






Unbounded-Time Safety Verification of Stochastic Differential Dynamics*

Shenghua Feng^{1,2} , Mingshuai Chen³ , Bai Xue^{1,2} , Sriram Sankaranarayanan⁴ , and Naijun Zhan^{1,2} 

¹ SKLCS, Institute of Software, CAS, Beijing, China

² University of Chinese Academy of Sciences, Beijing, China
{fengsh, xuebai, znj}@ios.ac.cn

³ Lehrstuhl für Informatik 2, RWTH Aachen University, Aachen, Germany
chenms@cs.rwth-aachen.de

⁴ University of Colorado, Boulder, USA
sriram.sankaranarayanan@colorado.edu



Abstract. In this paper, we propose a method for bounding the probability that a stochastic differential equation (SDE) system violates a safety specification over the infinite time horizon. SDEs are mathematical models of stochastic processes that capture how states evolve continuously in time. They are widely used in numerous applications such as engineered systems (e.g., modeling how pedestrians move in an intersection), computational finance (e.g., modeling stock option prices), and ecological processes (e.g., population change over time). Previously the safety verification problem has been tackled over finite and infinite time horizons using a diverse set of approaches. The approach in this paper attempts to connect the two views by first identifying a finite time bound, beyond which the probability of a safety violation can be bounded by a negligibly small number. This is achieved by discovering an exponential barrier certificate that proves exponentially converging bounds on the probability of safety violations over time. Once the finite time interval is found, a finite-time verification approach is used to bound the probability of violation over this interval. We demonstrate our approach over a collection of interesting examples from the literature, wherein our approach can be used to find tight bounds on the violation probability of safety properties over the infinite time horizon.

Keywords: Stochastic Differential Equations (SDEs) · Unbounded safety verification · Failure probability bound · Barrier certificates.

1 Introduction

In this paper, we investigate the problem of verifying probabilistic safety properties for continuous stochastic dynamics modeled by stochastic differential equations (SDEs).

* This work was partially funded by NSFC under grant No. 61625206, 61732001 and 61872341, by the ERC Advanced Project FRAPPANT under grant No. 787914, by the US NSF under grant No. CCF 1815983 and by the CAS Pioneer Hundred Talents Program under grant No. Y8YC235015.

The study of SDEs dates back to the 1900s when, e.g., Einstein used SDEs to model the phenomenon of Brownian motion [10]. Since then, SDEs have witnessed numerous applications including models of disturbances in engineered systems ranging from wind forces [37] to pedestrian motion [14]; models of financial instruments such as options [5]; and models of biological/ecological processes for instance predator-prey models [25]. In the meantime, SDEs are hard to reason about: they are defined using ideas from stochastic calculus that reimagine basic concepts such as integration in order to conform to the basic laws of probability and stochastic processes [24].

There are many important verification problems for SDEs. Prominent topics include the safety verification problem which seeks to know the probability that a given SDE with specified initial conditions will enter an unsafe region (or leave a safe region) over a given time horizon. Generally, safety verification can be performed over a finite-time horizon setting, wherein the probability is sought over a finite time interval $[0, T]$. On the other hand, the infinite-time horizon problem seeks a bound on the probability of satisfying a safety property over the unbounded time horizon $[0, \infty)$. A handful of methods have been proposed for verifying SDE systems, such as the barrier certificate-based methods over both the infinite time horizon [27] and finite time horizons [35], the moment optimization-based method over finite time horizons [33] and the Hamilton-Jacobi-based method over the infinite time horizon [16]. The novelty of our work lies in the reduction of infinite-time horizon verification problems to finite time problems.

