proceedings of the 6th IFAC Conference on Analysis and Design of Hybrid Systems* **51**(16), 91–96.

Zamani, M. (2014), Compositional approximations of interconnected stochastic hybrid systems, *in* 'Proceedings of the 53rd IEEE Conference on Decision and Control (CDC)', pp. 3395–3400.

Zamani, M. & Abate, A. (2014*a*), 'Approximately bisimilar symbolic models for randomly switched stochastic systems', *IEEE Control Systems Letters* **69**, 38–46.

Zamani, M. & Abate, A. (2014*b*), 'Symbolic models for randomly switched stochastic systems', *Systems & Control Letters* **69**, 38–46.

Zamani, M., Abate, A. & Girard, A. (2015), 'Symbolic models for stochastic switched systems: A discretization and a discretization-free approach', *Automatica* **55**(5), 183–196.

Zamani, M. & Arcak, M. (2018), 'Compositional abstraction for networks of control systems: A dissipativity approach', *IEEE Transactions on Control of Network Systems* **5**(3), 1003–1015.

Zamani, M., Mohajerin Esfahani, P., Majumdar, R., Abate, A. & Lygeros, J. (2014), 'Symbolic control of stochastic systems via approximately bisimilar finite abstractions', *IEEE Transactions on Automatic Control* **59**(12), 3135–3150.

Zamani, M., Rungger, M. & Mohajerin Esfahani, P. (2017), 'Approximations of stochastic hybrid systems: A compositional approach', *IEEE Transactions on Automatic Control* **62**(6), 2838–2853.

Zamani, M., Tkachev, I. & Abate, A. (2014), Bisimilar symbolic models for stochastic control systems without state-space discretization, *in* 'Proceedings of the 17th international conference on hybrid systems: computation and control', pp. 41–50.

Zamani, M., Tkachev, I. & Abate, A. (2017), 'Towards scalable synthesis of stochastic control systems', *Discrete Event Dynamic Systems* **27**(2), 341–369.

Zames, G. (1966), 'On the input-output stability of time-varying nonlinear feedback systems part one: Conditions derived using concepts of loop gain, conicity, and positivity', *IEEE Transactions on Automatic Control* **11**(2), 228–238.

Zhang, L., She, Z., Ratschan, S., Hermanns, H. & Hahn, E. M. (2010), Safety verification for probabilistic hybrid systems, *in* 'International Conference on Computer Aided Verification', pp. 196–211.

Zhang, W., Prabhakar, P. & Natarajan, B. (2017), Abstraction based reachability analysis for finite branching stochastic hybrid systems, *in* 'Proceedings of the 8th International Conference on Cyber-Physical Systems', pp. 121–130.

[1]School of Computing, Newcastle University, United Kingdom

*Email address*: {abolfazl.lavaei,sadegh.soudjani}@newcatle.ac.uk

[2]Department of Computer Science, University of Oxford, United Kingdom

*Email address*: aabate@cs.ox.ac.uk

[3]Department of Computer Science, University of Colorado Boulder, USA

[4]Department of Computer Science, LMU Munich, Germany

*Email address*: majid.zamani@colorado.edu