

Formal and Efficient Synthesis for Continuous-Time Linear Stochastic Hybrid Processes

Luca Laurenti, Morteza Lahijanian, Alessandro Abate, Luca Cardelli, and Marta Kwiatkowska

Abstract—Stochastic processes are expressive mathematical tools for modeling real-world systems that are subject to uncertainty. It is hence crucial to be able to formally analyze the behavior of these processes, especially in safety-critical applications. Most of the existing formal methods are not designed for continuous-time processes, and those that are typically suffer from state explosion in practice. This work introduces a theoretical framework and a scalable computational method for formal analysis and control synthesis for switched diffusions, a class of stochastic models with linear dynamics that are continuous in both time and space domains; the focus is on safety with possible extensions to other properties. The proposed framework first constructs a finite abstraction in the form of an uncertain Markov process through discretization of both time and space domains. The errors caused by the discretization in each domain are formally characterized and cast into the abstraction model. Then, a strategy that maximizes the probability of the safety property and is robust against the errors is synthesized over the abstraction model. Finally, this robust strategy is mapped to a switching strategy for the stochastic processes that guarantees the safety property. The framework is demonstrated in three case studies, including one that illustrates the trade-off of the error contribution by the time and space discretization parameters.

Index Terms—Formal verification, formal synthesis, probabilistic model checking, stochastic hybrid systems, switched diffusions, formal abstractions, temporal logics.

I. INTRODUCTION

SWITCHED stochastic processes are powerful models for mathematical representation of real-world systems. These models, through switching between several *Stochastic Differential Equations* (SDEs), enable the inclusion of uncertainty as well as control, both of which are intrinsic in the physical laws of systems and in their interactions with the world [2]. Therefore, they are frequently used for modeling and analysis of systems in a wide range of domains such as cyber-physical systems [3]–[5], biological systems [6], and chemical reaction networks [7]. Despite their wide applicability, however, a theoretical framework that provides formal guarantees as well as scalable computational algorithms for analysis and control of switched stochastic systems remains a major challenge. In particular, the need for such a framework is of immediate

importance for *safety-critical* applications such as autonomous cars and air traffic control [8]. In this work, we target this challenge for switched diffusions, a class of stochastic models where the dynamics are linear and continuous in both time and space domains.

Formal approaches to verification and synthesis for stochastic processes have been the focus of many studies in recent years [9]–[18]. The popularity of these methods is rooted in the formal guarantees that they provide over expressive and succinct specification languages, such as *linear temporal logic* (LTL), *probabilistic computation tree logic* (PCTL), and *continuous stochastic logic* (CSL) [19]. The classical accompanying challenge, though, is the *state explosion* problem, which is particularly exacerbated for systems with continuous domains. To overcome this problem, approaches based on *finite abstractions*, which are essentially coarse representations of the process, are proposed [9]–[13]. Existing works, by and large, focus on discrete-time, continuous-space stochastic processes and construct an abstraction in the form of a discrete-time, discrete-space Markov process. This usually results in discrepancies between the abstraction and original system, which are then captured through error bounds. In practice, there are two major limitations to these approaches. One is that most real-world systems evolve in continuous time, and the other is the (lack of) scalability of their computational frameworks. That is, on the one hand, the error bounds can be conservative if the abstraction is not fine enough, and on the other hand fine abstractions can result in state-space explosion.

In this manuscript, we introduce a theoretical and computational synthesis framework for safety properties of (continuous time and space) switched diffusions that is both formal and scalable. The framework consists of two stages: abstraction and synthesis. In the first stage, an appropriate discrete abstraction in the form of an uncertain Markov model that captures all possible behaviors of the system is constructed. This is achieved through a discretization of time and space domains, each introducing an error. These errors are formally characterized and represented as uncertain transition probabilities in the abstraction model. For the space domain, the framework in particular uses a suitable discretization based on the dynamics of the system that results in closed-form analytical solutions for the error term, leading to fast computations. In the second stage, a robust strategy that optimizes a safety property is synthesized 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. Finally, this robust strategy is correctly mapped to a switching strategy for

L. Laurenti, A. Abate, L. Cardelli, and M. Kwiatkowska are with the Dept. of Computer Science at University of Oxford, U.K. (email: {*firstname.lastname*}@cs.ox.ac.uk). M. Lahijanian is with the Dept. of Aerospace Engineering Sciences at University of Colorado Boulder, CO, USA (email: morteza.lahijanian@colorado.edu).

This work was supported in part by EPSRC Mobile Autonomy Program Grant EP/M019918/1, Royal Society grant RP120138, and the Turing Institute, London, UK.

L. Laurenti and M. Lahijanian have contributed equally to this work and both are the corresponding authors.

Limited portions of this work have appeared in [1].

the system of switched diffusions with the guarantee that the computed safety probability bounds also hold for this system.

The contributions of this work are fourfold. Firstly, we introduce a theoretical framework for formal analysis of continuous-time, continuous-space diffusion processes. This includes the characterization and derivation of the error bounds that result from time discretization. Secondly, we propose a novel space discretization technique that is dynamics-dependent and derive the closed-form analytical solution for the computation of the exact (achievable) error bounds. Thirdly, we introduce a strategy synthesis algorithm for uncertain Markov processes with safety (invariance) properties, which is robust against the embedded uncertainty. Fourthly, we perform an empirical analysis on the trade-off between time and space discretization parameters. One of the main outcomes of this work is a computational framework that is both formal and scalable for switched stochastic processes. The choice of the abstraction model and the derivation of discretization methods with tight error bounds (exact for space) leads to fine and compact abstractions, whose computation is fast without the need to rely on sampling-based approaches (as is currently practiced in the literature, e.g., [12]). Furthermore, due to duality between probabilistic safety and reachability problems [20], the proposed framework can be easily extended to verification and synthesis for more complex properties [14], [21], [22] expressed in, e.g., PCTL and CSL.

II. PROBLEM FORMULATION AND APPROACH

We consider a switched stochastic process that is continuous both in time and space and evolves according to:

$$d\mathbf{x}(t) = F(a)\mathbf{x}(t)dt + G(a)d\mathbf{w}(t), \quad (1)$$

$$\mathbf{x} \in \mathbb{R}^m, \quad a \in A, \quad t \in \mathbb{R}_{\geq 0},$$

where

- $A = \{a_1, \dots, a_{|A|}\}$ is a finite set of *actions*,
- $F : A \rightarrow \mathbb{R}^{m \times m}$ is the *drift* term,
- $G : A \rightarrow \mathbb{R}^{m \times r}$ is the *diffusion* term, and
- \mathbf{w} is an r -dimensional Wiener process.

Under each action $a \in A$, process \mathbf{x} evolves according to the SDE in (1). An action change at time t causes a switch in the drift and diffusion values, resulting in an update in the dynamics of the process only, i.e., no change in the value of $\mathbf{x}(t)$ (unlike general stochastic hybrid systems [23]). Then, the process continues evolving according to the updated SDE from $\mathbf{x}(t)$ until the next switch.

Assumption 1 (Weakly controllable). *Under each action $a \in A$, process (1) is weakly controllable, i.e., the controllability subspace associated to $(F(a), G(a))$ is full rank.*

This ensures that, under each action a , $\mathbf{x}(t)$ is a non-degenerate Gaussian random variable with values in \mathbb{R}^m [24].

A *sample path* or *trajectory* ω_X of \mathbf{x} is a time-unbounded execution of the process (1), i.e., a function $\omega_X : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^m$. We denote the set of all sample paths by Ω , the segment of a path ω_X that is limited to the time interval $[0, t)$ by ω_X^t , and the value of path ω_X at time t by $\omega_X(t)$. Let $\mathcal{F}_t^{\mathbf{w}}$ be the sigma-algebra generated by the Wiener process $\mathbf{w}(t')$ for

$t' \leq t$ [25]. A *switching strategy* σ_X is a $\mathcal{F}_t^{\mathbf{w}}$ -measurable function that selects actions for the process in (1). Under σ_X , the stochastic process \mathbf{x} is defined on the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, Prob)$, where $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ is a filtration such that $\mathcal{F}_t \subseteq \mathcal{F}_{t'}$, and $Prob$ is a probability measure [26], [27].

In this work, we assume that σ_X is also \mathcal{F}_t -adapted, i.e., its choice of action at time t depends on ω_X^t (the trajectory of \mathbf{x} up to time t) [26]. In order to avoid Zeno behavior (an infinite number of action switches over a finite time horizon), we also assume that σ_X assigns a finite number of switches over any finite time interval almost surely. Under such a strategy, both the existence and the uniqueness of a solution \mathbf{x} are guaranteed since (1) is a linear SDE. Moreover, \mathbf{x} remains bounded almost surely over a finite time interval [25], [27].

A. Problem Formulation

We are interested in the safety analysis of process \mathbf{x} , and our aim is to compute a switching strategy that maximizes the probability of remaining in a safe set.

Problem 1. *Given the switched stochastic process in (1), a compact and measurable safe set $X_{\text{safe}} \subset \mathbb{R}^m$, and a time duration $\tau \in \mathbb{R}_{\geq 0}$, find a switching strategy σ_X that maximizes the safety probability of process \mathbf{x} , namely the probability of remaining in X_{safe} given by*

$$P_{\text{safe}}(\mathbf{x}, X_{\text{safe}}, \tau \mid \sigma_X) = Prob(\omega_X(t) \in X_{\text{safe}} \forall t \in [0, \tau] \mid \sigma_X). \quad (2)$$

Safety analysis is one of the fundamental problems in quantitative verification, which can be reformulated as a non-trivial stochastic optimal control problem with a multiplicative cost comprising indicator functions [20], and in particular it has been studied for stochastic hybrid processes [28]. In this paper, the focus is on such analysis, as formulated in Problem 1. It is worth noting that solving this problem can also lead to the solution of the dual problems of reachability analysis [20], again a core problem in formal verification [19]. Therefore, the solution to the safety Problem 1 can be extended to the verification of stochastic hybrid systems over more general PCTL and CSL properties [14], [21], [22].

The measurability of the event in (2) is based on the fact that \mathbf{x} is almost surely continuous. This implies that \mathbf{x} is a *separable* stochastic process [29, Theorem 38.1], and for separable stochastic processes, (2) is well defined. We emphasize that the measurability of (2) has been established also for more general classes of stochastic processes [30] and that evidently (2) does not account for sets of paths of measure zero that might escape the safe set.

B. Proposed Approach

Analytical approaches to Problem 1 are generally infeasible. That is due to the fact that such approaches require solutions to partial differential equations (PDE) with absorbing boundaries, which usually cannot be obtained in closed form [31]. The switched process in (1) can be formulated as a bilinear stochastic control problem with control taking values over a finite set and acting on both drift and diffusion terms

[32]; however, literature mostly considers actions only in the drift term and studies exclusively classical stochastic control problems, which furthermore require, in general, numerical solutions with an error that is very difficult (if possible at all) to compute [26], [33].

In this work, we take a general approach that is both formal and computationally tractable through a discrete abstraction. We construct a finite model in the form of an uncertain Markov process that captures all possible behaviors of the process in (1). This construction involves a discretization of both time and space domains, which results in a model that approximates the behavior of the continuous process in (1). We quantify the error of this approximation and represent it as uncertainty in the Markov model. We then synthesize an optimal strategy on this model that is robust against the uncertainty and can be mapped to the process in (1). Our solution also provides accurate error bounds on the safety probability of the process in (1) under this strategy. We should remark that the resulting strategy, although robust against uncertainties, is in general only sub-optimal for process (1) due to the discretized nature of the abstraction model. In fact, as discussed in Sec. VI, optimality is obtained only in the limit of infinitely-fine time and space discretizations

We note that the proposed solution framework is not limited to finite safety time durations and is able to handle unbounded durations, i.e., $\tau \in \mathbb{R}_{\geq 0} \cup \{\infty\}$. However, due to the unbounded stochastic nature of the process in (1), the safety probability is zero for a compact X_{safe} and $\tau = \infty$; hence, only $\tau \in \mathbb{R}_{\geq 0}$ is considered in Problem 1 (see [12], [34]).

III. PRELIMINARIES

A. Stochastic Hybrid Systems

A switched stochastic process as presented in (1) can be represented as a continuous-time *stochastic hybrid system* (SHS), where each action in A is viewed as a discrete mode (or location), and the evolution of the system under a discrete action is determined by the linear dynamics in the corresponding discrete mode. Below, we provide a definition for a simple class of SHS that is adapted from [23] and models the process in (1).

Definition 1 (SHS). *A (continuous-time) linear stochastic hybrid system \mathcal{H} is a tuple $\mathcal{H} = (A, m, F, G)$, where*

- $A = \{a_1, \dots, a_{|A|}\}$ is a finite set of discrete modes,
- $m \in \mathbb{N}$ defines the dimension of the continuous state space \mathbb{R}^m in each mode. The hybrid state space is defined as $S = A \times \mathbb{R}^m$,
- $F = \{F(a) \in \mathbb{R}^{m \times m} \mid a \in A\}$ is a collection of drift terms,
- $G = \{G(a) \in \mathbb{R}^{m \times r} \mid a \in A\}$ is a collection of diffusion terms.

The evolution of stochastic hybrid system \mathcal{H} is a stochastic process $\mathbf{s}(t) = (\mathbf{a}(t), \mathbf{x}(t))$ with values in S . The term \mathbf{x} represents the evolution of the continuous component of \mathcal{H} , while \mathbf{a} describes the evolution of the discrete components over time. Let $Paths_{\mathcal{H}}^t$ be the sets of paths of \mathbf{s} over the time interval $[0, t)$. Then, a *switching strategy* at time t for \mathcal{H} is a

measurable function $\sigma_{\mathcal{H}} : Paths_{\mathcal{H}}^t \rightarrow A$ that assigns a discrete mode to each path up to time t .

Assumption 2 (Piecewise-constant $\sigma_{\mathcal{H}}$). *The switching strategy $\sigma_{\mathcal{H}}$ for \mathcal{H} is a piecewise-constant function that may change its value only at time instants $t = k \cdot \Delta t$, where $k \in \mathbb{N}$ and $\Delta t \in \mathbb{R}_{>0}$.*

This assumption considers $\sigma_{\mathcal{H}}$ to be a piecewise-constant function with constant sampling times. Hence, $\sigma_{\mathcal{H}}$ is a special case of the strategy σ_X defined in Sec. II.