In this paper, we propose a novel reduction-based method to verify unbounded-time safety properties of stochastic systems modeled as nonlinear polynomial SDEs. We employ a similar idea as in [11] (for verifying delay differential equations) that reduces the safety verification problem over the infinite time horizon to the one over a finite time interval. This is achieved by computing an *exponential stochastic barrier certificate* which witnesses an exponentially decreasing upper bound on the probability that a target system violates a given safety specification. Consequently, for any $\epsilon > 0$, we can identify a time instant T beyond which the violation (a.k.a. failure) probability is smaller than the negligibly small cutoff ϵ . The reduced bounded-time safety verification problem over $[0, T]$ can hence be tackled by any of the available methods. We furthermore present an alternative method to address the reduced finite-time horizon verification problem based on the discovery of a *time-dependent stochastic barrier certificate*. We show that both the exponential and the time-dependent stochastic barrier certificate can be synthesized by respectively solving a pertinent *semidefinite programming* (SDP) [38] optimization problem. Experimental results on some interesting examples taken from the literature demonstrated the effectiveness of the reduction and that our method often produces tighter bounds on the failure probability. Our approach has some broad similarities to related approaches in symbolic execution of probabilistic programs that conclude facts about infinitely many behaviors by analyzing finitely many paths in the program that account for a sufficient probability among all the behaviors [31].

Contributions. The main contributions of this work can be summarized as follows: (1) We reduce the unbounded-time safety verification of stochastic systems to a bounded one, based on an exponentially decreasing bound on the failure probability which guarantees the dominance of the overall failure probability by the truncated finite time hori-

zon. (2) We show how the obtained bound on the overall failure probability is tighter than that produced by existing methods for some interesting SDEs.

Related Work. The use of mathematical models of processes –ranging from finite state machines to various types of differential equations– has allowed us to reason about rich behaviors of Cyber-Physical Systems produced by the interaction between digital computers and physical plants [29]. In this regard, many modeling formalisms have been studied including finite state machines, ordinary differential equations (ODEs), timed automata, hybrid automata, etc. [8], on top of which a large variety of verification problems have been extensively investigated, e.g., safety verification through reachability analysis and temporal logic verification [3].

In the existing literature on formal verification, ODEs are often used to describe the behavior of deterministic continuous-time systems. However, these models have been shown over-simplistic in many applications that involve time delays, nondeterministic inputs and stochastic noises. SDEs hence arose as an important class of models that have been employed in practical domains covering, among others [24], financial models such as the famous Black-Scholes model used extensively in the theory of options pricing [5], wind disturbances [37], human pedestrian motion [14] and ecological models [25].

In what follows, we place our work in the context of formal verification techniques tailored for stochastic differential dynamics modeled as SDEs, and discuss contributions thereof that are highly related to our approach. Unbounded-time stochastic safety verification of SDE systems was first studied by Prajna et al. in [27,28], where a typical supermartingale was employed as a stochastic barrier certificate followed by computational conditions derived from Doob’s martingale inequality [15]. Thereafter, the stochastic barrier certificate-based method was extended to cater for bounded-time safety verification by Steinhardt and Tedrake [35] by leveraging a relaxed formulation called c -martingale for locally stable systems. The barrier certificate-based method by Prajna et al. (ibid.) for unbounded-time safety verification often leads to conservative bound on the failure probability. On the other hand, Steinhardt and Tedrake (ibid.) established impressive probability bounds but only for finite time horizons. In order to reduce the conservativeness, we propose a method of reducing the unbounded safety verification to a bounded one. Although our method in this paper is also based on the construction of stochastic barrier certificates, the gain of stochastic barrier certificates only helps to identify a finite time interval such that the violation probability of interest beyond this time interval is arbitrarily negligibly small. A time-dependent barrier certificate is further proposed to solve the resulting bounded-time safety verification. The Unbounded-time safety verification problem has also been studied by Koutsoukos and Riley [16], who linked the reachability probability to the viscosity solution of certain Hamilton-Jacobi partial differential equations, under restrictions on bounded state space and non-degenerate diffusion. Grid-based numerical approaches, e.g., the finite difference method in [16] and the level set method in [22], are traditionally used to solve these equations, leading to the fact that the Hamilton-Jacobi reachability method only scales well to systems of special structures. More recently, a novel constraint solving-based method has been proposed in [20] for algebraically over- and under-approximating the reachability probability, which is nevertheless limited to bounded-time safety verification. In addition to the above mentioned methods, we refer the readers to [7] for a

Dirichlet form-based method for stochastic hybrid systems featuring “nice” Markov properties, while to [39,6,18] and [1,17] respectively for related contributions in statistical and discrete/numerical methods for stochastic verification and control.

