2024 IEEE 63rd Conference on Decision and Control (CDC)
December 16-19, 2024. MiCo, Milan, Italy

Reactive Synthesis of Stochastic Control Systems: A Mode-triggered Safety Barrier Approach

Amy Nejati and Anne-Kathrin Schmuck

Abstract— **This paper develops a formal framework to systematically synthesize a controller for continuous-time nonlinear stochastic control systems which react to specification-mode changes, initiated by either the external environment or the system itself. Considering the particular challenging setting where mode switches affect the specification rather than the dynamics, our proposed scheme adopts a synthesis approach based on *control barrier certificates* to synthesize controllers that ensure compliance with *mode-triggered safety specifications*. Our method leverages the computational capabilities derived from state-space control techniques and combines them with the reactivity of logic control. We provide a robotic case study to illustrate the effectiveness of our proposed approach.**

I. Introduction

There has been a significant surge, over the past few years, in the adoption of formal control synthesis for complex dynamical systems, aiming to enforce safety classifications. However, a formal synthesis approach usually faces substantial challenges arising from some factors including (i) the continuous nature of state and input sets which demands specialized techniques for designing controllers capable of effectively navigating such continuous domains, and (ii) the inherent stochasticity in the underlying dynamics which introduces an additional layer of complexity. This becomes particularly crucial because, in specific scenarios, a physical system must adapt to specification changes to satisfy a given formal property. These shifts in specification modes represent logical context changes triggered by external environmental factors and can occur at various moments.

To tackle the challenges posed by the aforesaid computational complexity, one promising approach outlined in pertinent literature involves approximating complex models with simpler ones that possess *finite* state and input sets, commonly referred to as finite abstractions [1], [2]. In practical application, constructing such finite abstractions entails dividing the state and input sets of concrete models based on predefined discretization parameters [3]–[11]. However, a notable limitation of employing finite-abstraction techniques lies in their discretization-based nature, leading to susceptibility due to an *exponential curse of dimensionality*.

To overcome the difficulties associated with the state-explosion problem in abstraction-based approaches, a potential approach involves leveraging *barrier certificates* as a method free from discretization, facilitating the formal analysis of complex dynamical systems. Barrier certificates essentially are Lyapunov-like functions, spanning the system's state space while satisfying a set of inequalities pertaining to the function itself and its time derivative along the system's flow. The presence of such a function inherently provides a probabilistic guarantee over the safety of the system (see [12]–[19]). Expanding upon the aforementioned barrier concepts serves as a valuable tool in addressing other properties, including *reachability* specifications [20]. It is worth mentioning that, although abstraction-based techniques suffer from state explosion, they offer an advantage over barrier certificates in terms of their automated verification and synthesis process, which can also benefit from the inherent parallelism of abstraction methods [21].

Although prior work extensively covers finite abstraction and barrier certificate approaches for control systems, the consideration of *reactivity to environment-triggered logical specification changes* within these control frameworks has only recently been considered [22]–[25]. While [22]–[24] only consider non-stochastic dynamics, [25] focuses on scenarios where specification changes only occur *sequentially*. In our current work, we study safety specifications while allowing specification-mode changes to occur anytime – requiring the system to instantaneously react to the mode change by dynamically imposing new safety barriers—an aspect not addressed in the study in [25].

Given an operational scenario, our key objective is to synthesize a feedback control policy capable of responding to the external environment's *mode* choices m such that the dynamical system (e.g., a robot) avoids the *obstacle* O_m *promptly* with *high confidence* while handling the uncertainties in the system in a provably correct way. This poses several major challenges on the logical (higher) control level. The *first one* stems from the fact that a mode-change can be triggered at any point in time, requiring the system to immediately react to it. The *second challenge* arises from the scenario where the system might be within an obstacle when a new mode *aims to be activated*. To address this, a reachability analysis can be performed *before activating the mode* and deploying new safety constraints, ensuring that the system can safely exit the obstacle region and maintain the integrity of the safety certificates.

To implement such logical control objectives over stochastic dynamical systems, our framework employs mode-dependent *safety and reachability* barrier certificates to derive feedback control policies realizing the above-mentioned mode-dependent control tasks. This, however, results in a *third challenge* stemming from the fact that the probabilistic guarantees offered by barrier certificates inherently correlate with particular initial sets. As modes can change at any time,

A. Nejati is with the School of Computing, Newcastle University, UK. A.-K. Schmuck is with the Max Planck Institute for Software Systems, Kaiserslautern, Germany. Email: amy.nejati@newcastle.ac.uk, akschmuck@mpi-sws.org. The work is partially supported by DFG, Germany within projects 389792660 TRR 248–CPEC and SCHM 3541/1-1.

the entire state space might be initial for a subsequently activated barrier. To address this, we employ a coarse mode-dependent state-space discretization, which allows us to consider each resulting cell as an individual initial set. This results in a so-called *mode-dependent heatmap*, which associates each cell with a guaranteed probability of satisfying the safety property when the system starts from that cell. This allows us to construct a hybrid control policy which ensures the overall safety in a reactive and reliable manner.

It is worth noting that in our proposed framework, modes are mutually exclusive, with only one active mode at a time, eliminating the possibility of mode conflicts or overlaps. We do not impose specific time bounds, like dwell-time requirements, for exiting unsafe regions. Instead, we focus on ensuring provable safety at all times. While time-based constraints could be added, our approach prioritizes immediate safe behavior, making it ideal for environments where safety is more important than time limits. In addition, the mode transitions, triggered by external factors like task assignments or environmental changes, are considered *deterministic* in our setting. If transitions are stochastic (e.g., Markovian switching) [19], modeling transition probabilities can refine safety guarantees. Due to space constraints, proofs of some statements are omitted.