Definition 2 (SHS execution). *An execution of a SHS $\mathcal{H} = (A, m, F, G)$ under a switching strategy $\sigma_{\mathcal{H}}$ and an initial state $\mathbf{s}(0) = (\mathbf{a}(0), \mathbf{x}(0)) \in A \times \mathbb{R}^m$ is a stochastic process $\mathbf{s}(t)$, whose sample paths for the duration of $t \in [0, t_f)$ denoted by $\omega_{\mathcal{H}}^{t_f} \in Paths_{\mathcal{H}}^{t_f}$ are obtained according to the following algorithm:*

```

set  $t = 0$  and  $\omega_{\mathcal{H}}^0 = (\mathbf{a}(0), \mathbf{x}(0))$ ;
while  $t < t_f$ 
  set  $a = \sigma_{\mathcal{H}}(\omega_{\mathcal{H}}^t)$ ;
  compute  $\omega_X^{\Delta t}$  according to (1) with initial state  $\mathbf{x}(t)$ ;
  extend  $\omega_{\mathcal{H}}^t$  by  $(a, \omega_X^{\Delta t})$ ;
   $t = t + \Delta t$ ;
end

```

Given $\sigma_{\mathcal{H}}$, the probability space defined on process \mathbf{x} naturally extends to process \mathbf{s} . We can associate to $\mathbf{x}(t)$ a controlled transition probability measure T^d . For $a \in A$, let $\mathcal{B}(\mathbb{R}^m)$ be the Borel sigma-algebra on the state space \mathbb{R}^m . Then, starting from a continuous state $x \in \mathbb{R}^m$ in mode a at time $t \in \mathbb{R}_{\geq 0}$, the probability that $\mathbf{x}(t + \Delta t) \in B$, where $B \in \mathcal{B}(\mathbb{R}^m)$ is a region in the state space \mathbb{R}^m , with the assumption of no mode change in the duration Δt is given by

$$T^d(B \mid x, a, \Delta t) = \int_B \mathcal{N}(z \mid E_{\mathbf{x}}(\Delta t), Cov_{\mathbf{x}}(\Delta t)) dz, \quad (3)$$

where $\mathcal{N}(\cdot \mid E_{\mathbf{x}}(\Delta t), Cov_{\mathbf{x}}(\Delta t))$ is a multivariate normal density function with mean $E_{\mathbf{x}}(\Delta t)$ and covariance matrix $Cov_{\mathbf{x}}(\Delta t)$ given by

$$E_{\mathbf{x}}(\Delta t) = e^{F(a)\Delta t}x, \quad (4)$$

$$Cov_{\mathbf{x}}(\Delta t) = \int_0^{\Delta t} e^{F(a)(\Delta t-z)}G(a)G^T(a)(e^{F(a)(\Delta t-z)})^T dz. \quad (5)$$

Lemma 1. *The covariance matrix $Cov_{\mathbf{x}}(\Delta t)$ is full rank.*

Proof. The proof follows directly from Assumption 1. \square

B. Finite-state Markov Models

We utilize Markov models as abstraction structures.

Definition 3 (MDP). *A Markov Decision Process (MDP) is a tuple $\mathcal{M} = (Q, A, P)$, where:*

- Q is a finite set of states,
- A is a finite set of actions,
- $P : Q \times A \times Q \rightarrow [0, 1]$ is a transition probability function.

The set of actions available at $q \in Q$ is denoted by $A(q)$. The function P has the property $\sum_{q' \in Q} P(q, a, q') = 1$ for all $q \in Q$ paired with each $a \in A(q)$.

A path ω through an MDP is a sequence of states $\omega = q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \xrightarrow{a_2} \dots$ such that $a_i \in A(q_i)$ and $P(q_i, a_i, q_{i+1}) > 0$ for all $i \in \mathbb{N}$. We denote the last state of a finite path ω^{fin} by $\text{last}(\omega^{\text{fin}})$ and the set of all finite and infinite paths by $\text{Paths}^{\text{fin}}$ and Paths , respectively.

A *strategy* defines a choice of action at each state of the MDP. Its formal definition follows.

Definition 4 (Strategy). A strategy σ of an MDP model \mathcal{M} is a function mapping a finite path $\omega^{\text{fin}} = q_0 q_1 \dots q_n$ of \mathcal{M} onto an action in A . In other words, a strategy is a function $\sigma : \text{Paths}^{\text{fin}} \rightarrow A$ that specifies, for every finite path, the next action to be applied. If a strategy depends only on $\text{last}(\omega^{\text{fin}})$ and time step n , it is called a (time-dependent) Markov strategy. If a strategy depends only on $\text{last}(\omega^{\text{fin}})$, it is called a memoryless or stationary Markov strategy. The set of all strategies is denoted by Σ^1 .

Given strategy σ , a probability measure *Prob* over the set of all paths Paths is induced on the resulting Markov chain [19].

When modeling with MDPs, it might be difficult to determine exact values of transition probabilities between states. In such cases, an interval for each value may be considered. The model that allows the inclusion of these intervals is known as *bounded-parameter* [35] or *interval MDP* (IMDP) [22].

Definition 5 (IMDP). An interval Markov decision process (IMDP) is a tuple $\mathcal{I} = (Q, A, \check{P}, \hat{P})$, where Q and A are as in Def. 3, and

- $\check{P} : Q \times A \times Q \rightarrow [0, 1]$ is a function, where $\check{P}(q, a, q')$ is the infimum (lower bound) of the transition probabilities from state q to state q' under action $a \in A(q)$,
- $\hat{P} : Q \times A \times Q \rightarrow [0, 1]$ is a function, where $\hat{P}(q, a, q')$ is the supremum (upper bound) of the transition probabilities from state q to state q' under action $a \in A(q)$.

For all $q, q' \in Q$ and $a \in A(q)$, it holds that $\check{P}(q, a, q') \leq \hat{P}(q, a, q')$ and $\sum_{q' \in Q} \check{P}(q, a, q') \leq 1 \leq \sum_{q' \in Q} \hat{P}(q, a, q')$. Let $\mathcal{D}(Q)$ denote the set of discrete probability distributions over Q . Given $q \in Q$ and $a \in A(q)$, we call $\gamma_q^a \in \mathcal{D}(Q)$ a *feasible distribution* reachable from q by a if

$$\check{P}(q, a, q') \leq \gamma_q^a(q') \leq \hat{P}(q, a, q')$$

for each state $q' \in Q$. We denote the set of all feasible distributions for state q and action a by Γ_q^a . In IMDPs, the notions of paths and strategies are extended from MDPs in a straightforward manner. The additional concept is that of *adversary*, which makes the choice of a feasible distribution².

Definition 6 (Adversary). Given an IMDP \mathcal{I} , an adversary is a function $\pi : \text{Paths}^{\text{fin}} \times A \rightarrow \mathcal{D}(Q)$ that, for each finite path

$\omega^{\text{fin}} \in \text{Paths}^{\text{fin}}$ and action $a \in \text{last}(\omega^{\text{fin}})$, assigns a feasible distribution $\pi(\omega^{\text{fin}}, a) \in \Gamma_{\text{last}(\omega^{\text{fin}})}^a$. The set of all adversaries is denoted by Π .

Given a finite path ω^{fin} , a strategy σ , and an adversary π , the IMDP evolution proceeds as follows. At state $q = \text{last}(\omega^{\text{fin}})$, first an action $a \in A(q)$ is chosen by σ . Then, π resolves the uncertainties and chooses one feasible distribution $\gamma_q^a \in \Gamma_q^a$. Finally, the next state q' is chosen according to the distribution γ_q^a , and the path ω^{fin} is extended by q' .

Given a strategy σ and an adversary π , a probability measure *Prob* over the set of all paths Paths (under σ and π) is induced by the resulting Markov chain [12].

C. Polytopes and their Post Images

Let $m \in \mathbb{N}$ and consider the m -dimensional Euclidean space \mathbb{R}^m . A full dimensional *polytope* P is defined as the convex hull of at least $m + 1$ affinely independent points in \mathbb{R}^m [36]. The set of vertices of P is the set of points $v_1^P, \dots, v_{n_P}^P \in \mathbb{R}^m$, $n_P \geq m + 1$, whose convex hull gives P and with the property that, for any $i = 1, \dots, n_P$, point v_i^P is not in the convex hull of the remaining points $v_1^P, \dots, v_{i-1}^P, v_{i+1}^P, \dots, v_{n_P}^P$. A polytope is completely described by its set of vertices,

$$P = \text{conv}(v_1^P, \dots, v_{n_P}^P), \quad (6)$$

where *conv* denotes the convex hull. Alternatively, P can be described as the bounded intersection of at least $m + 1$ closed half spaces. In other words, there exists a $k \geq m + 1$, $h_i \in \mathbb{R}^m$, and $l_i \in \mathbb{R}$, $i = 1, \dots, k$ such that

$$P = \{x \in \mathbb{R}^m \mid h_i^T x \leq l_i, i = 1, \dots, k\}. \quad (7)$$

The above definition can be written as the matrix inequality $Hx \leq L$, where $H \in \mathbb{R}^{k \times m}$ and $L \in \mathbb{R}^k$. Forms (6) and (7) are referred to as V - and H -representations of the polytope.

Given a matrix $M \in \mathbb{R}^{m \times m}$, the post image of polytope P by M is defined as [12]:

$$\text{Post}(P|M) = \{Mx \mid x \in P\}.$$

This post image is a polytope itself under the linear transformation M and can be computed as $\text{Post}(P|M) = \text{conv}(\{Mv_i^P, 1 \leq i \leq n_P\})$.

IV. DISCRETE ABSTRACTIONS

In the first step of our approach to Problem 1, we model the process in (1) as a SHS and then construct a discrete abstraction of it in the form of an IMDP. This abstraction is based on discretization of both time and space domains. That is, we sample time at a fixed time interval $\Delta t > 0$ and partition the given safe set into polytopical regions. Such a discretization, however, introduces an error, which needs to be accounted for as detailed below.

A. Stochastic Hybrid System Modeling

We model the switched stochastic process in (1) as a SHS \mathcal{H} as defined in Section III-A. Each action of the process in (1) is associated with a mode of \mathcal{H} , and the drift and diffusion terms in mode a are $F(a)$ and $G(a)$, respectively. In this model, each

¹In this work, we focus on (time-dependent) Markov strategies as they are sufficient for optimality of safety (and reachability) properties for MDPs and IMDPs [12], [20].

²In the verification literature for MDPs, the notions of strategy, adversary, or policy are equivalent and used interchangeably. Their semantics are however distinguished over IMDPs.




Figure 1: Hybrid system model of a switched diffusion process with two actions.

mode of \mathcal{H} contains a copy of the safe set $X_{\text{safe}} \subset \mathbb{R}^m$, and hence, the hybrid safe set is

$$S_{\text{safe}} = \{ (a, x) \mid x \in X_{\text{safe}} \text{ and } a \in A \}. \quad (8)$$

Therefore, the safety probability of process $\mathbf{x}(t)$ in (2) becomes

$$P_{\text{safe}}(\mathbf{s}, S_{\text{safe}}, \tau \mid \sigma_{\mathcal{H}}) = \text{Prob}(\omega_{\mathcal{H}}(t) \in S_{\text{safe}} \forall t \in [0, \tau] \mid \sigma_{\mathcal{H}}). \quad (9)$$

for the process $\mathbf{s}(t)$ of the hybrid system \mathcal{H} , and Problem 1 becomes equivalent to finding $\sigma_{\mathcal{H}}$ that maximizes (9).

Figure 1 illustrates a hybrid system representation of a switched stochastic process with two actions. The hybrid system has two modes denoted by a_1 and a_2 . The blue surfaces are the copies of set X_{safe} in each mode, and the black continuous trajectory is a sample path of the hybrid system with two switches. Note that a segment of the path in mode a_1 exits and re-enters the safe set. Lastly, the values of $\mathbf{x}(t)$ on the trajectory do not change at the times of switching (illustrated as dashed vertical lines); only the modes $\mathbf{a}(t)$ change.

B. IMDP Abstraction

We abstract the hybrid system \mathcal{H} to an IMDP $\mathcal{I} = (Q, A, \hat{P}, \hat{P})$ with the purpose of safety analysis. To this end, we perform a discretization of the hybrid state space S by distinguishing between the safe and unsafe states and exclusively focusing on S_{safe} as described below.

For each mode $a \in A$, we partition the corresponding safe set X_{safe} into a set of cells (regions) that are non-overlapping except for at most trivial sets of measure zero (their boundaries). We assume that each region is a bounded polytope. We denote by $Q^a = \{q_1^a, \dots, q_{|Q^a|}^a\}$ the resulting set of regions in mode a . To each cell q_i^a , we associate a state of the IMDP \mathcal{I} . We overload the notation by using q_i^a for both a region in S , i.e., $q_i^a \subseteq \{a\} \times X_{\text{safe}} \subseteq S_{\text{safe}} \subseteq S$, and a state of \mathcal{I} , i.e., $q_i^a \in Q$; the exact meaning of it should be clear from the context. Therefore, S_{safe} can be represented by $Q_{\text{safe}} = \bigcup_{a \in A} Q^a$. We define the set of IMDP states to be

$$Q = Q_{\text{safe}} \cup \{q_{\text{unsafe}}\},$$

where q_{unsafe} is an IMDP state that represents the remaining hybrid states in $S \setminus Q_{\text{safe}}$.

We define the set of actions of \mathcal{I} to be the set of modes A of \mathcal{H} and allow all actions to be available at each state of \mathcal{I} , i.e., $A(q) = A$ for all $q \in Q$. To capture the safe behavior of \mathcal{H} by \mathcal{I} , we define the one-step safe transition probability from a continuous state $x \in X_{\text{safe}}$ to region $q \in Q_{\text{safe}}$ under action (mode) $a \in A$ to be

$$P_{\text{safe}}(q \mid x, a, \Delta t) = \text{Prob}(\mathbf{x}(\Delta t) \in q \wedge \forall t \in [0, \Delta t] (\mathbf{a}(t), \mathbf{x}(t)) \in S_{\text{safe}} \mid \mathbf{x}(0) = x, \mathbf{a}(t) = a \forall t \in [0, \Delta t]). \quad (10)$$

In other words, $P_{\text{safe}}(q \mid x, a, \Delta t)$ is the probability of reaching q from x in Δt without leaving the safe set. By the application of the law of conditional probabilities we get:

$$P_{\text{safe}}(q \mid x, a, \Delta t) = T^d(q \mid x, a, \Delta t) T^c(\Delta t \mid x, q, a), \quad (11)$$

where

$$T^d(q \mid x, a, \Delta t) = \text{Prob}(\mathbf{x}(\Delta t) \in q \mid \mathbf{x}(0) = x, \mathbf{a}(t) = a \forall t \in [0, \Delta t])$$

is the *discrete transition kernel* that determines the probability of ending up in q after Δt , as introduced in (3), and

$$T^c(\Delta t \mid x, q, a) = \text{Prob}((\mathbf{a}(t), \mathbf{x}(t)) \in S_{\text{safe}} \forall t \in [0, \Delta t] \mid \mathbf{x}(0) = x, \mathbf{x}(\Delta t) \in q, \mathbf{a}(t) = a \forall t \in [0, \Delta t])$$

is the *continuous transition kernel* that considers the safety of the continuous trajectories over the time interval Δt . Note that T^d considers all the paths that end up in q from x . This can be viewed as an error caused by sampling time at intervals of Δt , i.e., a time discretization error. The term T^c corrects this error by considering those paths that exit and then re-enter the safe set during the sampling time interval. As shown in Section IV-D, the computation of T^c reduces to quantify excursion probabilities for Gaussian processes, which, in general, cannot be obtained analytically [37]. Hence, in what follows we will derive sound upper and lower bounds for T^c .