Finally, we mention a relation between the ideas in this paper and previously proposed ideas for (non-stochastic) ODEs due to Sogokon et al. [34]. The key similarity lies in the use of a non-negative matrix through which a vector of functions whose derivatives are related to their current value. Whereas Sogokon et al. explored this idea for ODEs, we do so for SDEs. Another significant difference, in our work, is that we use the super-martingale functions to identify a time horizon $[0, T]$ and bound the probability of safety violation beyond T .

The remainder of this paper is structured as follows. Sect. 2 introduces stochastic differential dynamics modeled by SDEs and the unbounded-time safety verification problem of interest. Sect. 3 elucidates the reduction of unbounded safety verification to bounded ones based on the witness of stochastic barrier certificates. Sect. 4 presents the SDP formulation for discovering such barrier certificates over the reduced bounded time interval. After demonstrating our method on several examples in Sect. 5, we conclude the paper in Sect. 6.

2 Problem Formulation

Notations. Let \mathbb{R} be the set of real numbers. For a vector $x \in \mathbb{R}^n$, x_i refers to its i -th component and $|x|$ denotes the ℓ^2 -norm. Particularly, $\mathbf{0}$ and $\mathbf{1}$ denote respectively the vector of zeros and ones of appropriate dimension, and the comparison between vectors, e.g., $x \leq \mathbf{0}$, is component-wise. We define for $\delta > 0$, $\mathfrak{B}(x, \delta) \triangleq \{x' \in \mathbb{R}^n \mid |x' - x| \leq \delta\}$ as the δ -closed ball centered at x . We abuse the notation $|\cdot|$ for an $m \times n$ matrix M as $|M| \triangleq \sqrt{\sum_{i=1}^m \sum_{j=1}^n |M_{ij}|^2}$. The exponential of a square matrix $M \in \mathbb{R}^{n \times n}$, denoted by e^M , is the $n \times n$ matrix given by the power series $e^M \triangleq \sum_{k=0}^{\infty} \frac{1}{k!} M^k$. For a set $\mathcal{X} \subseteq \mathbb{R}^n$, $\partial\mathcal{X}$, $\bar{\mathcal{X}}$ and \mathcal{X}° denote respectively the boundary, the closure and the interior of \mathcal{X} . Let C^k be the space of functions on \mathbb{R} with continuous derivatives up to order k ; a function $f(t, x): \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$ is in $C^{1,2}(\mathbb{R} \times \mathbb{R}^n)$ if $f \in C^1$ w.r.t. $t \in \mathbb{R}$ and $f \in C^2$ w.r.t. $x \in \mathbb{R}^n$.

Let (Ω, \mathcal{F}, P) be a probability space, where Ω is a sample space, $\mathcal{F} \subseteq 2^\Omega$ is a σ -algebra on Ω , and $P: \mathcal{F} \rightarrow [0, 1]$ is a probability measure on the measurable space (Ω, \mathcal{F}) . A *random variable* X defined on the probability space (Ω, \mathcal{F}, P) is an \mathcal{F} -measurable function $X: \Omega \rightarrow \mathbb{R}^n$; its *expectation* (w.r.t. P) is denoted by $E[X]$. Every random variable X induces a probability measure $\mu_X: \mathcal{B} \rightarrow [0, 1]$ on \mathbb{R}^n , defined as $\mu_X(B) \triangleq P(X^{-1}(B))$ for Borel sets B in the Borel σ -algebra \mathcal{B} on \mathbb{R}^n . μ_X is called the *distribution of X* ; its *support set* is $\text{supp}(\mu_X) \triangleq \bigcup_{\mu_X(B) > 0} \bar{B}$, which will also be referred to as the support of X .