II. System Description

Notation. Sets of nonnegative integers $\mathbb{N}_0 := \{0, 1, 2, \ldots\}$ and positive integers $\mathbb{N} := \{1, 2, 3, \ldots\}$ are denoted by these symbols, respectively. The symbols $\mathbb{R}$, $\mathbb{R}^+$, and $\mathbb{R}_0^+$ represent sets of real numbers, positive real numbers, and nonnegative real numbers, respectively. When considering a matrix $A \in \mathbb{R}^{m\times m}$, the symbol $\mathsf{Tr}(A)$ signifies the trace of A, which corresponds to the sum of all its diagonal elements. We utilize $x = [x_1; \ldots; x_N]$ to symbolize the vector associated with a dimension $\sum_i n_i$, given N vectors $x_i \in \mathbb{R}^{n_i}$, with $n_i \in \mathbb{N}$ and $i \in \{1, \ldots, N\}$. The identity matrix of size n is denoted as $\mathbb{I}_n$. We define $b(X)$ and $cl(X)$ for a set $X \in \mathbb{R}^n$ as its *boundary* and topological *closure*, respectively. We express the relationship between a system Σ and a property Ψ as $\Sigma \vDash \Psi$, signifying that Σ satisfies Ψ.

Preliminaries. We work within a probability space $(\Omega, \mathsf{F}_\Omega, \mathbb{P}_\Omega)$, where Ω represents the sample space, F_Ω is a sigma-algebra on Ω encompassing subsets of Ω as events, and $\mathbb{P}_\Omega$ is a probability measure assigning probabilities to events. We assume that the triple $(\Omega, \mathsf{F}_\Omega, \mathbb{P}_\Omega)$ is equipped with a filtration $\mathbb{F} = (\mathsf{F}_s)_{s\geq 0}$ meeting the usual conditions of completeness and right continuity. Within this space, we denote an r-dimensional $\mathbb{F}$-Brownian motion by $(\mathbb{W}_s)_{s\geq 0}$.

Continuous-Time Stochastic Control Systems. The focus of this work lies in the study of continuous-time nonlinear stochastic control systems (ct-SCS), aligning with the formal definition outlined subsequently.

Definition 2.1: A continuous-time nonlinear stochastic control system (ct-SCS) is described by the tuple

$$\Sigma = (X, U, \mathcal{U}, f, \delta), \tag{1}$$

where:

- $X \subseteq \mathbb{R}^n$ denotes the state set of ct-SCS;
- $U \subseteq \mathbb{R}^{\bar{m}}$ denotes the input set of ct-SCS;
- $\mathcal{U}$ is a subset of sets of $\mathbb{F}$-progressively measurable processes [26] taking values in $\mathbb{R}^{\bar{m}}$;
- $f\colon X \times U \to X$ represents the drift term which is globally Lipschitz continuous: there exist constants $\mathscr{L}_x, \mathscr{L}_u, \in \mathbb{R}_{\geq 0}$ such that $\|f(x,u) - f(x',u')\| \leq \mathscr{L}_x\|x - x'\| + \mathscr{L}_u\|u - u'\|$ for all $x, x' \in X$, and for all $u, u' \in U$;
- $\delta\colon \mathbb{R}^n \to \mathbb{R}^{n\times \mathsf{r}}$ represents the diffusion term which is globally Lipschitz continuous.

A continuous-time stochastic control system Σ fulfills

$$\mathsf{d}x(t) = f(x(t), \nu(t))\,\mathsf{d}t + \delta(x(t))\,\mathsf{d}\mathbb{W}_t, \tag{2}$$

$\mathbb{P}$-almost surely ($\mathbb{P}$-a.s.), where the stochastic process $x : \Omega\times\mathbb{R}_0^+ \to X$ is the *solution process* of Σ. Describing a state-feedback policy for ct-SCS, denoted by (2), entails a function $\nu : X \to \mathcal{U}$. To represent the value of the solution process at a specific time $t \in \mathbb{R}_0^+$ under policy ν, originating from an initial condition $x_0 = x(0)$ $\mathbb{P}$-a.s., we utilize $x_{x_0\nu}(t)$, where x_0 stands as an F_0-measurable random variable.

Safety and Reachability Control Problems. Consider the continuous-time stochastic system depicted in (2), alongside sets $\mathsf{I}, \mathsf{S}, \mathsf{T} \subseteq X$, denoted as the "initial", "safe", and "target" sets, respectively. In this context, a bounded-horizon *safety property* asserts that a solution process $x_{x_0\nu}$ starting from an initial condition $x_0 \in \mathsf{I}$ under policy ν stays in the safe set S within the time interval $[0, \mathcal{T}] \subset \mathbb{R}_0^+$. This requirement can be extended to infinite time horizons as $\mathcal{T} \to \infty$ signifying "*always* S". The primary goal when computing a feedback-policy ν for this objective is to ensure that the probability that a solution process remains in S within an infinite horizon is lower bounded by $\theta_1 \in [0, 1]$ (cf. Theorem 4.2), *i.e.,*

$$\mathbb{P}\Big\{x_{x_0\nu}(t) \models (\mathsf{I}, \mathsf{S}) \text{ for all } t \in \mathbb{R}_0^+ \,|\, x_0 \in \mathsf{I}\Big\} \geq \theta_1.$$