We define the transitions in \mathcal{I} by using the safe transition probabilities in (11). The caveat is that the states of \mathcal{I} are regions in \mathcal{H} , and there are uncountably many possible (continuous) initial states in each region, giving rise to a range of safe feasible transition probabilities to the other regions. Therefore, the exact transition probability from one region to another cannot be known, but its range is given by min and max of (11) over all the possible points in the initial region. That is, for $q_i, q_j \in Q_{\text{safe}}$ and $a \in A$, the safe feasible transition probability is bounded from below by

$$\begin{aligned} \gamma_{q_i}^a(q_j) &\geq \min_{x \in q_i} P_{\text{safe}}(q_j \mid x, a, \Delta t) \\ &= \min_{x \in q_i} T^d(q_j \mid x, a, \Delta t) T^c(\Delta t \mid x, q_j, a) \\ &\geq \left(\min_{x \in q_i} T^d(q_j \mid x, a, \Delta t) \right) \left(\min_{x \in q_i} T^c(\Delta t \mid x, q_j, a) \right), \end{aligned} \quad (12)$$

$$(13)$$

and from above by

$$\begin{aligned} \gamma_{q_i}^a(q_j) &\leq \max_{x \in q_i} P_{\text{safe}}(q_j | x, a, \Delta t) \\ &= \max_{x \in q_i} T^d(q_j | x, a, \Delta t) T^c(\Delta t | x, q_j, a) \\ &\leq \left(\max_{x \in q_i} T^d(q_j | x, a, \Delta t) \right) \left(\max_{x \in q_i} T^c(\Delta t | x, q_j, a) \right). \end{aligned} \quad (14)$$

Therefore, we can define the transition probability bounds \check{P} and \hat{P} of \mathcal{I} according to these bounds. Ideally, they should be the exact bounds given in (12) and (14), but they are usually difficult to compute as discussed in Sec. IV-D. Instead, we set the values of \check{P} and \hat{P} according to the possibly looser bounds in (13) and (15), i.e.,

$$\check{P}(q_i, a, q_j) = \left(\min_{x \in q_i} T^d(q_j | x, a, \Delta t) \right) \left(\min_{x \in q_i} T^c(\Delta t | x, q_j, a) \right), \quad (16)$$

$$\hat{P}(q_i, a, q_j) = \left(\max_{x \in q_i} T^d(q_j | x, a, \Delta t) \right) \left(\max_{x \in q_i} T^c(\Delta t | x, q_j, a) \right), \quad (17)$$

for all $a \in A$ and $q_i, q_j \in Q_{\text{safe}}$.

Similarly, we define the bounds of the feasible transition probabilities to the unsafe state as

$$\gamma_{q_i}^a(q_{\text{unsafe}}) \geq 1 - \max_{x \in q_i} P_{\text{safe}}(X_{\text{safe}} | x, a, \Delta t), \quad (18)$$

$$\gamma_{q_i}^a(q_{\text{unsafe}}) \leq 1 - \min_{x \in q_i} P_{\text{safe}}(X_{\text{safe}} | x, a, \Delta t), \quad (19)$$

and set the bounds in \mathcal{I} to be

$$\check{P}(q_i, a, q_{\text{unsafe}}) = 1 - \max_{x \in q_i} P_{\text{safe}}(X_{\text{safe}} | x, a, \Delta t), \quad (20)$$

$$\hat{P}(q_i, a, q_{\text{unsafe}}) = 1 - \min_{x \in q_i} P_{\text{safe}}(X_{\text{safe}} | x, a, \Delta t), \quad (21)$$

for all $a \in A$ and $q_i \in Q_{\text{safe}}$. Finally, we make the unsafe state q_{unsafe} absorbing, i.e.,

$$\check{P}(q_{\text{unsafe}}, a, q_{\text{unsafe}}) = \hat{P}(q_{\text{unsafe}}, a, q_{\text{unsafe}}) = 1,$$

for all $a \in A$.

Note that the ranges of transition probabilities can be viewed as errors caused by the space discretization. The benefit of using IMDP as the abstraction model is that its structure allows one to encode both time and space discretization errors as uncertainty into the abstraction. In the next three subsections, we show an efficient method of computation for \check{P} and \hat{P} , and in Section V, we show how a safe strategy that is robust against the uncertainty can be synthesized over \mathcal{I} . Finally, we prove that the synthesized strategy on \mathcal{I} can be soundly mapped on to the process in (1) in Section VI.

C. Bounds on the Discrete Transition Kernel

Here, we focus on the values of

$$\min_{x \in q_i} T^d(q_j | x, a, \Delta t), \quad \max_{x \in q_i} T^d(q_j | x, a, \Delta t), \quad (22)$$

and introduce an efficient method for their exact computations. To this end, we first define a *hyper-rectangle* in \mathbb{R}^m to be an m -dimensional rectangle defined by the intervals

$$[v_l^{(1)}, v_u^{(1)}] \times [v_l^{(2)}, v_u^{(2)}] \times \cdots \times [v_l^{(m)}, v_u^{(m)}],$$

where vectors $v_l, v_u \in \mathbb{R}^m$ capture the lower and upper values of the vertices of the rectangle in each dimension, and $v^{(i)}$ denotes the i -th component of vector v . Then, we characterize T^d analytically as follows.

Proposition 1. *For process \mathbf{x} under action $a \in A$ for a duration Δt , let $\mathcal{T}_a = \Lambda_a^{-\frac{1}{2}} V_a^T$ be a transformation function (matrix), where $\Lambda_a = V_a^T \text{Cov}_{\mathbf{x}}(\Delta t) V_a$ is a diagonal matrix whose entries are eigenvalues of $\text{Cov}_{\mathbf{x}}(\Delta t)$ and V_a is the corresponding orthonormal (eigenvector) matrix. For a polytopic region $q \subset \mathbb{R}^m$, if $\text{Post}(q | \mathcal{T}_a)$ is a hyper-rectangle given by $[v_l^{(1)}, v_u^{(1)}] \times \cdots \times [v_l^{(m)}, v_u^{(m)}]$, then it holds that*

$$T^d(q | x, a, \Delta t) = \frac{1}{2^m} \prod_{i=1}^m \left(\text{erf}\left(\frac{y^{(i)} - v_l^{(i)}}{\sqrt{2}}\right) - \text{erf}\left(\frac{y^{(i)} - v_u^{(i)}}{\sqrt{2}}\right) \right), \quad (23)$$

where $\text{erf}(\cdot)$ is the error function, and $y^{(i)}$ is the i -th component of vector $y = \mathcal{T}_a E_{\mathbf{x}}(\Delta t)$.

Proof. For a fixed $a \in A$ and duration Δt , recall that $T^d(q | x, a, \Delta t)$ is given by (3). By applying a whitening transformation through the transformation matrix $\mathcal{T}_a = \Lambda_a^{-\frac{1}{2}} V_a^T$, we obtain that $\mathcal{T}_a \text{Cov}_{\mathbf{x}}(\Delta t) \mathcal{T}_a^T = \mathbf{I}$, where \mathbf{I} is the identity matrix. Thus, by working in the transformed space induced by \mathcal{T}_a , we obtain

$$T^d(q | x, a, \Delta t) = \int_{\text{Post}(q | \mathcal{T}_a)} \mathcal{N}(z | \mathcal{T}_a E_{\mathbf{x}}(\Delta t), \mathbf{I}) dz.$$

Under the assumption that $\text{Post}(q | \mathcal{T}_a)$ is a hyper-rectangle, the above multidimensional integral can be separated and expressed as a product of m integrals of uni-dimensional normal distributions:

$$\begin{aligned} T^d(q | x, a, \Delta t) &= \int_{\text{Post}(q | \mathcal{T}_a)} \mathcal{N}(z | \mathcal{T}_a E_{\mathbf{x}}(\Delta t), \mathbf{I}) dz \\ &= \int_{v_l^{(1)}}^{v_u^{(1)}} \cdots \int_{v_l^{(m)}}^{v_u^{(m)}} \mathcal{N}(z_1 | y^{(1)}, 1) \cdots \mathcal{N}(z_m | y^{(m)}, 1) dz_1 \cdots dz_m \\ &= \prod_{i=1}^m \frac{1}{2} \left(\text{erf}\left(\frac{y^{(i)} - v_l^{(i)}}{\sqrt{2}}\right) - \text{erf}\left(\frac{y^{(i)} - v_u^{(i)}}{\sqrt{2}}\right) \right), \end{aligned}$$

where $y = \mathcal{T}_a E_{\mathbf{x}}(\Delta t)$. \square

A direct consequence of Proposition 1 is that the optimization problems in (22) can be performed on (23) through a linear transformation, as stated by the following corollary.

Corollary 1. *For polytopic regions $q_i, q_j \subset \mathbb{R}^m$ and process \mathbf{x} under action a for the duration of Δt , assume $\text{Post}(q_j | \mathcal{T}_a)$ is a hyper-rectangle given by*

$$[v_l^{(1)}, v_u^{(1)}] \times \cdots \times [v_l^{(m)}, v_u^{(m)}],$$

and define $q'_i = \text{Post}(q_i | e^{F(a)\Delta t})$ and

$$f(y) = \frac{1}{2^m} \prod_{i=1}^m \left(\text{erf}\left(\frac{y^{(i)} - v_l^{(i)}}{\sqrt{2}}\right) - \text{erf}\left(\frac{y^{(i)} - v_u^{(i)}}{\sqrt{2}}\right) \right).$$

Then, it holds that

$$\begin{aligned} \min_{x \in q_i} T^d(q_j | x, a, \Delta t) &= \min_{y \in \text{Post}(q'_i | \mathcal{T}_a)} f(y), \\ \max_{x \in q_i} T^d(q_j | x, a, \Delta t) &= \max_{y \in \text{Post}(q'_i | \mathcal{T}_a)} f(y). \end{aligned}$$

The above proposition and corollary show that, with a particular geometry for q_j , an analytical form can be obtained for the discrete kernel. This is an important observation because it enables efficient computation for the min and max values of the kernel. Therefore, we use a space discretization to satisfy the condition in Proposition 1, as described below.

For each mode $a \in A$, we define the linear transformation function (matrix) of

$$\mathcal{T}_a = \Lambda_a^{-\frac{1}{2}} V_a^T,$$

where $\Lambda_a = V_a^T \text{Cov}_x(\Delta t) V_a$ is a diagonal matrix whose entries are the eigenvalues of $\text{Cov}_x(\Delta t)$, and V_a is the corresponding orthonormal (eigenvector) matrix. The discretization of X_{safe} in mode a is achieved by using a grid in the transformed space by \mathcal{T}_a . That is, we first transform X_{safe} by \mathcal{T}_a , and then discretize it using a grid, and then transform it back to the original space using \mathcal{T}_a^{-1} . This method of discretization guarantees that, for each $q^a \in Q^a$, $\text{Post}(q^a | \mathcal{T}_a)$ is a hyper-rectangle. Hence, we can use the result of Proposition 1 and Corollary 1 for the computation of the values in (22).

It is worth noting that, for an arbitrary geometry of X_{safe} , it may not be possible to obtain a discretization such that $\bigcup_{q^a \in Q^a} q^a = X_{\text{safe}}$. Nevertheless, for the purpose of safety analysis, a discretization that under-approximates X_{safe} , i.e., $\bigcup_{q^a \in Q^a} q^a \subseteq X_{\text{safe}}$, is sufficient in each mode a . For a better approximation, the grid can be non-uniform, allowing in particular for smaller cells near the boundary of X_{safe} , as in [11].

In the next theorem, we use the result of Corollary 1 and the Karush-Kuhn-Tucker (KKT) conditions [38] to compute the exact values for (22).

Theorem 1. For polytopic regions $q_i, q_j \subset \mathbb{R}^m$ and transformation matrix \mathcal{T}_a , let $\text{Post}(q_j | \mathcal{T}_a)$ be a hyper-rectangle defined by intervals

$$[v_l^{(1)}, v_u^{(1)}] \times [v_l^{(2)}, v_u^{(2)}] \times \dots \times [v_l^{(m)}, v_u^{(m)}],$$

and

$$\text{Post}(q'_i | \mathcal{T}_a) = \{y \in \mathbb{R}^m \mid Hy \leq b\},$$

where $q'_i = \text{Post}(q_i | e^{F(a)\Delta t})$, $H \in \mathbb{R}^{k \times m}$, $b \in \mathbb{R}^m$, and $k \geq m + 1$, and call

$$f(y) = \frac{1}{2^m} \prod_{i=1}^m \left(\text{erf}\left(\frac{y^{(i)} - v_l^{(i)}}{\sqrt{2}}\right) - \text{erf}\left(\frac{y^{(i)} - v_u^{(i)}}{\sqrt{2}}\right) \right). \quad (24)$$

Introduce the following conditions:

- **Condition 1:** y is at the center of $\text{Post}(q_j | \mathcal{T}_a)$, i.e., $y = (\frac{v_u^{(1)} + v_l^{(1)}}{2}, \dots, \frac{v_u^{(m)} + v_l^{(m)}}{2})$.
- **Condition 2:** y is a vertex of $\text{Post}(q'_i | \mathcal{T}_a)$.
- **Condition 3:** y is on the boundary of $\text{Post}(q'_i | \mathcal{T}_a)$, where $r \geq 1$ of the k half-spaces that define $\text{Post}(q'_i | \mathcal{T}_a)$ intersect, and

$$\nabla f(y) = \bar{H}^T \mu,$$

for vector $\mu = (\mu_1, \dots, \mu_r)$ of non-negative constants, and submatrix $\bar{H} \in \mathbb{R}^{r \times m}$ that contains only the rows of H that correspond to the r -intersecting half-spaces at y .

- **Condition 4:** y is as in Condition 3, and

$$\nabla f(y) = -\bar{H}^T \mu,$$

for vector $\mu = (\mu_1, \dots, \mu_r)$ of non-negative constants, and \bar{H} is defined as in Condition 3.

Then, it follows that the point $y \in \text{Post}(q'_i | \mathcal{T}_a)$ that satisfies Condition 1 necessarily maximizes $f(y)$. If Condition 1 cannot be satisfied, then the maximum is necessarily given by one of the points that satisfy Condition 2 or 3. Furthermore, the point $y \in \text{Post}(q'_i | \mathcal{T}_a)$ that minimizes $f(y)$ necessarily satisfies Condition 2 or 4.

The proof of Theorem 1 can be found in Appendix A.

Theorem 1 identifies the arguments (points y in $\text{Post}(q'_i | \mathcal{T}_a)$) that give rise to the optimal values of the discrete transition kernel in (22). Then, the actual optimal values of T^d can be computed by (24) as guaranteed by Corollary 1. Therefore, from Theorem 1, an algorithm can be constructed to generate a set of finite candidate points based on Conditions 1-4 and to obtain the exact values of (22) by plugging those points into (24).

In short, Condition 1 maximizes the unconstrained problem and gives rise to the global maximum. Hence, if the center of q_j is contained in $\text{Post}(q'_i | \mathcal{T}_a)$, no further check is required for maximum. If not, the maximum is given by a point on the boundary of $\text{Post}(q'_i | \mathcal{T}_a)$. It is either a vertex (Condition 2) or a boundary point that satisfies Condition 3. The minimum is always given by a boundary point, which can be either a vertex or a boundary point that satisfies Condition 4. Note that Conditions 3 and 4 are similar and both state that the optimal value of T^d is given by a point where the gradient of T^d becomes linearly dependent on the vectors that are defined by the intersecting half-spaces of $\text{Post}(q'_i | \mathcal{T}_a)$ at that point. Each of these two conditions defines a system of m equations and $r < m$ variables, which may have a solution only if some of the equations are linear combinations of the others.

In order to give a better insight on the conditions and the result of Theorem 3, an illustration on a 2-dimensional system is provided in the example below.

Example 1. Consider a 2-dimensional system ($m = 2$), where $\text{Post}(q_j | \mathcal{T}_a)$ is a rectangle given by $[v_l^{(1)}, v_u^{(1)}] \times [v_l^{(2)}, v_u^{(2)}]$, and $\text{Post}(q'_i | \mathcal{T}_a)$ is defined by the following four inequalities for $y \in \mathbb{R}^2$:

$$h_{i1} y^{(1)} + h_{i2} y^{(2)} \leq b^{(i)}, \quad i \in \{1, 2, 3, 4\}.$$

Theorem 1 guarantees that if point $\bar{y} = (\frac{v_l^{(1)} + v_u^{(1)}}{2}, \frac{v_l^{(2)} + v_u^{(2)}}{2})$ satisfies all the four inequalities (Condition 1), then

$$\max_{x \in q_i} T^d(q_j | x, a, \Delta t) = f(\bar{y}).$$

If not, then the maximum is given by a point on the lines

$$h_{i1} y^{(1)} + h_{i2} y^{(2)} = b^{(i)}, \quad i \in \{1, 2, 3, 4\}.$$

That point is either a vertex, i.e., an intersection of two lines, (Condition 2) or where

$$\frac{\partial f(y)}{\partial y^{(1)}} = h_{i1}\mu_i, \quad \text{and} \quad \frac{\partial f(y)}{\partial y^{(2)}} = h_{i2}\mu_i$$

(Condition 3), resulting in

$$\frac{\partial f(y)}{\partial y^{(2)}} = \frac{h_{i2}}{h_{i1}} \frac{\partial f(y)}{\partial y^{(1)}}.$$

Note that $\frac{h_{i2}}{h_{i1}}$ is the slope of the line perpendicular to the i -th boundary line. This means that a non-vertex candidate point on line i is where the gradient of $f(y)$ becomes perpendicular to the line. It is a valid candidate if it is on the line segment that satisfies the four inequalities, i.e., on the boundary of $\text{Post}(q'_i|T)$. Let Y be the set containing all such valid candidate points and the 4 vertices, i.e., $4 \leq |Y| \leq 8$. Then,

$$\max_{x \in q_i} T^d(q_j | x, a, \Delta t) = \max_{\bar{y} \in Y} f(\bar{y}).$$

For the minimum, all four vertices need to be considered (Condition 2) in addition to the valid points where the gradient of $-f(y)$ becomes perpendicular to the boundary line (Condition 4). Let Y' denote the set of these points, i.e., $4 \leq |Y'| \leq 8$. Then,

$$\min_{x \in q_i} T^d(q_j | x, a, \Delta t) = \min_{\bar{y} \in Y'} f(\bar{y}).$$

It is important to note that Theorem 1 enables to find the values for (22) without the need for sampling-based methods and/or for over-approximations of the error, as currently done in the literature, e.g., [1], [9], [12]. Thus, in addition to computing the exact values for those quantities, it also speeds up the process of building the abstraction.

D. Bounds on the Continuous Transition Kernel

Here, we derive bounds for the quantities

$$\min_{x_1 \in q_i} T^c(\Delta t | x_1, q_j, a), \quad \max_{x_1 \in q_i} T^c(\Delta t | x_1, q_j, a). \quad (25)$$

To this end, we first need to introduce the notion of bridge of a stochastic process [39]. Intuitively, the *bridge* of a stochastic process is also a stochastic process for which both initial and final states are known. Formally, for $\mathbf{x}(t)$ with given initial and final states $x_1, x_2 \in \mathbb{R}^m$, fixed action $a \in A$, and duration $\Delta t \in \mathbb{R}_{\geq 0}$, the bridge $\mathbf{b}_x^{x_1, x_2, a, \Delta t}(t)$ (or simply $\mathbf{b}_x(t)$ if the context is clear) is defined as

$$\begin{aligned} \mathbf{b}_x^{x_1, x_2, a, \Delta t}(t) &= \mathbf{x}(t), \quad \text{conditioned on} \\ \mathbf{x}(0) &= x_1, \quad \mathbf{x}(\Delta t) = x_2, \quad \mathbf{a}(t) = a \quad \forall t \in [0, \Delta t]. \end{aligned}$$

The following proposition guarantees that $\mathbf{b}_x(t)$ is a Gaussian process and derives its expectation and covariance from those of \mathbf{x} [39], [40].

Proposition 2 (Proposition 4 in [41]). *Let \mathbf{x} be the stochastic process described by (1). For $a \in A$, $x_1, x_2 \in \mathbb{R}^m$, $\mathbf{b}_x^{x_1, x_2, a, \Delta t}$ is a Gaussian process with a probability measure $\text{Prob}^{\mathbf{b}_x}$ such*

that $\mathbf{b}_x(t) \sim \mathcal{N}(E_{\mathbf{b}_x}(t), \text{Cov}_{\mathbf{b}_x}(t))$, and with the following expectation and covariance

$$E_{\mathbf{b}_x}(t) = \text{Cov}_{\mathbf{x}}(\Delta t, t) \text{Cov}_{\mathbf{x}}(\Delta t)^{-1} (x_2 - E_{\mathbf{x}}(\Delta t)) + E_{\mathbf{x}}(t), \quad (26)$$

$$\text{Cov}_{\mathbf{b}_x}(t) = \text{Cov}_{\mathbf{x}}(t) - \text{Cov}_{\mathbf{x}}(\Delta t, t) \text{Cov}_{\mathbf{x}}(\Delta t)^{-1} \text{Cov}_{\mathbf{x}}(\Delta t, t), \quad (27)$$

where $\text{Cov}_{\mathbf{x}}(t_1, t_2)$ is the covariance of process \mathbf{x} at times t_1 and t_2 given by

$$\begin{aligned} \text{Cov}_{\mathbf{x}}(t_1, t_2) &= E \left[(\mathbf{x}(t_1) - E_{\mathbf{x}}(t_1)) (\mathbf{x}(t_2) - E_{\mathbf{x}}(t_2))^T \right] \\ &= e^{F(a)t_2} \text{Cov}_{\mathbf{x}}(0) (e^{F(a)t_1})^T + \\ &\quad \int_0^{\min(t_1, t_2)} e^{F(a)(t_2-u)} G(a) G(a)^T \\ &\quad (e^{F(a)(t_1-u)})^T du. \quad (28) \end{aligned}$$

Process \mathbf{b}_x describes the evolution of \mathbf{x} during $[0, \Delta t]$ given that $\mathbf{x}(0)$ and $\mathbf{x}(\Delta t)$ are known. Since \mathbf{b}_x is Gaussian, we can use the theory of Gaussian processes to derive bounds for $T^c(\Delta t | x_1, q_j, a)$. Theorem 2 focuses on the upper bound.

Theorem 2. *For polytopic regions $q_i, q_j \subseteq X_{\text{safe}}$, it holds that*

$$\max_{x_1 \in q_i} T^c(\Delta t | x_1, q_j, a) \leq \max_{\bar{x} \in \bar{q}} \int_{X_{\text{safe}}} \mathcal{N}(z | \bar{x}, \text{Cov}_{\mathbf{b}_x}(\frac{\Delta t}{2})) dz, \quad (29)$$

where

$$\begin{aligned} \bar{q} &= \text{Post}(q_i | A_1) \oplus \text{Post}(q_j | A_2), \\ A_1 &= e^{F(a)\frac{\Delta t}{2}} - \text{Cov}_{\mathbf{x}}(\Delta t, \frac{\Delta t}{2}) \text{Cov}_{\mathbf{x}}(\Delta t)^{-1} e^{F(a)\Delta t}, \\ A_2 &= \text{Cov}_{\mathbf{x}}(\Delta t, \frac{\Delta t}{2}) \text{Cov}_{\mathbf{x}}(\Delta t)^{-1}, \end{aligned}$$

and \oplus is the Minkowski sum.

Proof. Let $\mathcal{D} = \{t_1, t_2, \dots\} \subset [0, \Delta t]$ be a countable and dense subset of $[0, \Delta t]$ such that $\frac{\Delta t}{2} \in \mathcal{D}$. The upper bound on the continuous kernel from region q_i to q_j is the maximum safety probability of the bridge over all the possible initial points in q_i and final points in q_j :

$$\begin{aligned} \max_{x_1 \in q_i} T^c(\Delta t | q_j, x_1, a) &= \max_{(x_1, x_2) \in (q_i, q_j)} \text{Prob}^{\mathbf{b}_x}(\mathbf{b}_x(t_i) \in X_{\text{safe}}, \forall t_i \in \mathcal{D}) \\ &\leq \max_{(x_1, x_2) \in (q_i, q_j)} \text{Prob}^{\mathbf{b}_x}(\mathbf{b}_x(\frac{\Delta t}{2}) \in X_{\text{safe}}) \\ &= \max_{(x_1, x_2) \in (q_i, q_j)} \int_{X_{\text{safe}}} \mathcal{N}(z | E_{\mathbf{b}_x}(\frac{\Delta t}{2}), \text{Cov}_{\mathbf{b}_x}(\frac{\Delta t}{2})) dz, \end{aligned}$$

where in particular the first identity is due to the separability of \mathbf{b}_x in $[0, \Delta t]$ and the inequality is due to the fact that $\frac{\Delta t}{2} \in \mathcal{D}$. In the above optimization, the only term that depends on x_1 and x_2 is $E_{\mathbf{b}_x}(\frac{\Delta t}{2})$. By rearranging the terms in (26), we have

$$\begin{aligned} E_{\mathbf{b}_x}(\frac{\Delta t}{2}) &= \left(e^{F(a)\frac{\Delta t}{2}} - \text{Cov}_{\mathbf{x}}(\Delta t, \frac{\Delta t}{2}) \text{Cov}_{\mathbf{x}}(\Delta t)^{-1} e^{F(a)\Delta t} \right) \\ &\quad x_1 + \left(\text{Cov}_{\mathbf{x}}(\Delta t, \frac{\Delta t}{2}) \text{Cov}_{\mathbf{x}}(\Delta t)^{-1} \right) x_2. \end{aligned}$$

Therefore, $E_{\mathbf{b}_x}(\frac{\Delta t}{2})$ is the sum of linear transformations of x_1 and x_2 . Since $x_1 \in q_i$ and $x_2 \in q_j$, the above optimization can be defined as a maximization over all the points in the Minkowski sum of the linear transformations of q_i and q_j . \square

Theorem 2 shows that the upper bound of T^c is given by a constrained maximization (over a convex region \bar{q}) of the integral of a normal distribution over X_{safe} . An efficient method of computation for this maximization problem is similar to the one introduced in Section IV-C. That is to solve the problem in the transformed space where $Cov_{\mathbf{b}_x}(\frac{\Delta t}{2})$ is the identity matrix. Then, through a grid discretization of the transformed X_{safe} (the discretization needs to be over-approximating if a precise representation is not possible), the integral of the normal distribution over X_{safe} can be expressed as a summation of closed-form functions, similar to (24). Then, a straightforward extension of Theorem 1 can be used to compute the max value of this summation. For details, see Proposition 3 and the ensuing discussion.