Similarly, a bounded-horizon *reachability property* stipulates that a solution process $x_{x_0\nu}$ starting from an initial condition $x_0 \in \mathsf{I}$ under policy ν reaches the target set T within the time interval $[0, \mathcal{T}] \subset \mathbb{R}_0^+$ if there exists a time instant $t \in [0, \mathcal{T}]$ for which $x(t) \in \mathsf{T}$. This requirement can be extended to infinite time horizons as $\mathcal{T} \to \infty$ signifying "*eventually* T". The main goal when computing a feedback-policy ν is to ensure that the probability that a solution process reaches T starting from I within an infinite horizon is lower bounded by $\theta_2 \in [0, 1]$ (cf. Theorem 4.5), *i.e.,*

$$\mathbb{P}\Big\{x_{x_0\nu}(t) \models (\mathsf{I}, \mathsf{T}) \text{ for some } t \in \mathbb{R}_0^+ \,|\, x_0 \in \mathsf{I}\Big\} \geq \theta_2.$$

III. Hybrid Control Interaction

The aim of this paper is to synthesize a hybrid safety controller for a ct-SCS Σ which reacts to the type of specification changes. We achieve this by a separation of concerns: while the (high-level) *reactive decisions* are formalized over a *mode-dependent discrete partition* of the state space, the underlying ct-SCS is actuated via (switching) feedback policies $\nu : X \to \mathcal{U}$ in continuous time and space. In particular, every (externally triggered) mode-change is translated into a reachability and a safety control problem, implemented through policies based on control barrier certificates.

Mode-dependent Heat Maps. Towards a formalization of the above intuition, we first define mode-dependent heat

maps. To this end, let I_m be a finite index set for mode $m \in M$. Then $Q_m = \{q_i \subseteq X \mid i \in I_m\}$ is a finite partition of X s.t. $\bigcup_{i\in I_m} q_i = X$ and $q_i \cap q_j = \emptyset$ for all $i, j \in I_m$, $i \neq j$. A heat map for mode m is a tuple (Q_m, Γ_m), where $\Gamma_m : q_i \to [0, 1]$ associates each element in Q_m with a corresponding probability.

We call the heat map (Q_m, Γ_m) *implementable* for ct-SCS Σ with respect to the safety specification induced by the mode-dependent obstacles O_m, if for any $i \in I_m$ there exists a feedback policy $\nu_{m,i}$ solving the respective safety control problem:

$$\mathbb{P}\Big\{x_{x_0\nu}(t) \models (q_i, X\backslash \mathsf{O}_m) \text{ for all } t\in\mathbb{R}_0^+ \,\big|\, x_0 \in q_i, \nu_{m,i}\Big\} \geq \Gamma_m(q_i).$$

Intuitively, $\Gamma_m(q_i)$ is the probability with which the control policy $\nu_{m,i}$ can guarantee the system to stay outside of O_m when starting from some $x_0 \in q_i$. It is clear that $\Gamma_m(q_i) = 0$ if $q_i \cap \mathsf{O}_m \neq \emptyset$.

Reachability Control. Given an *implementable* heat-map (Q_m, Γ_m) for safety specification over the ct-SCS Σ, there might be partitions q_i such that $q_i \cap \mathsf{O}_m \neq \emptyset$. In this case, we consider an additional *reachability* control problem which first navigates the system to the partition $q_j^{\uparrow} \in Q_m$ with the *highest safety probability* and only upon reaching $q_j^{\uparrow}$ activates the associated safety control policy $\nu_{m,j}^{\uparrow}$. This requires us to construct a control policy $\tilde{\nu}_{m,i}$ such that the following reachability control problem has a solution:

$$\mathbb{P}\Big\{x_{x_0\tilde{\nu}}(t) \models (q_i, q_j^{\uparrow}) \text{ for some } t\in\mathbb{R}_0^+ \mid x_0 \in q_i, \tilde{\nu}_{m,i}\Big\} \geq \tilde{\theta}_i.$$

Hybrid Control Policies. With this, for any partition $q_i \in Q_m$, $q_i \cap \mathsf{O}_m \neq \emptyset$, we define a hybrid policy $\overline{\nu}_{m,i} := \tilde{\nu}_{m,i} \cdot \nu_{m,j}^{\uparrow}$ which ensures that $\nu_{m,j}^{\uparrow}$ is activated when $q_j^{\uparrow}$ is entered, which is guaranteed by the construction of $\tilde{\nu}_{m,i}$. For all other partitions $q_i \in Q_m$, probabilistic safety is ensured by defining $\overline{\nu}_{m,i} := \nu_{m,i}$. In other words, the overall hybrid safety controller triggers policy $\nu_{m,i}$ whenever mode m is activated while the system is in region q_i.

Problem Statement. Given the above discussion, the problem addressed in this paper can be summarized as follows.

Problem 3.1: Let Σ be a ct-SCS with state space X and $\mathsf{O}_m \subseteq X$, with $m \in M$ being a finite set of mode-dependent obstacles. For every mode $m \in M$, compute a heat map (Q_m, Γ_m) and associated hybrid policies $\overline{\nu}_{m,i}$ such that the system under control is guaranteed to eventually stay safe for all modes with an overall probability $\theta \in [0, 1]$ (see Sec. V).

IV. Control Barrier Certificates

In this section, we begin by introducing two fundamental concepts of control barrier certificates for ct-SCS, to address, respectively, the *safety and reachability* problems, specified in Section III. Based on the preceding discussion, we subsequently employ these concepts to establish lower bounds on two key probabilities: one that pertains to preventing the system from entering designated unsafe areas (referred to as the safety problem), while the other concerns the likelihood of the system reaching predefined target regions (referred to as the reachability problem). We slightly abuse the notation and omit the index i for the sake of simplicity.