In order to derive the lower bound of $T^c(\Delta t | q_j, x_1, a)$, we employ the 1-norm distance of a point $x \in \mathbb{R}^m$ to a set $X \subseteq \mathbb{R}^m$, defined as

$$\|x - X\|_1 = \min_{x' \in X} \|x - x'\|_1 = \min_{x' \in X} \sum_{i=1}^m |x^{(i)} - x'^{(i)}|,$$

where $x^{(i)}$ denotes the i -th component of vector x . Hence, the minimum distance from a region $q \subseteq X_{\text{safe}}$ to the boundary of X_{safe} can be written as

$$\epsilon_q = \min_{x \in q} \|x - \partial X_{\text{safe}}\|_1,$$

where ∂X_{safe} is the boundary of X_{safe} . Then, we can bound T^c from below by

$$\begin{aligned} \min_{x_1 \in q_i} T^c(\Delta t | x_1, q_j, a) &= 1 - \max_{(x_1, x_2) \in (q_i, q_j)} \text{Prob}^{\mathbf{b}_x}(\exists t \in [0, \Delta t] \text{ s.t. } \mathbf{b}_x(t) \notin X_{\text{safe}}) \\ &\geq 1 - \max_{(x_1, x_2) \in (q_i, q_j)} \text{Prob}^{\mathbf{b}_x}(\exists t \in [0, \Delta t] \text{ s.t. } \|\mathbf{b}_x(t) - x_1\|_1 \geq \epsilon_{q_i} \\ &\quad \wedge \|\mathbf{b}_x(t) - x_2\|_1 \geq \epsilon_{q_j}). \end{aligned} \quad (30)$$

The inequality in (30) holds because the set of paths of \mathbf{b}_x that remain within a distance ϵ_{q_i} from x_1 or within a distance ϵ_{q_j} from x_2 is guaranteed to be safe during $[0, \Delta t]$, and hence the complement of this set contains all the paths that are not safe during $[0, \Delta t]$. Theorem 3 (below) shows how to compute a bound on (30), but in order to arrive to this theorem, the following terms need to be defined first.

For $i \in \{1, \dots, m\}$, let $L_{\mathbf{b}^{(i)}}$ denote the maximum distance that the expectation of the i -th component of the bridge can travel. Formally,

$$L_{\mathbf{b}^{(i)}} = \sup_{(x_1, x_2, t_1, t_2) \in (q_i, q_j, [0, \Delta t], [0, \Delta t])} |E_{\mathbf{b}_x^{(i)}}(t_1) - E_{\mathbf{b}_x^{(i)}}(t_2)|,$$

where $\mathbf{b}_x^{(i)}$ is the i -th component of \mathbf{b}_x , $E_{\mathbf{b}_x}(0) = x_1$, and $E_{\mathbf{b}_x}(\Delta t) = x_2$. For $i \in \{1, \dots, m\}$, let $\bar{\mathbf{b}}_x^{(i)}(t)$ be the zero-mean Gaussian process

$$\bar{\mathbf{b}}_x^{(i)}(t) = \mathbf{b}_x^{(i)}(t) - E_{\mathbf{b}_x^{(i)}}(t).$$

Intuitively, $\bar{\mathbf{b}}_x^{(i)}(t)$ represents the i -th component of the bridge with zero mean. Furthermore, for $i \in \{1, \dots, m\}$, we use $K_{\mathbf{b}}^{d,i} > 0$ to denote the constant that bounds the expected maximum variation of the process $\bar{\mathbf{b}}_x^{(i)}(t)$. That is,

$$\sup_{t_1, t_2 \in [0, \Delta t]} d_i(t_1, t_2) \leq K_{\mathbf{b}}^{d,i} \Delta t,$$

where

$$d_i(t_1, t_2) = \sqrt{E\left[\left(\bar{\mathbf{b}}_x^{(i)}(t_1) - \bar{\mathbf{b}}_x^{(i)}(t_2)\right)^2\right]}.$$

We can now state the following theorem, which shows how to bound the inequality in (30).

Theorem 3. For regions $q_i, q_j \subseteq X_{\text{safe}}$, $i \in \{1, \dots, m\}$, let

$$\eta_i = \frac{\epsilon^*}{m} - \left(L_{\mathbf{b}^{(i)}} + 12 \int_0^{\frac{1}{2} K_{\mathbf{b}}^{d,i} \Delta t} \ln\left(\frac{2K_{\mathbf{b}}^{d,i} \Delta t}{z} + 1\right)^{\frac{1}{2}} dz\right),$$

where $\epsilon^* = \max\{\epsilon_{q_i}, \epsilon_{q_j}\}$. Assume $\eta_i > 0, \forall i \in \{1, \dots, m\}$. Then, it holds that

$$\min_{x_1 \in q_i} T^c(\Delta t | x_1, q_j, a) \geq 1 - 2 \sum_{i=1}^m e^{-\frac{\eta_i^2}{2\xi_i}},$$

where $\xi_i = \sup_{t \in [0, \Delta t]} Cov_{\bar{\mathbf{b}}_x^{(i)}}(t)$.

The proof of Theorem 3 is given in Appendix B. It makes use of the Borell-TIS inequality [37] and of Dudley's theorem [42] to bound (30).

Theorem 3 requires the computation of constant $L_{\mathbf{b}^{(i)}}$, which is always lower-bounded by

$$L_{\mathbf{b}^{(i)}} \geq \max_{(x_1, x_2) \in (q_i, q_j)} |x_1^{(i)} - x_2^{(i)}|.$$

For particular models, this bound can be tightened. For example, if F and G are diagonal matrices, F is stable (i.e., all of its eigenvalues are negative), and X_{safe} is centered on the origin, then $L_{\mathbf{b}^{(i)}} = 0$. Intuitively, this is because $E_{\mathbf{b}_x}(t)$ always points towards the origin; hence, it is always moving away from the unsafe set.

The term η_i in Theorem 3 is a constant, which represents the difference between the larger distance between q_i or q_j and the boundary of X_{safe} , scaled by m , and a term that bounds the expectation of the supremum of $\mathbf{b}_x^{(i)}$ (loosely speaking, it is the maximum distance that the system can travel in the duration of Δt in expectation) given by the Dudley's entropy integral [37]. There are two scenarios when η_i can be non-positive: (i) both q_i and q_j are very close to the boundary, and (ii) q_i and q_j are far from each other. In either case, we set the lower bound of T^c to zero. Intuitively, when both q_i and q_j are very close to the boundary, the chances that $\mathbf{b}_x^{(i)}$ for all i remains safe are very low. When q_i and q_j are far from each other, the lower bound of the transition probability between them is close to zero for small Δt .

E. Bounds on the Transition Probabilities to the Unsafe Region

Here, we focus on the transition probabilities to the unsafe state q_{unsafe} in (20) and (21) and we consider the quantities

$$\max_{x \in q_i} P_{\text{safe}}(X_{\text{safe}} | x, a, \Delta t), \quad \min_{x \in q_i} P_{\text{safe}}(X_{\text{safe}} | x, a, \Delta t).$$

We can also efficiently compute bounds for these quantities by using the results obtained above. The following proposition shows this efficient method of computation.

Proposition 3. *Let Q^a and \bar{Q}^a be two sets of polytopical regions in mode a such that*

$$\bigcup_{q \in Q^a} q \subseteq X_{\text{safe}} \subseteq \bigcup_{q \in \bar{Q}^a} q,$$

and $\text{Post}(q|\mathcal{T}_a)$ is a hyper-rectangle defined by

$$[v_{l,q}^{(1)}, v_{u,q}^{(1)}] \times \cdots \times [v_{l,q}^{(m)}, v_{u,q}^{(m)}]$$

for every $q \in Q^a \cup \bar{Q}^a$, and call

$$f(y, q) = \frac{1}{2^m} \prod_{i=1}^m \left(\text{erf}\left(\frac{y^{(i)} - v_{l,q}^{(i)}}{\sqrt{2}}\right) - \text{erf}\left(\frac{y^{(i)} - v_{u,q}^{(i)}}{\sqrt{2}}\right) \right). \quad (31)$$

Then, it holds that

$$\max_{x \in q_i} P_{\text{safe}}(X_{\text{safe}} | x, a, \Delta t) \leq \max_{y \in \text{Post}(q'_i|\mathcal{T}_a)} \sum_{q \in \bar{Q}^a} f(y, q), \quad (32)$$

$$\min_{x \in q_i} P_{\text{safe}}(X_{\text{safe}} | x, a, \Delta t) \geq \alpha \min_{y \in \text{Post}(q'_i|\mathcal{T}_a)} \sum_{q \in Q^a} f(y, q), \quad (33)$$

where $q'_i = \text{Post}(q_i | e^{F(a)\Delta t})$,

$$\alpha = \begin{cases} \max\{0, 1 - 2 \sum_{i=1}^m e^{-\frac{\eta_{\mathbf{x},i}^2}{2\xi_{\mathbf{x}}}}\} & \text{if } \eta_{\mathbf{x},i} > 0, \\ & \forall i \in \{1, \dots, m\} \\ 0 & \text{otherwise,} \end{cases}$$

$$\eta_{\mathbf{x},i} = \epsilon_{q_i} - \left(L_{\mathbf{x}}^{(i)} + 12 \int_0^{\frac{1}{2} K_{\mathbf{x}}^{d,i} \Delta t} \ln \left(\frac{K_{\mathbf{x}}^{d,i} \Delta t}{2z} + 1 \right)^{\frac{1}{2}} dz \right),$$

and $\xi_{\mathbf{x},i}$, $K_{\mathbf{x},i}^d$ and $L_{\mathbf{x}}^{(i)}$ are the constants introduced in Section IV-D, but computed for process $\mathbf{x}^{(i)}$.

The proof of this proposition is in Appendix C.

Intuitively, Proposition 3 states that, with a particular choice of discretization, i.e., a grid in the transformed space, the safe transition probability to X_{safe} is equal to the sum of the transition probabilities to the discrete regions, where each discrete transition kernel is given by the close-form function $f(y, q)$ in (31). If X_{safe} cannot be precisely discretized with a grid (in the transformed space), then the upper and lower bounds of the transition probabilities are given by the over- and under-approximating grids (\bar{Q}^a and Q^a), respectively. In both cases, the bounds need to be scaled by the time discretization error, which is captured by α in (33) for the lower bound. For the upper bound in (32), the maximum time discretization error is taken to be one. Theoretically, a tighter upper bound can be obtained by extending Theorem 2 for the time discretization error. In the experimental evaluations, however, it made no practical improvement to the bounds whilst drastically increasing the burden of computations.

Note that, for the computation of the values in (32) and (33), a straightforward extension of Theorem 1 can be used. In the

extended version, the point y at which $\sum_{q \in \bar{Q}^a} \nabla f(y, q) = 0$ needs to be considered instead of the center point in Condition 1. Condition 2 remains the same, and in Conditions 3 and 4, $\nabla f(y)$ needs to be replaced with $\sum_{q \in \bar{Q}^a} \nabla f(y, q)$ and $\sum_{q \in Q^a} \nabla f(y, q)$, respectively. This leads to an efficient computation of the transition probability bounds to q_{unsafe} , finalizing the construction of the abstraction \mathcal{I} .

V. SYNTHESIS

Given the IMDP abstraction \mathcal{I} , our goal is to synthesize a strategy that is robust against all the introduced uncertainties (errors) by the discretization of time and space domains. These uncertainties can be viewed as the nondeterministic choice of a feasible transition probability from one IMDP state to another under a given action. Therefore, we interpret the evolution of the IMDP as a 2-player stochastic game, where Player 1 chooses an action $a \in A$ at state $q \in Q$, and Player 2 chooses a feasible transition probability distribution $\gamma_q^a \in \Gamma_q^a$. This game is adversarial, where the objectives of Players 1 and 2 are to maximize and minimize the probability of remaining in the safe set, respectively. Hence, the goal becomes to synthesize a strategy for Player 1 that is robust against all adversaries.

It is worth noting that, in this 2-player stochastic game, the choice for Player 2 is from a continuous set, and hence the classical algorithms for 2-player games, e.g., [43] cannot be used. In [12], a synthesis algorithm for reachability for IMDPs is introduced. Since reachability and safety problems are dual, one can adapt that algorithm to solve for the safety problem. The dual reachability problem of safety is to find a strategy that minimizes the probability of reaching the unsafe state under all possible adversaries. We introduce a synthesis algorithm that directly solves the safety problem in this section.

Recall that τ is the required time duration for the system to remain safe. Let $k_\tau = \lceil \frac{\tau}{\Delta t} \rceil$ be the equivalent number of time steps, where $\lceil \cdot \rceil$ is the ceiling function. Also note that, given the IMDP abstraction \mathcal{I} , under a strategy σ , the probability of remaining safe from each state is necessarily a range for all the available choices of Player 2 or adversaries. Let $\check{p}_\sigma^k(q)$ and $\hat{p}_\sigma^k(q)$ denote the lower and upper bounds of the probability of remaining safe in k time steps starting from state $q \in Q$ under strategy σ , respectively. Derived from Bellman equation, we can compute the optimal lower bound by k_τ recursive evaluations of

$$\check{p}_{\sigma^*}^k(q) = \max_{a \in A(q)} \min_{\gamma_q^a \in \Gamma_q^a} \sum_{q' \in Q} \gamma_q^a(q') \check{p}_{\sigma^*}^{k-1}(q') \quad (34)$$

with initial values of $\check{p}_{\sigma^*}^0(q) = 1$ for $q \in Q_{\text{safe}}$ and $\check{p}_{\sigma^*}^0(q_{\text{unsafe}}) = 0$.

The minimization over the adversaries can be computed iteratively through an ordering of the states of \mathcal{I} [12], [35]. Let $O^\uparrow = o_1, \dots, o_{|Q|}$, where $o_i \in Q$, be an ascending ordering of the states in Q with respect to the safety probability. Then, the adversary that minimizes the safety probability in one transition is the one that assigns as much transition probability mass as possible to the states early in the ordering O^\uparrow . For the state-action pair (q, a) , let r be the state index $1 \leq r \leq |Q|$

in the ordering O^\dagger that maximizes the following expression without letting it exceed 1:

$$\sum_{i=1}^{r-1} \hat{P}(q, a, o_i) + \sum_{i=r}^{|Q|} \check{P}(q, a, o_i).$$

Then, the minimizing adversary at (q, a) is:

$$\gamma_q^{*a}(o_i) = \begin{cases} \hat{P}(q, a, o_i) & \text{if } i < r \\ \check{P}(q, a, o_i) & \text{if } i > r \end{cases}, \quad (35)$$

$$\gamma_q^{*a}(o_r) = 1 - \sum_{i=1, i \neq r}^{|Q|} \gamma_q^{*a}(o_i).$$

Once these adversaries are obtained for all $a \in A(q)$, then a maximization over the actions can be performed to complete the computation for one time step in (34). This operation can be achieved by a matrix-vector multiplication followed by a maximization, as detailed in [12]. After k_τ iterations of this algorithm, the robust strategy σ^* and the lower bound safety probability $\check{p}_{\sigma^*}^{k_\tau}(q)$ for each $q \in Q$ are obtained. The upper bounds are then given by recursive evaluations of

$$\hat{p}_{\sigma^*}^k(q) = \max_{\gamma_q^{*\sigma^*(q)} \in \Gamma_q^{\sigma^*(q)}} \sum_{q' \in Q} \gamma_q^{\sigma^*(q)}(q') \hat{p}_{\sigma^*}^{k-1}(q'), \quad (36)$$

with the initial values of $\hat{p}_{\sigma^*}^0(q) = 1$ for $q \in Q_{\text{safe}}$ and $\hat{p}_{\sigma^*}^0(q_{\text{unsafe}}) = 0$. The maximization over adversaries is obtained through (35) with a descending ordering O^\downarrow of the states.