A. Safety Certificates

Definition 4.1: Given a ct-SCS $\Sigma = (X, U, \mathcal{U}, f, \delta)$, let $\mathsf{I}, X_u \subseteq X$ represent the initial and unsafe sets of ct-SCS, respectively. A twice-differentiable function $\mathcal{B}_s : X \to \mathbb{R}_0^+$ is said to be a *safety* control barrier certificate (S-CBC) for Σ if there exist $\varepsilon_s, \lambda_s \in \mathbb{R}^+$ with $\lambda_s > \varepsilon_s$, such that

$$\forall x \in \mathsf{I}: \qquad \mathcal{B}_s(x) \leq \varepsilon_s, \tag{3}$$

$$\forall x \in X_u: \qquad \mathcal{B}_s(x) \geq \lambda_s, \tag{4}$$

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

$$\mathcal{L}\mathcal{B}_s(x) \leq 0, \tag{5}$$

with $\mathcal{L}\mathcal{B}_s$ being the *infinitesimal generator* of the stochastic process acting on function $\mathcal{B}_s$ [27], defined as

$$\mathcal{L}\mathcal{B}_s(x) = \partial_x \mathcal{B}_s(x) f(x, \nu) + \frac{1}{2}\mathsf{Tr}(\delta(x)\delta(x)^\top \partial_{x,x}\mathcal{B}_s(x)). \tag{6}$$

Through the utilization of the S-CBC outlined in Definition 4.1, the subsequent theorem establishes a quantifiable lower bound on the probability associated with the system's avoidance of specific unsafe regions [28].

Theorem 4.2: Given $\Sigma = (X, U, \mathcal{U}, f, \delta)$, let $\mathcal{B}_s$ designated as an S-CBC for Σ according to Definition 4.1. Then the probability that the solution process of Σ, initiating from any initial state $x_0 \in \mathsf{I}$, under policy $\nu(\cdot)$, never reaches X_u within an infinite time horizon is formally quantified as

$$\mathbb{P}\Big\{x_{x_0\nu}(t) \models (\mathsf{I}, \mathsf{S}) \text{ for all } t \in \mathbb{R}_0^+ \mid x_0\Big\} \geq 1 - \frac{\varepsilon_s}{\lambda_s}, \tag{7}$$

with $\mathsf{S} = X\backslash X_u$.

Remark 4.3: Note that there are no restrictions on the shape of the partition cells within Q_m when computing S-CBC according to Definition 4.1 and constructing the heat map (Q_m, Γ_m). However, the finer partitions result in less conservative S-CBC corresponding to those cells, albeit at the expense of increased computational complexity in computing S-CBC for a larger number of partition cells.

B. Reachability Certificates

As discussed in Section III, when a mode aims to be activated and the system remains within the obstacle of that mode, it necessitates first solving a *reachability problem*, where the target set is defined as the cell with the highest safety probability based on the computed heatmap. To address this issue, inspired by [20], [29], we now define a notion for control barrier certificates to quantify a lower bound on the probability that the continuous-time stochastic control system reaches some target regions.

Definition 4.4: Given a ct-SCS $\Sigma = (X, U, \mathcal{U}, f, \delta)$, let $\mathsf{I}, \mathsf{T} \subseteq X$ be initial and target sets of ct-SCS, respectively. A twice-differentiable function $\mathcal{B}_r : X \to \mathbb{R}_0^+$ is said to be a *reachability* control barrier certificate (R-CBC) for Σ if there exist $\varepsilon_r, \lambda_r, \psi \in \mathbb{R}^+$ with $\lambda_r > \varepsilon_r$, such that

$$\forall x \in \mathsf{I}: \qquad \mathcal{B}_r(x) \leq \varepsilon_r, \tag{8}$$

$$\forall x \in b(X)\backslash b(\mathsf{T}): \qquad \mathcal{B}_r(x) \geq \lambda_r, \tag{9}$$

and $\forall x \in cl(X\backslash \mathsf{T})$, $\exists \tilde{\nu} \in U$, such that

$$\mathcal{LB}_r(x) \leq -\psi, \tag{10}$$

with $\mathcal{LB}_r$ being the *infinitesimal generator* of the stochastic process acting on function $\mathcal{B}_r$, as defined in (6).

We now leverage the R-CBC in Definition 4.4 and quantify a lower bound on the probability that the system reaches target set T in an infinite time horizon [29].

Theorem 4.5: Given $\Sigma = (X, U, \mathcal{U}, f, \delta)$, let $\mathcal{B}_r$ be an R-CBC for Σ as in Definition 4.4. Then the probability that the solution process of Σ, commencing from $x_0 \in \mathsf{I}$ under policy $\tilde{\nu}(\cdot)$, reaches T in an infinite horizon is quantified as

$$\mathbb{P}\Big\{x_{x_0\tilde{\nu}}(t) \models (\mathsf{I}, \mathsf{T}) \text{ for some } t \in \mathbb{R}_0^+ \mid x_0\Big\} \geq 1 - \frac{\varepsilon_r}{\lambda_r}. \tag{11}$$

C. Computation of Safety and Reachability CBC

Here, we translate the conditions outlined in Definitions 4.1 and 4.4 into an optimization problem leveraging the sum-of-squares (SOS) methodology [30]. We accordingly provide a systematic approach for computing both S-CBC and R-CBC, along with their associated control policies tailored for the system Σ. To effectively employ SOS optimization techniques, we assume that Σ has continuous state and input sets X, U, with polynomial drift and diffusion terms f, δ.

The ensuing lemma provides a set of sufficient conditions validating the existence of an S-CBC and its control policy.