The complexity of the above strategy synthesis algorithm is polynomial in the size of the IMDP \mathcal{I} , and the obtained strategy is a (time-dependent) Markov strategy for finite k_τ [12]. Even though this extension is beyond the scope of this work, it is worth noting that the Bellman equations in (34) and (36) are guaranteed to converge as $k \rightarrow \infty$ [12], [44]. Therefore, the safety computations can also be performed for unbounded (safety) time durations, i.e., $\tau \rightarrow \infty$, in which case the obtained strategy becomes stationary (memoryless).

VI. CORRECTNESS GUARANTEES AND COMPLETENESS ANALYSIS

The computed strategy σ^* and probability bounds $\check{p}_{\sigma^*}^{k_\tau}$ and $\hat{p}_{\sigma^*}^{k_\tau}$ on \mathcal{I} also hold for the hybrid system \mathcal{H} . Let $\mathcal{L} : \text{Paths}_{\mathcal{H}}^{\text{fin}} \rightarrow \text{Paths}_{\mathcal{H}}^{\text{fin}}$ be a function that maps the sample paths of the hybrid system \mathcal{H} to the finite paths of the IMDP \mathcal{I} through sampling time at intervals Δt , i.e., $\mathcal{L}(\omega_{\mathcal{H}}^{k\Delta t}) = q_0, q_1, \dots, q_k$ for all $k \geq 0$, where q_i is the state in \mathcal{I} that corresponds to the region $q_i^{\mathbf{a}(i\Delta t)}$ in mode $\mathbf{a}(i\Delta t)$ of \mathcal{H} such that $\mathbf{x}(i\Delta t) \in q_i^{\mathbf{a}(i\Delta t)}$ for all $0 \leq i \leq k$. Then, the IMDP strategy σ^* correctly maps to a piecewise-constant switching strategy $\sigma_{\mathcal{H}}^*$ of \mathcal{H} through

$$\sigma_{\mathcal{H}}^*(\omega_{\mathcal{H}}^{k\Delta t}) = \sigma^*(\mathcal{L}(\omega_{\mathcal{H}}^{k\Delta t})). \quad (37)$$

Using this construction of the switching strategy $\sigma_{\mathcal{H}}^*$, the following theorem shows that the safety probability bounds $\check{p}_{\sigma^*}^{k_\tau}$ and $\hat{p}_{\sigma^*}^{k_\tau}$ are guaranteed to hold for the safety of process \mathbf{s} in \mathcal{H} under $\sigma_{\mathcal{H}}^*$ for the duration of τ .

Theorem 4. *Let \mathbf{s} be the execution associated to the hybrid system \mathcal{H} and \mathcal{I} be the IMDP abstraction of \mathcal{H} for sampling time $\Delta t > 0$, as described in Section IV. Furthermore, let σ^* be the strategy on \mathcal{I} computed by solving (34) and (36) with safety probability bounds $\check{p}_{\sigma^*}^{k_\tau}$ and $\hat{p}_{\sigma^*}^{k_\tau}$, where $k_\tau = \lceil \frac{\tau}{\Delta t} \rceil$ for time duration $\tau \in \mathbb{R}_{\geq 0}$. For switching strategy $\sigma_{\mathcal{H}}^*$ constructed from σ^* per (37) and initial hybrid state $\mathbf{s}(0) = s_0 = (a_0, x_0)$ with $x_0 \in q_0$ in mode $a_0 \in A$, it holds that*

$$P_{\text{safe}}(\mathbf{s}, S_{\text{safe}}, \tau \mid \sigma_{\mathcal{H}}^*, s_0) \in [\check{p}_{\sigma^*}^{k_\tau}(q_0), \hat{p}_{\sigma^*}^{k_\tau}(q_0)], \quad (38)$$

where $P_{\text{safe}}(\mathbf{s}, S_{\text{safe}}, \tau \mid \sigma_{\mathcal{H}}^*, s_0)$ is the safety probability for the process \mathbf{s} initialized at s_0 under $\sigma_{\mathcal{H}}^*$.

Proof. By Theorems 1, 2 and 3, it holds that $\Gamma_{q_0}^a$ contains all the feasible distributions of $\mathbf{s} = (\mathbf{a}, \mathbf{x})$ during $[0, \Delta t]$ given $\mathbf{a}(t) = a$ and $\mathbf{x}(0) = x$, for all $x \in q_0$. Bounds $\check{p}_{\sigma^*}^{k_\tau}(q_0)$ and $\hat{p}_{\sigma^*}^{k_\tau}(q_0)$ are computed by solving (34) and (36). Thus, given that the system is Markov, by induction over all the discrete time intervals we have that $\check{p}_{\sigma^*}^{k_\tau}(q_0)$ is the lower bound over all possible distributions of \mathbf{s} during $[0, \tau]$, and $\hat{p}_{\sigma^*}^{k_\tau}(q_0)$ is the upper bound on the same set. \square

Theorem 4 guarantees that the safety probability of the process \mathbf{s} is contained in the safety probability interval computed on the abstraction \mathcal{I} . The size of this interval depends on the difference of the one-step transition probability bounds of \check{P} and \hat{P} in \mathcal{I} , which can be viewed as the error induced by time and space discretization of \mathcal{H} cast into the abstraction \mathcal{I} . Both of these errors can be tuned: for time, the discretization variable is Δt ; for space, the discretization variable is the largest volume of the discrete cells (partitions) denoted by $\mathcal{V}_{\text{max}}^q$. Similar to the proof in [1] for switching diffusions, it can be shown that, as both $\Delta t \rightarrow 0$ and $\mathcal{V}_{\text{max}}^q \rightarrow 0$, the error of the abstraction goes to 0, and the IMDP abstraction becomes an MDP, i.e., $\check{P}(q, a, q') \rightarrow P(q, a, q') \leftarrow \hat{P}(q, a, q')$ for all $q, q' \in Q$ and $a \in A(q)$; hence, $\check{p}_{\sigma^*}^{k_\tau}(q_0) \rightarrow P_{\text{safe}}(\mathbf{s}, S_{\text{safe}}, \tau \mid \sigma_{\mathcal{H}}^*, s_0) \leftarrow \hat{p}_{\sigma^*}^{k_\tau}(q_0)$. As a consequence, the synthesized strategy is optimal for $\Delta t \rightarrow 0$ and $\mathcal{V}_{\text{max}}^q \rightarrow 0$.

Note that the underlying assumption in Theorem 4 is that partitions $\bigcup_{q \in Q_{\text{safe}}} q = S_{\text{safe}}$. If Q_{safe} is a conservative under approximation of S_{safe} , i.e., $\bigcup_{q \in Q_{\text{safe}}} q \subset S_{\text{safe}}$, the lower bound in (38) still holds, but the upper bound may be smaller than the actual upper bound. Nevertheless, both bounds always hold for Q_{safe} , i.e., $P_{\text{safe}}(\mathbf{s}, Q_{\text{safe}}, \tau \mid \sigma_{\mathcal{H}}^*, s_0) \in [\check{p}_{\sigma^*}^{k_\tau}(q_0), \hat{p}_{\sigma^*}^{k_\tau}(q_0)]$. Also note that, for an initial continuous state $\mathbf{x}(0) = x_0$, there are $|A|$ choices for the initialization of the hybrid system \mathcal{H} , i.e., the choice of a_0 in $s_0 = (a_0, x_0)$. Let q_0^a denote the region that contains x_0 in mode $a \in A$. Then, the optimal choice of mode for s_0 is the one that corresponds to the IMDP state with maximum lower bound, i.e., $a_0 = \arg \max_{a \in A} \check{p}_{\sigma^*}^{k_\tau}(q_0^a)$.

An interesting question in the analysis of the proposed framework is how changes in the discretization variables affect the error. In the case of time discretization, Theorem 3 guarantees that the error due to the continuous transition kernel T^c for a single transition step goes to zero exponentially with Δt . However, a smaller Δt causes the distribution of the discrete transition kernel T^d to have a smaller variance, which can lead to an increase in the space discretization error due to

\mathcal{V}_{\max}^q . Therefore, in order to obtain a smaller overall error, a finer space discretization is required, causing an increase in the number of states in the abstraction IMDP. This gives rise to a trade-off in error contribution by time and space discretization, which is empirically studied in the next section. A thorough analysis on how to tune the discretization parameters is the subject of future work.

Finally, it is important to note that, in order to obtain an IMDP abstraction of process \mathbf{x} , a hybrid system modeling of \mathbf{x} is not necessarily required. One can always partition X_{safe} and compute the transition probability bounds \tilde{P} and \hat{P} through evaluations of T^c and T^d as in (20) and (21). The computations for T^c and T^d , however, are difficult for arbitrary geometry of discrete cells. As a matter of fact, this step is known to be the bottleneck of verification and synthesis for continuous stochastic processes. In our approach, we overcome this burden by a particular choice of geometry for the discretization that is dependent on the dynamics of the stochastic process under each action. The hybrid system modeling allows this unique discretization of X_{safe} under each action (mode). Therefore, we are able to obtain closed-form solutions for T^c and T^d , allowing the exact and fast computation for their values. This of course comes at the cost of increasing the number of states in the abstraction \mathcal{I} , which is a much smaller burden given the low (polynomial) computational complexity of the IMDP synthesis algorithm.

VII. EXPERIMENTAL RESULTS

We implemented our abstraction and synthesis algorithms in MATLAB and tested the performance of the framework on a set of case studies. We first considered a stochastic process with a single mode to illustrate the efficacy and scalability of the approach in the formal analysis of SDEs (Case Study 1). Then, we analyzed the trade-off between space and time discretization parameters and their effects on the overall error (Case Study 2). Finally, we considered a strategy synthesis problem on a two-mode switched stochastic process (Case Study 3). The case studies presented in this section have been distilled as particular instances of a model for a multi-room heating system, where the continuous dynamics represent the evolution of the temperature in a set of rooms (with number equal to the state space dimension), and the discrete modes represent the different heating actions, which may be turned on or off according to the selected strategy [20]. All the computations were performed on an Intel Dual Core i5 machine with 8 GB of RAM using a single thread program. Nevertheless, the implementation of the abstraction construction is highly parallelizable, allowing significant speedups.

A. Case Study 1 - Verification

We consider a stochastic process in the form of (1) with $A = \{a_1\}$,

$$F(a_1) = \begin{pmatrix} -1 & 0 \\ 0 & -0.5 \end{pmatrix}, \quad G(a_1) = \mathbf{I},$$

and $X_{\text{safe}} = [-8, 8] \times [-8, 8]$. We are interested in the computation of the safety probability for every possible initial

state of this model for the duration $\tau = 1$. In order to perform this analysis, we abstract the model to an IMDP. We use a sampling time of $\Delta t = 0.1$ and discretize X_{safe} in the transformed space induced by transformation function \mathcal{T}_{a_1} , i.e., $Post(X_{\text{safe}}|\mathcal{T}_{a_1})$, by a grid such that the size of each cell in the original space is $\Delta x = 0.5$ per side. This results in an IMDP abstraction with $|Q| = 1025$ states, and the discrete safety time steps $k_\tau = 10$. We then run our IMDP (synthesis) algorithm to obtain the lower and upper bounds of the probability of safety, i.e., $\check{p}^{10}(q)$ and $\hat{p}^{10}(q)$, for each possible initial state $q \in Q$.

Figure 2a shows X_{safe} and the space discretization, where the lower bound of the probability of safety \check{p}^{10} is marked as a shade of gray. The obtained upper bound probabilities were $\hat{p}^{10}(q) \geq 0.97$ for all $q \in Q$. It took a total of 175 seconds to compute the abstraction and the safety probability bounds.

In order to thoroughly analyze the error introduced by the discretization variables, we consider the local error term

$$\varepsilon_q = \hat{p}^{k_\tau}(q) - \check{p}^{k_\tau}(q),$$

and introduce the error metrics of ε_{med} , which is the median of ε_q for all $q \in Q$, and ε_{ave} , which is the average of ε_q per unite volume given by

$$\varepsilon_{\text{ave}} = \frac{\sum_{q \in Q} \varepsilon_q \mathcal{V}_q}{\sum_{q \in Q} \mathcal{V}_q},$$

where \mathcal{V}_q is the area (volume) of the region associated to state q . For the chosen discretization variables $\Delta t = 0.1$ and $\Delta x = 0.5$, we obtain $\varepsilon_{\text{med}} = 0.08$ and $\varepsilon_{\text{ave}} = 0.35$. The smaller median error (than the average error) means that at least half of the cells have small errors $\varepsilon_q \leq 0.08$, and a small number of the cells have very large errors. This can be explained by analyzing Figure 2a. It is easy to observe that $\check{p}^{10}(q) \approx 0$ for the cells $q \in Q$ near ∂X_{safe} , the boundary of X_{safe} , i.e., $\varepsilon_q \approx 1$. In turn, for the cells $q \in Q$ that are not close to the boundary, $\varepsilon_q \approx 0$. This observation is aligned with the theory and the intuition as expressed in (32) and (33). That is, \mathbf{x} is a continuous-time diffusion process with a stable drift term $F(a_1)$ (the real eigenvalues are negative); hence, the process gets a strong pull towards the center, resulting in large upper bounds for safety probability in snapshots of Δt (discrete transition kernel), as suggested by (32). However, during this time, the probability of its continuous trajectory leaving X_{safe} is high due to the stochastic nature of the process if \mathbf{x} starts near ∂X_{safe} , resulting in near-zero lower bound for the continuous transition kernel, i.e., $\alpha \approx 0$ in (33). This suggests that the use of smaller cells near ∂X_{safe} can potentially reduce the spread of this effect (large error) to the nearby states, i.e., adaptive (non-uniform) gridding, as in [11].

B. Case Study 2 - Analysis of Error Trade-off

In order to explore the trade-offs between the time and space discretization parameters and their effects on the overall error, we consider the same system as in Case Study 1 (Section VII-A) and perform safety analysis for the duration $\tau = 1$ with various values for $\Delta t \in \{0.05, 0.10, 0.15, 0.20, 0.25\}$ and $\Delta x \in \{0.20, 0.32, 0.40, 0.50, 0.64\}$. The obtained results are shown in Figure 3, illustrating the change in the distribution of ε_q for the different values of Δt and Δx .