Lemma 4.6: Considering the semi-algebraic nature of sets X, I, X_u, U, defined by vectors of polynomial inequalities $g(x), g_0(x), g_u(x), g_\nu(x) \in \mathbb{R}_0^+$, let the following conditions hold: there exist sum-of-square polynomial $\mathcal{B}_s(x)$, constants $\varepsilon_s, \lambda_s \in \mathbb{R}^+$, polynomials $l_{\nu_z}(x)$ corresponding to the z^{th} input in $\nu_z = (\nu_1, \nu_2, \ldots, \nu_{\bar{m}}) \in U \subseteq \mathbb{R}^{\bar{m}}$, and vectors of sum-of-squares polynomials $l(x)$, $l_0(x)$, $l_u(x)$, $\hat{l}_\nu(x)$ of appropriate dimensions, ensuring the following expressions are sum-of-squares polynomials:

$$-\mathcal{B}_s(x) - l_0^\top(x) g_0(x) + \varepsilon_s, \tag{12}$$

$$\mathcal{B}_s(x) - l_u^\top(x) g_u(x) - \lambda_s, \tag{13}$$

$$-\mathcal{LB}_s(x) - \sum_{z=1}^{\bar{m}} (\nu_z - l_{\nu_z}(x)) - l^\top(x) g(x) - \hat{l}_\nu^\top(x) g_\nu(x). \tag{14}$$

Then, $\mathcal{B}_s(x)$ fulfills conditions (3)-(5) in Definition 4.1 and $\nu_z \geq l_{\nu_z}$, is the corresponding safety controller.

One can use a similar argument as Lemma 4.6 and compute an SOS polynomial R-CBC and its corresponding *reachability* controller according to Definition 4.4, by defining vectors of polynomial inequalities $g_u(x)$ for $b(X)\backslash b(\mathsf{T})$ in condition (9) and $g(x)$ for $cl(X\backslash \mathsf{T})$ in condition (10).

V. Probabilistic Guarantee for Overall Mode-triggered Specification

In this section, we offer our main result as a probabilistic guarantee for the overall mode-triggered specification to address Problem 3.1. In particular, the following theorem provides a safety probability under the assumption that any mode can be triggered at any time, *i.e.,* the following guarantee holds for the system being in any (not previously known) state x_0 and any mode m being triggered at this point. If another mode m' gets triggered at a future state x', the same reasoning holds again based on the location of x' and its associated barrier-based controller.

Theorem 5.1: Consider a ct-SCS $\Sigma = (X, U, \mathcal{U}, f, \delta)$ with state space X and $\mathsf{O}_m \subseteq X$, with $m \in M = \{1, \ldots, p\}$ being a finite set of mode-dependent obstacles. Let there exist S-CBC $\mathcal{B}_s^{m,i}$ (corresponding to heat map (Q_m, Γ_m)) and R-CBC $\mathcal{B}_r^{m,i}$ for each mode according to Definitions 4.1 and 4.4, respectively. Now one can guarantee that the overall mode-triggered specification is satisfied with an overall probability of at least $\theta = 1 - \sum_{m=1}^{p} (\frac{\varepsilon_r^{m,i}}{\lambda_r^{m,i}} + \frac{\varepsilon_s^{m,i}}{\lambda_s^{m,i}})$.

Proof. By defining events

$$\mathcal{R}_{m,i} := \Big\{x_{x_0\tilde{\nu}}(t) \models (q_i, q_j^{\uparrow}) \text{ for some } t \in \mathbb{R}_0^+ \mid x_0 \in q_i, \tilde{\nu}_{m,i}\Big\},$$

$$\mathcal{S}_{m,i} := \Big\{x_{x_0\nu}(t) \models (q_i, X\backslash \mathsf{O}_m) \text{ for all } t \in \mathbb{R}_0^+ \mid x_0 \in q_i, \nu_{m,i}\Big\},$$

one has $\mathbb{P}\{\mathcal{R}_{m,i}\} \geq 1 - \frac{\varepsilon_r^{m,i}}{\lambda_r^{m,i}}$ and $\mathbb{P}\{\mathcal{S}_{m,i}\} \geq 1 - \frac{\varepsilon_s^{m,i}}{\lambda_s^{m,i}}, m \in \{1, \ldots, p\}$. We are interested in occurrences of events $\mathcal{R}_{m,i}$ and $\mathcal{S}_{m,i}$ to ensure the satisfaction of both safety and reachability properties, that can be computed as:

$$\mathbb{P}\{\mathcal{R}_{m,i} \cap \mathcal{S}_{m,i}\} = 1 - \mathbb{P}\{\bar{\mathcal{R}}_{m,i} \cup \bar{\mathcal{S}}_{m,i}\}, \tag{15}$$

where $\bar{\mathcal{R}}_{m,i}$ and $\bar{\mathcal{S}}_{m,i}$ are the complement of $\mathcal{R}_{m,i}$ and $\mathcal{S}_{m,i}$, respectively. Since

$$\mathbb{P}\{\bar{\mathcal{R}}_{m,i} \cup \bar{\mathcal{S}}_{m,i}\} \leq \mathbb{P}\{\bar{\mathcal{R}}_{m,i}\} + \mathbb{P}\{\bar{\mathcal{S}}_{m,i}\},$$

and by leveraging (15), one can conclude that

$$\begin{aligned}\mathbb{P}\{\mathcal{R}_{m,i} \cap \mathcal{S}_{m,i}\} &\geq 1 - \mathbb{P}\{\bar{\mathcal{R}}_{m,i}\} - \mathbb{P}\{\bar{\mathcal{S}}_{m,i}\} \\ &\geq 1 - (\frac{\varepsilon_r^{m,i}}{\lambda_r^{m,i}} + \frac{\varepsilon_s^{m,i}}{\lambda_s^{m,i}}).\end{aligned} \tag{16}$$