Figure 2: Verification and synthesis results for case studies 1 and 3. The safe set X_{safe} is a rectangle, and the shades of gray correspond to the lower bound of the safety probability, where black and white correspond to values 0 and 1, respectively.

Figure 3a is the box plot of ε_q , showing the distribution of $\varepsilon_q \forall q \in Q$ for each pair $(\Delta x, \Delta t)$ along with ε_{med} (indicated with a red line) and ε_{ave} (indicated with a black star). A general observation is that, by decreasing Δx , the distribution of the error becomes narrower and its average value decreases. In Figure 3b, we plot ε_{ave} as a function of Δx and Δt . For each Δt , ε_{ave} tends to decrease linearly with Δx . This is intuitive because a smaller Δx means a smaller error due to space discretization for the same Δt . However, if we fix Δx and consider different values of Δt , we have a substantially different behavior: the error curve becomes parabolic. That is, for each Δx , there is an optimal value of Δt that minimizes ε_{ave} , and this value of Δt is not necessarily the smallest one. This is due to the fact that, as discussed in Section VI, decreasing Δt causes a smaller variance in the discrete transition kernel, which means smaller cells are required to obtain a smaller overall error. Moreover, the smaller the Δx is, the smaller the optimal value of Δt becomes. This means that, for a given finite Δx , if Δt is chosen to be too small, some of the dynamics of the system may not be appropriately reflected in the abstraction, hence resulting in large errors.

This case study shows that the choice of values for discretization parameters Δt and Δx in order to obtain a desired overall error is not trivial. One possible approach in tuning these parameters is to adapt a refinement technique to iteratively refine Δt and Δx to arrive at the desired overall error value. The proposed framework enables such an approach due to the low computation cost of abstraction and synthesis algorithms. Another approach is to obtain a conservative approximation of the optimal value of Δx for a given Δt before performing the computations, as shown in [1]. However, since such approaches tend to be very conservative compared to our approach, this results in a very small Δx to guarantee a given overall error, leading to state explosion.

C. Case Study 3 - Switching Strategy Synthesis

We add a second mode a_2 to the process in Case Study 1 in Section VII-A, i.e., $A = \{a_1, a_2\}$, where

$$F(a_2) = \begin{pmatrix} -0.5 & 0.1 \\ 0 & -1 \end{pmatrix}, \quad G(a_2) = \mathbf{I}.$$

We are interested in synthesizing a switching strategy that maximizes the probability of safety at every possible initial state with the duration of $\tau = 1$ for the same X_{safe} as in Case Study 1. In abstraction, we again use a sampling time $\Delta t = 0.1$. For space discretization, we use an adaptive grid such that the resulted cells have the maximum and minimum cells sizes of $\Delta x_{\text{max}} = 0.5$ and $\Delta x_{\text{min}} = 0.1$ in the original space. Our adaptive grid algorithm first over-approximates $\text{Post}(X_{\text{safe}}|\mathcal{T}_{a_i})$ for $i \in \{1, 2\}$ by using a uniform grid with the allowed maximum-sized cells. Then, it refines the cells that are partly unsafe up to the resolution of the minimum-sized cells to under-approximate the safe set. Figures 2b and 2c show the discretization of modes a_1 and a_2 , respectively. Note that, in mode a_2 , the cells are rotated which indicate that the transformation function \mathcal{T}_{a_2} includes a rotation in addition to translation. The IMDP has a total of $|Q| = 2947$ states with $|Q^{a_1}| = 1024$, $|Q^{a_2}| = 1922$, and one q_{unsafe} . The total abstraction took 54 minutes.

We then run our synthesis algorithm to obtain the robust strategy σ^* with its corresponding safety probability bounds $\tilde{p}_{\sigma^*}^{10}(q)$ and $\hat{p}_{\sigma^*}^{10}(q)$ for all $q \in Q$. In Figures 2b and 2c the lower bounds $\tilde{p}_{\sigma^*}^{10}$ are shown as shades of gray. The upper bounds were $\hat{p}_{\sigma^*}^{10}(q) \geq 0.93 \forall q \in Q$. The synthesis computation took 15 seconds.

Call $\bar{\sigma}$ the strategy that always picks action a_1 . This strategy gives rise to the results obtained in Case Study 1 and presented in Figure 2a. Comparing it with the results of the strategy σ^* in Figure 2b, it becomes evident that σ^* improves the safety probability bounds, especially for the states near the top and bottom sides of X_{safe} . This means that σ^* takes advantage of the drift in mode a_2 to keep the process safe in those sections of X_{safe} . More precisely, the median and average errors obtained for mode a_1 under σ^* are $\varepsilon_{\text{med}} = 0.01$ and $\varepsilon_{\text{ave}} = 0.28$, which show a reduction from the errors under $\bar{\sigma}$ where $\varepsilon_{\text{med}} = 0.08$ and $\varepsilon_{\text{ave}} = 0.35$.

VIII. CONCLUSIONS

In this work, we proposed a theoretical and computational framework for analysis and synthesis for switched stochastic systems that evolve in continuous time. The framework mitigates the problem of state explosion through a suitable choice




Figure 3: Error distribution for the process in Case Study 2 with different values of Δt and Δx . (a) Box plot of ε_q (blue boxes), ε_{med} (red lines), and ε_{ave} (black stars). (b) 3-D plot of ε_{ave} as a function of Δt and Δx .

of the abstraction model and the derivation of discretization methods for both time and space domains that result in tight error bounds (exact for space). The latter is specifically enabled by stochastic hybrid system modeling and a novel dynamic-dependent space discretization. This leads to fine and compact abstractions, whose computation is fast. Even though the framework is presented for synthesis problems with safety properties, it can be extended, in a straightforward manner, to verification and synthesis for more complex and even multi-objective properties expressed in, e.g., PCTL and CSL.

One of the main results of this study is the inherent trade-off in the error contribution by the time and space discretization parameters. An empirical analysis of this trade-off is performed in this work. An interesting direction for future work is a thorough analysis on how the discretization parameters can be tuned to obtain a desired overall error. One potential candidate is sequential adaptive gridding in space discretization and iterative refinement of the parameters.

REFERENCES

- [1] L. Laurenti, A. Abate, L. Bortolussi, L. Cardelli, M. Ceska, and M. Kwiatkowska, “Reachability computation for switching diffusions: Finite abstractions with certifiable and tuneable precision,” in *Int Conf. on Hybrid Systems: Computation and Control*. ACM, 2017, pp. 55–64.
- [2] G. Yin and C. Zhu, *Hybrid switching diffusions: properties and applications*. Springer New York, 2010, vol. 63.
- [3] C. G. Cassandras and J. Lygeros, *Stochastic hybrid systems*. CRC Press, 2006, vol. 24.
- [4] R. Luna, M. Lahijanian, M. Moll, and L. E. Kavragi, “Asymptotically optimal stochastic motion planning with temporal goals,” in *Int’l Workshop on the Algorithmic Foundations of Robotics (WAFR)*, Istanbul, Turkey, Aug. 2014, pp. 335–352.
- [5] S. Haesaert, N. Cauchi, and A. Abate, “Certified policy synthesis for general Markov decision processes: An application in building automation systems,” *Perfor. Eval.*, vol. 117, pp. 75–103, 2017.
- [6] A. A. Julius, Á. Halász, M. S. Sakar, H. Rubin, V. Kumar, and G. J. Pappas, “Stochastic modeling and control of biological systems: the lactose regulation system of *escherichia coli*,” *IEEE Transactions on Automatic Control*, vol. 53, no. Special Issue, pp. 51–65, 2008.
- [7] L. Cardelli, M. Kwiatkowska, and L. Laurenti, “A stochastic hybrid approximation for chemical kinetics based on the linear noise approximation,” in *International Conference on Computational Methods in Systems Biology*. Springer, 2016, pp. 147–167.
- [8] W. Glover and J. Lygeros, “A stochastic hybrid model for air traffic control simulation,” in *International Workshop on Hybrid Systems: Computation and Control*. Springer, 2004, pp. 372–386.
- [9] A. Abate, “Probabilistic bisimulations of switching and resetting diffusions,” in *Conf. on Decision and Control*. IEEE, 2010, pp. 5918–5923.
- [10] M. Lahijanian, S. B. Andersson, and C. Belta, “Approximate Markovian abstractions for linear stochastic systems,” in *Int. Conf. on Decision and Control*. IEEE, 2012, pp. 5966–5971.
- [11] S. Esmail Zadeh Soudjani and A. Abate, “Adaptive and sequential gridding procedures for the abstraction and verification of stochastic processes,” *SIAM J. on Applied Dyn. Sys.*, vol. 12, no. 2, pp. 921–956, 2013.
- [12] M. Lahijanian, S. B. Andersson, and C. Belta, “Formal verification and synthesis for discrete-time stochastic systems,” *IEEE Transactions on Automatic Control*, vol. 60, no. 8, pp. 2031–2045, 2015.
- [13] N. Cauchi, L. Laurenti, M. Lahijanian, A. Abate, M. Kwiatkowska, and L. Cardelli, “Efficiency through uncertainty: Scalable formal synthesis for stochastic hybrid systems,” in *International Conference on Hybrid Systems: Computation and Control*. ACM, Apr. 2019, pp. 240–251.
- [14] S. Summers and J. Lygeros, “Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem,” *Automatica*, vol. 46, no. 12, pp. 1951 – 1961, 2010.
- [15] A. P. Vinod, B. Homchaudhuri, and M. M. Oishi, “Forward stochastic reachability analysis for uncontrolled linear systems using fourier transforms,” in *Int. Conf. on Hybrid Systems: Computation and Control*. ACM, 2017, pp. 35–44.
- [16] H. Bohnenkamp, P. R. d’Argenio, H. Hermanns, and J.-P. Katoen, “Modest: A compositional modeling formalism for hard and softly timed systems,” *IEEE Trans. on Soft. Eng.*, vol. 32, no. 10, pp. 812–830, 2006.
- [17] M. Zamani, P. M. 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.
- [18] H. Kushner and P. G. Dupuis, *Numerical methods for stochastic control problems in continuous time*. Springer Science & Business Media, 2013, vol. 24.
- [19] C. Baier, J.-P. Katoen *et al.*, *Principles of model checking*. MIT press Cambridge, 2008, vol. 26202649.
- [20] 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.
- [21] I. Tkachev, A. Mereacre, J.-P. Katoen, and A. Abate, “Quantitative model-checking of controlled discrete-time markov processes,” *Information and Computation*, vol. 253, pp. 1–35, 2017.
- [22] E. M. Hahn, V. Hashemi, H. Hermanns, M. Lahijanian, and A. Turrini, “Multi-objective robust strategy synthesis for interval Markov decision processes,” in *Int. Conf. on Quantitative Evaluation of SysTems (QEST)*. Berlin, Germany: Springer, Sep. 2017, pp. 207–223.
- [23] M. H. Davis, “Piecewise-deterministic Markov processes: A general class of non-diffusion stochastic models,” *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 353–388, 1984.
- [24] J. Zabczyk, “Controllability of stochastic linear systems,” *Systems & Control Letters*, vol. 1, no. 1, pp. 25–31, 1981.
- [25] B. Øksendal, *Stochastic differential equations*. Springer, 2003.
- [26] V. S. Borkar *et al.*, “Controlled diffusion processes,” *Probability surveys*, vol. 2, pp. 213–244, 2005.

- [27] A. Budhiraja, "An ergodic control problem for constrained diffusion processes: Existence of optimal markov control," *SIAM journal on control and optimization*, vol. 42, no. 2, pp. 532–558, 2003.
- [28] L. M. Bujorianu, *Stochastic reachability analysis of hybrid systems*. Springer Science & Business Media, 2012.
- [29] P. Billingsley, *Probability and measure*. John Wiley & Sons, 1995.
- [30] M. L. Bujorianu, "Extended stochastic hybrid systems and their reachability problem," in *International Workshop on Hybrid Systems: Computation and Control*. Springer, 2004, pp. 234–249.
- [31] X. Koutsoukos and D. Riley, "Computational methods for reachability analysis of stochastic hybrid systems," in *Int. Workshop on Hybrid Systems: Comput. and Control*. Springer, 2006, pp. 377–391.
- [32] A. Fleig and R. Guglielmi, "Optimal control of the fokker–planck equation with space-dependent controls," *Journal of Optimization Theory and Applications*, vol. 174, no. 2, pp. 408–427, 2017.
- [33] M. Annunziato and A. Borzi, "A fokker–planck control framework for multidimensional stochastic processes," *Journal of Computational and Applied Mathematics*, vol. 237, no. 1, pp. 487–507, 2013.
- [34] I. Tkachev and A. Abate, "Characterization and computation of infinite-horizon specifications over markov processes," *Theoretical Computer Science*, vol. 515, pp. 1–18, 2014.
- [35] R. Givan, S. Leach, and T. Dean, "Bounded-parameter Markov decision processes," *Artificial Intelligence*, vol. 122, no. 1-2, pp. 71–109, 2000.
- [36] B. Grünbaum, V. Klee, M. A. Perles, and G. C. Shephard, *Convex polytopes*. Springer, 1967, vol. 16.
- [37] R. J. Adler and J. E. Taylor, *Random fields and geometry*. Springer Science & Business Media, 2009.
- [38] D. P. Bertsekas, *Constrained optimization and Lagrange multiplier methods*. Academic press, 2014.
- [39] Y. Chen and T. Georgiou, "Stochastic bridges of linear systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 526–531, 2016.
- [40] T. Sottinen and A. Yazigi, "Generalized Gaussian bridges," *Stochastic Processes and their Applications*, vol. 124, no. 9, pp. 3084–3105, 2014.
- [41] D. Gasbarra, T. Sottinen, and E. Valkeila, "Gaussian bridges," in *Stochastic analysis and applications*. Springer, 2007, pp. 361–382.
- [42] R. M. Dudley, "The sizes of compact subsets of Hilbert space and continuity of Gaussian processes," *Journal of Functional Analysis*, vol. 1, no. 3, pp. 290–330, 1967.
- [43] L. S. Shapley, "Stochastic games," *Proceedings of the national academy of sciences*, vol. 39, no. 10, pp. 1095–1100, 1953.
- [44] D. Wu and X. Koutsoukos, "Reachability analysis of uncertain systems using bounded-parameter Markov decision processes," *Artificial Intelligence*, vol. 172, no. 8-9, pp. 945–954, 2008.

APPENDIX A

Proof of Theorem 1. We first consider the maximum case and then discuss the minimum case. The KKT conditions guarantee that if $y \in \text{Post}(q'_i|\mathcal{T}_a)$ is a local maximum for f , then there must exist a vector of constants $\mu = (\mu_1, \dots, \mu_k)$ such that $\nabla f(y) = H^T \mu$, $\mu_i \geq 0$ for all $i \in \{1, \dots, k\}$, and $\mu_i (\sum_{j=1}^m H^{(i,j)} y^{(j)} - b_i) = 0$, where $H^{(i,j)}$ is the component in the i -th row and j -th column of matrix H . Note that we have a constant μ_i , $i \in \{1, \dots, k\}$, for each of the half-spaces defining $\text{Post}(q'_i|\mathcal{T}_a)$. Thus, there are three possible cases:

Case 1: x^* is not in the boundary of $\text{Post}(q'_i|\mathcal{T}_a)$. In this case the KKT conditions imply that y is a maximum only if $\nabla f(y) = 0$. For a normal distribution with identity covariance, this point is exactly $y = (\frac{v_u^{(1)} + v_l^{(1)}}{2}, \dots, \frac{v_u^{(m)} + v_l^{(m)}}{2})$. If $y \in \text{Post}(q'_i|\mathcal{T}_a)$, then this is the global maximum, because it is the global maximum of the unconstrained problem.

Case 2: x^* is a vertex of $\text{Post}(q'_i|\mathcal{T}_a)$. We call a vertex an intersection of m half-spaces. As a consequence, we have that the KKT conditions are satisfied in y , vertex of $\text{Post}(q'_i|\mathcal{T}_a)$, if and only if $\nabla f(y) = \bar{H}^T \mu$, where \bar{H} is the submatrix that contains only the m rows of H representing the half-spaces interesting at y , and vector μ contains only the m corresponding constants. Thus, we have a system of m equations and m variables that has solution for $\mu_i \in \mathbb{R}$.

However, since the set of vertices is finite, it is generally faster to just include all the vertices as possible candidate solutions instead of solving the system of equations.

Case 3: y is in the boundary of $\text{Post}(q'_i|\mathcal{T}_a)$, but is not a vertex. In this case only $r < m$ of the half-spaces in H intersect at y . Thus, if y is a maximum then $\nabla f(y) = \bar{H}^T \mu$, where \bar{H} is the submatrix of H containing the $r < m$ half-spaces intersecting at y , and μ contains only the r corresponding constants. Note that this is a system with more equations than variables. Therefore, only when some of constraints become linearly dependent, there may be a solution for $y \in \text{Post}(q'_i|\mathcal{T}_a)$, if at all.

The minimum case is identical except that condition $\nabla f(y) = H^T \mu$ is replaced with $\nabla f(y) = -H^T \mu$. \square

APPENDIX B

Proof of Theorem 3. By definition of $T^c(\Delta t|q_j, x_1, a)$ we have that

$$\begin{aligned} & \min_{x_1 \in q_i} T^c(\Delta t|x_1, q_j, a) \\ & \geq 1 - \max_{(x_1, x_2) \in (q_i, q_j)} \text{Prob}^{\mathbf{b}_x}(\exists t \in [0, \Delta t], \mathbf{b}_x(t) \notin X_{\text{safe}}) \\ & \geq 1 - \max_{(x_1, x_2) \in (q_i, q_j)} \text{Prob}^{\mathbf{b}_x}(\exists t \in [0, \Delta t] \text{ s.t. } \|\mathbf{b}_x(t) - x_1\|_1 \geq \epsilon_{q_i} \\ & \quad \wedge \|\mathbf{b}_x(t) - x_2\|_1 \geq \epsilon_{q_j}). \end{aligned}$$

The last inequality exploits the fact that each path of \mathbf{b}_x that reaches a state outside X_{safe} during $[0, \Delta t]$ travels a distance of at least $\epsilon_q = \max\{\epsilon_{q_i}, \epsilon_{q_j}\}$. Without any loss of generality, assume $\epsilon_q = \epsilon_{q_i}$. Then, we have

$$\begin{aligned} & \min_{x_1 \in q_i} T^c(\Delta t|x_1, q_j, a) \\ & = 1 - \max_{(x_1, x_2) \in (q_1, q_2)} \text{Prob}^{\mathbf{b}_x}(\sup_{t \in [0, \Delta t]} \|\mathbf{b}_x(t) - x_1\|_1 \geq \epsilon_q) \\ & \geq 1 - \max_{(x_1, x_2) \in (q_1, q_2)} \text{Prob}^{\mathbf{b}_x}(\forall i \in \{1, \dots, m\} (\sup_{t \in [0, \Delta t]} |\mathbf{b}_x^{(i)}(t) - x_1^{(i)}| \geq \frac{\epsilon_q}{m})) \\ & \geq 1 - \max_{(x_1, x_2) \in (q_1, q_2)} \sum_{i \in \{1, \dots, m\}} \text{Prob}^{\mathbf{b}_x}(\sup_{t \in [0, \Delta t]} |\mathbf{b}_x^{(i)}(t) - x_1^{(i)}| \geq \frac{\epsilon_q}{m}). \end{aligned}$$

The last inequality holds due to the union bound. By the linearity of Gaussian processes, we can write $\mathbf{b}_x(t) = E_{\mathbf{b}_x}(t) + \bar{\mathbf{b}}_x(t)$, where $\bar{\mathbf{b}}_x(t)$ is a normally distributed random variable with zero-mean and covariance equal to the one of $\mathbf{b}_x(t)$. By the application of the triangle inequality, we then have

$$\begin{aligned} & \min_{x_1 \in q_i} T^c(\Delta t|x_1, q_j, a) \\ & \geq 1 - \max_{(x_1, x_2) \in (q_i, q_j)} \sum_{i \in \{1, \dots, m\}} \text{Prob}^{\bar{\mathbf{b}}_x}((\sup_{t \in [0, \Delta t]} |E_{\mathbf{b}_x}^{(i)}(t) - x_1^{(i)}| + |\bar{\mathbf{b}}_x^{(i)}(t_1)| \geq \frac{\epsilon_q}{m})) \\ & = 1 - \max_{(x_1, x_2) \in (q_i, q_j)} \sum_{i \in \{1, \dots, m\}} \text{Prob}^{\bar{\mathbf{b}}_x}(\sup_{t \in [0, \Delta t]} |\bar{\mathbf{b}}_x^{(i)}(t_1)|_1 \geq \\ & \quad \frac{\epsilon_q}{m} - \sup_{t_1 \in [0, \Delta t]} |E_{\mathbf{b}_x}^{(i)}(t_1) - x_1^i|) \\ & \geq 1 - 2 \sum_{i \in \{1, \dots, m\}} \text{Prob}^{\bar{\mathbf{b}}_x}(\sup_{t \in [0, \Delta t]} \bar{\mathbf{b}}_x^{(i)}(t_1) \geq \bar{\epsilon}_{q,i}), \end{aligned}$$

where $\bar{\epsilon}_{q,i} = \frac{\epsilon_q}{m} - L_{\mathbf{b}^{(i)}}$ for $L_{\mathbf{b}^{(i)}} = \sup_{(x_1, x_2, t) \in [q_i, q_j, [0, \Delta t]]} |E_{\mathbf{b}_x^{(i)}}(t) - x_1^{(i)}|$. The last inequality holds because $Cov_{\mathbf{b}_x}(t, s)$ is independent of x_1, x_2 . Thus, the only term depending on x_1, x_2 is ϵ_q . Since $\bar{\mathbf{b}}_x^{(i)}(t)$ is a uni-dimensional Gaussian Process (GP), we can bound its supremum during $[0, \Delta t]$ by using the Borell-TIS inequality [37]. For $D > E \left[\sup_{t \in [0, \Delta t]} \bar{\mathbf{b}}_x^{(i)}(t) \right]$, the Borell-TIS inequality [37] guarantees that

$$Prob^{\bar{\mathbf{b}}_x} \left(\sup_{t \in [0, \Delta t]} \bar{\mathbf{b}}_x^{(i)}(t) > D \right) \leq e^{-\frac{(D - E \left[\sup_{t \in [0, \Delta t]} \bar{\mathbf{b}}_x^{(i)}(t) \right])^2}{2\xi^{(i)}}},$$

where $\xi^{(i)} = \sup_{t \in [0, \Delta t]} Cov_{\bar{\mathbf{b}}_x^{(i)}}(t)$.

In order to bound $E \left[\sup_{t \in [0, \Delta t]} \bar{\mathbf{b}}_x^{(i)}(t) \right]$, we consider the pseudometric d_i and constant $K_{\mathbf{b}}^{d,i}$ as defined in Sec. IV-D. By the Dudley's entropy integral [37], we obtain

$$E \left[\sup_{t \in [0, \Delta t]} \bar{\mathbf{b}}_x^{(i)}(t) \right] \leq 12 \int_0^{\frac{K_{\mathbf{b}}^{d,i} \Delta t}{2}} \ln \left(\frac{2K_{\mathbf{b}}^{d,i} \Delta t}{x} + 1 \right)^{\frac{1}{2}} dx.$$

By substituting this expression in the Borell-TIS inequality and picking $D = \frac{\bar{\epsilon}_{q,i}}{m}$, we arrive to the result. \square

APPENDIX C

Proof of Proposition 3. For the upper bound, we have that for $q_i \in Q_{\text{safe}}$ and $a \in A$,

$$\begin{aligned} \max_{x \in q_i} P_{\text{safe}}(X_{\text{safe}}, \Delta t | x, a) &\leq \left(\max_{x \in q_i} T^d(X_{\text{safe}} | x, a, \Delta t) \right) \left(\max T^c(\Delta t | x, X_{\text{safe}}, a) \right) \\ &\leq \max_{x \in q_i} T^d(X_{\text{safe}} | x, a, \Delta t) \\ &\leq \max_{x \in q_i} \int_{X_{\text{safe}}} \mathcal{N}(z | E_x(\Delta t), Cov_x(\Delta t)) dz \\ &\leq \max_{y \in Post(q'_i | \mathcal{T}_a)} \sum_{q \in \bar{Q}^a} \int_{Post(q | \mathcal{T}_a)} \mathcal{N}(z | y, \mathbf{I}) dz \\ &= \max_{y \in Post(q'_i | \mathcal{T}_a)} \sum_{q \in \bar{Q}^a} f(y, q). \end{aligned}$$

For the lower bound, we have that

$$\min_{x \in q_i} P_{\text{safe}}(X_{\text{safe}}, \Delta t | x, a) \geq \left(\min_{x \in q_i} T^d(X_{\text{safe}} | x, a, \Delta t) \right) \left(\min_{x \in q_i} T^c(\Delta t | x, Q_{\text{safe}}, a) \right).$$


Similarly to the upper bound, we have that

$$\min_{x \in q_i} T^d(X_{\text{safe}} | x, a, \Delta t) \geq \min_{y \in Post(q'_i | \mathcal{T}_a)} \sum_{q \in \bar{Q}^a} f(y, q).$$

For T^c , by using a similar reasoning as in the one in Theorem 3, we have that

$$\min_{x \in q_i} T^c(\Delta t | x, Q_{\text{safe}}, a) \geq \begin{cases} \max\{0, 1 - 2 \sum_{i=1}^m e^{-\frac{\eta_{\mathbf{x},i}^2}{2\xi_{\mathbf{x}}}}\} & \text{if } \eta_{\mathbf{x},i} > 0, \forall i \in \{1, \dots, m\} \\ 0 & \text{otherwise.} \end{cases}$$

\square



Luca Laurenti is a Research Associate in the Department of Computer Science, Oxford. He received his PhD from the Department of Computer Science, University of Oxford (UK) (2018) where he was a member of the Trinity College, B.S. degree in information engineering from University La Sapienza, Rome, and M.S. in information engineering from University of Perugia.



Morteza Lahijanian is an Assistant Professor in the Dept. of Aerospace Engineering Sciences at University of Colorado Boulder, USA and the director of Assured, Robust, and Interactive Autonomous (ARIA) Systems group. He received his B.S. in Bio-engineering from University of California, Berkeley (2005), M.S. in Mechanical Engineering from Boston University (2009), and Ph.D. in Mechanical Engineering from Boston University (2013). He conducted his postdoctoral research in the Dept. of Computer Science at Rice University. Prior to joining CU Boulder, he was a research scientist in the Dept. of Computer Science at University of Oxford, UK.



Alessandro Abate (S'02–M'08–SM'19) is Professor of Verification and Control in the Department of Computer Science at Oxford, and a fellow of the Alan Turing Institute for Data Sciences in London. He received a Laurea in Electrical Engineering in October 2002 from the University of Padova (IT), an MS in May 2004 and a PhD in December 2007, both in Electrical Engineering and Computer Sciences, at UC Berkeley (USA). He has been an International Fellow in the CS Lab at SRI International in Menlo Park (USA), and a PostDoctoral Researcher at Stanford University (USA), in the Department of Aeronautics and Astronautics. From June 2009 to mid 2013 he has been an Assistant Professor at the Delft Centre for Systems and Control, TU Delft (NL).



Luca Cardelli has a PhD in computer science from the University of Edinburgh. He worked at Bell Labs, Murray Hill, from 1982 to 1985, and at Digital Equipment Corporation, Systems Research Center in Palo Alto, from 1985 to 1997, before assuming a position at Microsoft Research, in Cambridge UK, where he was head of the Programming Principles and Tools and Security groups until 2012. Since 2014 he is also a Royal Society Research Professor at the University of Oxford. He is a Fellow of the Royal Society, an ACM Fellow, an Elected Member of the Academia Europaea, an Elected Member of AITO, and a long-standing member of EATCS.



Marta Kwiatkowska is Professor of Computing Systems and Fellow of Trinity College, University of Oxford. Prior to this she was Professor in the School of Computer Science at the University of Birmingham, Lecturer at the University of Leicester and Assistant Professor at the Jagiellonian University in Cracow, Poland. Kwiatkowska has made fundamental contributions to the theory and practice of model checking for probabilistic systems, focusing on automated techniques for verification and synthesis from quantitative specifications. She led the development of the PRISM model checker (www.prismmodelchecker.org), the leading software tool in the area and winner of the HVC Award 2016. Kwiatkowska awarded an honorary doctorate from KTH Royal Institute of Technology in Stockholm in 2014 and the Royal Society Milner Medal in 2018. She is a Fellow of ACM and Member of Academia Europaea.