We now proceed with showing the overall mode-triggered specification covering the *whole modes*. By defining events $\mathcal{E}_{m,i} = \{\mathcal{R}_{m,i} \cap \mathcal{S}_{m,i}\}$, we have $\mathbb{P}\{\mathcal{E}_{m,i}\} \geq 1 - (\frac{\varepsilon_r^{m,i}}{\lambda_r^{m,i}} + \frac{\varepsilon_s^{m,i}}{\lambda_s^{m,i}}), m \in \{1, \ldots, p\}$ according to (16). We now compute the occurrence of all events $\mathcal{E}_m$ to ensure the satisfaction of the property for all modes, namely $\mathbb{P}\{\mathcal{E}_{1,i} \cap \cdots \cap \mathcal{E}_{p,i}\}$, as:

$$\mathbb{P}\{\mathcal{E}_{1,i} \cap \cdots \cap \mathcal{E}_{p,i}\} = 1 - \mathbb{P}\{\bar{\mathcal{E}}_{1,i} \cup \cdots \cup \bar{\mathcal{E}}_{p,i}\}, \tag{17}$$

where $\bar{\mathcal{E}}_{m,i}$ are complements of $\mathcal{E}_{m,i}, \forall m \in \{1, \ldots, p\}$. Since

$$\mathbb{P}\{\bar{\mathcal{E}}_{1,i} \cup \cdots \cup \bar{\mathcal{E}}_{p,i}\} \leq \mathbb{P}\{\bar{\mathcal{E}}_{1,i}\} + \cdots + \mathbb{P}\{\bar{\mathcal{E}}_{p,i}\},$$

and by leveraging (17), one can deduce that

$$\begin{aligned}\mathbb{P}\{\mathcal{E}_{1,i} \cap \ldots \cap \mathcal{E}_{p,i}\} &\geq 1 - (\mathbb{P}\{\bar{\mathcal{E}}_{1,i}\} + \cdots + \mathbb{P}\{\bar{\mathcal{E}}_{p,i}\}) \\ &\geq 1 - \sum_{m=1}^{p} (\frac{\varepsilon_r^{m,i}}{\lambda_r^{m,i}} + \frac{\varepsilon_s^{m,i}}{\lambda_s^{m,i}}),\end{aligned}$$

which concludes the proof. ■

Remark 5.2: It is worth noting that if the system is not in the obstacle O_m when mode m is activated, then $\mathbb{P}\{\mathcal{R}_{m,i} \cap \mathcal{S}_{m,i}\} \geq 1 - \frac{\varepsilon_s^{m,i}}{\lambda_s^{m,i}}$ for that specific mode. In an ideal scenario, the lower bound guarantee proposed in Theorem 5.1 can be improved to $1 - \sum_{m=1}^{p} \frac{\varepsilon_s^{m,i}}{\lambda_s^{m,i}}$.

Remark 5.3: By employing the designed controllers and running Monte Carlo simulations, the empirical probabilities

slightly exceed our formal computations. This is expected due to the conservative nature of using fixed-degree polynomial barrier functions, which ensures formal guarantees. To mitigate this conservatism, higher-degree polynomials could be used for the barrier certificates or controllers, though this would increase computational complexity.

VI. Experimental Evaluation

In order to showcase the applicability of S-CBC and R-CBC to *reactive* logical control of stochastic dynamical systems, we consider a robotic case study with the following dynamics:

$$\mathsf{d}x(t) = (x(t) + \nu(t))\,\mathsf{d}t + 0.5x(t)\mathsf{d}\mathbb{W}_t,$$

with $x = [x_1; x_2]$ being the coordinate of the location of the robot, and $\nu = [\nu_1; \nu_2]$ being the control input. The regions of interest encompass the state space $X = [0,4]\times[0,4]$, consisting of two unsafe regions $\mathsf{O}_1 = [0,1.2]\times[2.8,4], \mathsf{O}_2 = [2.8,4]\times[0,1.2]$, along with the surrounding walls. We assume the robot is initially located in $\mathsf{I} = [3.6,3.8]\times[1.8,2]$ with mode $m = 1$ being activated. The robot then traverses the state space X, while avoiding O_1. When the robot is in the region $[2.8,4]\times[0,1.2]$, we assume that mode $m = 2$ is intended to be activated. Given that the robot is inside O_2, we first need to solve a reachability problem aiming to guide the robot away from O_2. We use the constructed partition cell with the highest safety probability as the target set for this purpose. Now since mode $m = 2$ is activated, the robot can traverse the state space while ensuring it avoids O_2.

We leverage SOSTOOLS [31] and the SDP solver SeDuMi [32] to construct all required S-CBC, R-CBC, and their associated safety and reachability controllers, as outlined in Definitions 4.1, 4.4. Alternatively, Julia-based SOS optimization tools can be used as well [33]. Firstly, upon activation of mode $m = 1$, to ensure avoidance of O_1 and surrounding walls, as per Lemma 4.6, we compute an S-CBC of an order 2 along with its associated safety controller as

$$\begin{aligned}\mathcal{B}_{s_1} =& 0.000015x_1^2 - 0.00015x_1 + 0.000019x_1x_2 - 0.00028x_2\\ &+ 0.000057x_2^2 + 0.00055,\\ \nu_1 =& -1.6x + 0.7\,\mathbb{I}_2.\end{aligned}$$

Furthermore, the corresponding constants in Definition 4.1 fulfilling conditions (3)-(5) are quantified as $\varepsilon_{s_1} = 9.099\times 10^{-7}$ and $\lambda_{s_1} = 3.807\times 10^{-5}$. Now by employing Theorem 4.2, we guarantee that the robot avoids O_1 and the surrounding walls under the designed policy with a probability of at least $1 - \frac{\varepsilon_{s_1}}{\lambda_{s_1}} = 98\%$. It is worth noting that we reported the barrier function values up to 6 decimal places, although they were computed with a precision of up to 16 decimal digits.

Given that the robot is in the region $[2.8,4]\times[0,1.2]$ (*i.e.* location of O_2) and mode $m = 2$ aims to be activated, one should first solve a *reachability* problem aiming to guide the robot reaching the partition cell with the highest safety probability from the heatmap in Fig. 1 (yellow part). Note that this heatmap is formally constructed by designing S-CBC for each partition cell and computing safety probabilities according to Theorem 4.2. We construct an R-CBC of an order 2 and its corresponding reachability controller for

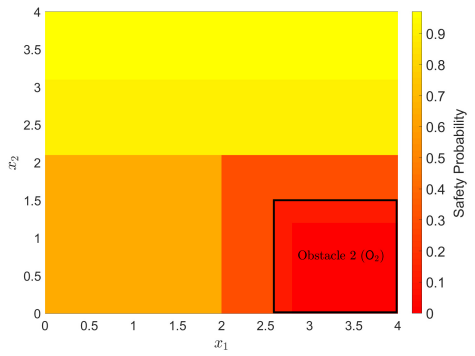

Fig. 1. Probability heatmap: Illustrating the safety probability of the robot starting from various initial sets within the state space $X = [0,4]\times[0,4]$. The red box marks obstacle $\mathsf{O}_2 = [2.8,4]\times[0,1.2]$ with a clear zero-safe probability. This heatmap is constructed via Definition 4.1 and Theorem 4.2.

steering the robot from O_2 to the dedicated partition cell while avoiding the surrounding walls:

$$\begin{aligned}\mathcal{B}_{r_2} =& 0.00065x_1^2 - 0.00667x_1 + 0.00142x_1x_2 - 0.00773x_2\\ &+ 0.00098x_2^2 + 0.01728,\\ \tilde{\nu}_2 =& -1.7x + 2.2\,\mathbb{I}_2.\end{aligned}$$

The corresponding constants in Definition 4.4 are quantified as $\varepsilon_{r_2} = 3.151\times10^{-5}, \lambda_{r_2} = 9.957\times10^{-4}$, and $\psi_{r_2} = 10^{-6}$. Now under Theorem 4.5, we guarantee that the robot goes from O_1 to partition cell $[0,4]\times[3,4]$ under the designed policy with a probability of at least $1 - \frac{\varepsilon_{r_2}}{\lambda_{r_2}} = 97\%$.

Finally, given that the mode $m = 2$ is activated, we leverage the S-CBC and its associated controller, constructed for the partition cell with the highest probability, to ensure that the robot avoids O_2 and the surrounding walls. In particular, the S-CBC and its corresponding safety controller are constructed as

$$\begin{aligned}\mathcal{B}_{s_2} =& 0.0000307x_1^2 - 0.000051x_1 + 0.000029x_1x_2\\ &- 0.000049x_2 + 0.000073x_2^2 + 0.00083,\\ \nu_2 =& -1.5x + 0.6\,\mathbb{I}_2,\end{aligned}$$

with $\varepsilon_{s_2} = 6.938\times10^{-8}, \lambda_{s_2} = 3.089\times10^{-6}$. Under Theorem 4.2, we guarantee that the robot avoids O_2 and the surrounding walls under the synthesized policy with a probability of at least $1 - \frac{\varepsilon_{s_2}}{\lambda_{s_2}} = 98\%$.

Now under the main result of Theorem 5.1, one can guarantee that the *overall mode-triggered safety-reachability specification* is fulfilled with an overall probability of at least $1 - \sum_{m=1}^{2}(\frac{\varepsilon_r^m}{\lambda_r^m} + \frac{\varepsilon_s^m}{\lambda_s^m}) = 93\%$. Note that there is no reachability probability for mode $m = 1$ (cf. Remark 5.2). Closed-loop state trajectories of the robot, obtained using the synthesized controllers across 10 noise realizations, are depicted in Fig. 2, thus affirming the theoretical guarantee provided.

References

[1] P. Tabuada, *Verification and control of hybrid systems: a symbolic approach*. Springer Science & Business Media, 2009.

[2] C. Belta, B. Yordanov, and E. A. Gol, *Formal methods for discrete-time dynamical systems*. Springer, 2017, vol. 89.

[3] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, "Probabilistic reachability and safety for controlled discrete-time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, 2008.

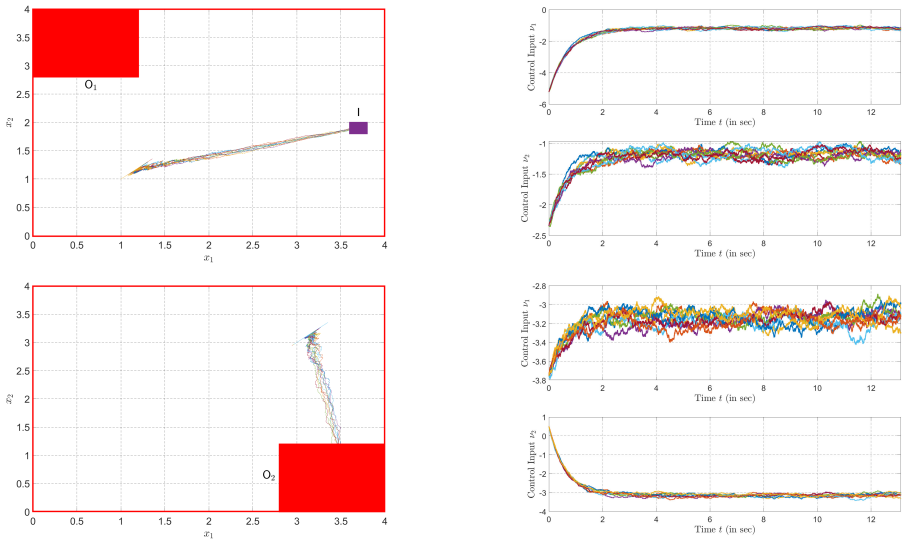

Fig. 2. Closed-loop state trajectories of the robot (left figures) under the synthesized *safety and reachability* controllers (right figures) with 10 noise realizations. Red boxes represent unsafe regions, while the purple box denotes the initial region. When mode $m = 2$ is activated (bottom figure), the robot located at O_2 is enforced by a *reachability* controller to reach the partition cell with the highest safety probability from the heatmap in Fig. 1 (yellow area).

[4] A. A. Julius and G. J. Pappas, "Approximations of stochastic hybrid systems," *IEEE Transactions on Automatic Control*, vol. 54, no. 6, pp. 1193–1203, 2009.

[5] M. Zamani, P. Mohajerin Esfahani, R. Majumdar, A. Abate, and J. Lygeros, "Symbolic control of stochastic systems via approximately bisimilar finite abstractions," *IEEE Transactions on Automatic Control*, vol. 59, no. 12, pp. 3135–3150, 2014.

[6] I. Tkachev, A. Mereacre, J.-P. Katoen, and A. Abate, "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*, 2013, pp. 293–302.

[7] A. Nejati, S. Soudjani, and M. Zamani, "Compositional abstraction-based synthesis for continuous-time stochastic hybrid systems," *European Journal of Control*, vol. 57, pp. 82–94, 2021.

[8] A. Lavaei, "Abstraction-based synthesis of stochastic hybrid systems," in *Proceedings of the 27th ACM International Conference on Hybrid Systems: Computation and Control*, 2024, pp. 1–11.

[9] A. Lavaei and M. Zamani, "From dissipativity theory to compositional synthesis of large-scale stochastic switched systems," *IEEE Transactions on Automatic Control*, vol. 67, no. 9, pp. 4422–4437, 2022.

[10] A. Nejati and M. Zamani, "Compositional construction of finite MDPs for continuous-time stochastic systems: A dissipativity approach," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 1962–1967, 2020.

[11] A. Lavaei and E. Frazzoli, "Scalable synthesis of finite MDPs for large-scale stochastic switching systems," in *61st Conference on Decision and Control (CDC)*, 2022, pp. 7510–7515.

[12] S. Prajna, A. Jadbabaie, and G. J. Pappas, "A framework for worst-case and stochastic safety verification using barrier certificates," *IEEE Transactions on Automatic Control*, vol. 52, no. 8, pp. 1415–1428, 2007.

[13] R. Wisniewski and M. L. Bujorianu, "Stochastic safety analysis of stochastic hybrid systems," in *Proceedings of the 56th IEEE Conference on Decision and Control*, 2017, pp. 2390–2395.

[14] C. Huang, X. Chen, W. Lin, Z. Yang, and X. Li, "Probabilistic safety verification of stochastic hybrid systems using barrier certificates," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 16, no. 5s, p. 186, 2017.

[15] A. Nejati, S. Soudjani, and M. Zamani, "Compositional construction of control barrier functions for continuous-time stochastic hybrid systems," *Automatica*, vol. 145, 2022.

[16] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proceedings of the 18th European Control Conference (ECC)*, 2019, pp. 3420–3431.

[17] A. Clark, "Control barrier functions for stochastic systems," *Automatica*, vol. 130, 2021.

[18] A. Lavaei and E. Frazzoli, "Scalable synthesis of safety barrier certificates for networks of stochastic switched systems," *IEEE Transactions on Automatic Control*, 2024.

[19] ——, "Compositional controller synthesis for interconnected stochastic systems with markovian switching," in *American Control Conference (ACC)*, 2022, pp. 4838–4843.

[20] S. Prajna and A. Rantzer, "Primal–dual tests for safety and reachability," in *International Workshop on Hybrid Systems: Computation and Control*, 2005, pp. 542–556.

[21] A. Lavaei, S. Soudjani, A. Abate, and M. Zamani, "Automated verification and synthesis of stochastic hybrid systems: A survey," *Automatica*, vol. 146, 2022.

[22] D. Gundana and H. Kress-Gazit, "Event-based signal temporal logic tasks: Execution and feedback in complex environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 001–10 008, 2022.

[23] L. Lindemann, G. J. Pappas, and D. V. Dimarogonas, "Reactive and risk-aware control for signal temporal logic," *IEEE Transactions on Automatic Control*, vol. 67, no. 10, pp. 5262–5277, 2022.

[24] S. P. Nayak, L. N. Egidio, M. Della Rossa, A.-K. Schmuck, and R. Jungers, "Context-triggered abstraction-based control design," *IEEE Open Journal of Control Systems*, 2023.

[25] A. Nejati, S. P. Nayak, and A.-K. Schmuck, "Context-triggered games for reactive synthesis over stochastic systems via control barrier certificates," in *27th International Conference on Hybrid Systems: Computation and Control*, 2024.

[26] I. Karatzas and S. Shreve, *Brownian motion and stochastic calculus*. springer, 2014, vol. 113.

[27] B. Oksendal, *Stochastic differential equations: An introduction with applications*. Springer Science & Business Media, 2013.

[28] H. J. Kushner, *Stochastic Stability and Control*, ser. Mathematics in Science and Engineering. Elsevier Science, 1967.

[29] M. Anand, V. Murali, A. Trivedi, and M. Zamani, "k-inductive barrier certificates for stochastic dynamical systems," in *25th International Conference on Hybrid Systems: Computation and Control*, 2022.

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

[31] A. Papachristodoulou, J. Anderson, G. Valmorbida, S. Prajna, P. Seiler, and P. Parrilo, "SOSTOOLS version 3.00 sum of squares optimization toolbox for MATLAB," *arXiv:1310.4716*, 2013.

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

[33] B. Legat, C. Coey, R. Deits, J. Huchette, and A. Perry, "Sum-of-squares optimization in julia," in *JuMP Developers Meetup/Workshop*, 2017.