A (continuous-time) *stochastic process* is a parametrized collection of random variables $\{X_t\}_{t \in T}$ where the parameter space T is interpreted as, unless explicitly notated in this paper, the halfline $[0, \infty)$. We sometimes further drop the brackets in $\{X_t\}$ when it is clear from the context. A collection $\{\mathcal{F}_t \mid t \geq 0\}$ of σ -algebras of sets in \mathcal{F} is a *filtration* if $\mathcal{F}_t \subseteq \mathcal{F}_{t+s}$ for $t, s \in [0, \infty)$. Intuitively, \mathcal{F}_t carries the information known to

an observer at time t . A random variable $\tau: \Omega \rightarrow [0, \infty)$ is called a *stopping time* w.r.t. some filtration $\{\mathcal{F}_t \mid t \geq 0\}$ of \mathcal{F} if $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t \geq 0$. A stochastic process $\{X_t\}$ adapted to a filtration $\{\mathcal{F}_t \mid t \geq 0\}$ is called a *supermartingale* if $E[X_t] < \infty$ for any $t \geq 0$ and $E[X_t \mid \mathcal{F}_s] \leq X_s$ for all $0 \leq s \leq t$. That is, the conditional expected value of any future observation, given all the past observations, is no larger than the most recent observation.

Stochastic Differential Dynamics. We consider a class of dynamical systems featuring stochastic differential dynamics governed by time-homogeneous SDEs of the form¹

$$dX_t = b(X_t) dt + \sigma(X_t) dW_t, \quad t \geq 0 \quad (1)$$

where $\{X_t\}$ is an n -dimensional continuous-time stochastic process, $\{W_t\}$ denotes an m -dimensional Wiener process (standard Brownian motion), $b: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a vector-valued polynomial flow field (called the *drift coefficient*) modeling deterministic evolution of the system, and $\sigma: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ is a matrix-valued polynomial flow field (called the *diffusion coefficient*) that encodes the coupling of the system to Gaussian white noise dW_t .

Suppose there exists a Lipschitz constant D s.t. $|b(x) - b(y)| + |\sigma(x) - \sigma(y)| \leq D|x - y|$ holds for all $x, y \in \mathbb{R}^n$. Then, given an initial state (a random variable) X_0 , an SDE of the form (1) has a unique *solution* which is a stochastic process $X_t(\omega) = X(t, \omega): [0, \infty) \times \Omega \rightarrow \mathbb{R}^n$ satisfying the stochastic integral equation (à la Itô's interpretation)

$$X_t = X_0 + \int_0^t b(X_s) ds + \int_0^t \sigma(X_s) dW_s. \quad (2)$$

The solution $\{X_t\}$ in Eq. (2) is also referred to as an (*Itô*) *diffusion process*, and will be denoted by X_t^{0, X_0} (or simply $X_t^{X_0}$), if necessary, to indicate the initial condition X_0 at $t = 0$.

A great deal of information about a diffusion process can be encoded in a partial differential operator termed the *infinitesimal generator*, which generalizes the Lie derivative that captures the evolution of a function along the diffusion process:

Definition 1 (Infinitesimal generator [24]). Let $\{X_t\}$ be a (time-homogeneous) diffusion process in \mathbb{R}^n . The infinitesimal generator \mathcal{A} of X_t is defined by

$$\mathcal{A}f(s, x) = \lim_{t \downarrow 0} \frac{E^{s, x} [f(s + t, X_t)] - f(s, x)}{t}, \quad x \in \mathbb{R}^n.$$

The set of functions $f: \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$ s.t. the limit exists at (s, x) is denoted by $\mathcal{D}_{\mathcal{A}}(s, x)$, while $\mathcal{D}_{\mathcal{A}}$ denotes the set of functions for which the limit exists for all $(s, x) \in \mathbb{R} \times \mathbb{R}^n$.

In subsequent sections, the readers may find applications of the operator \mathcal{A} to a vector-valued function in a component-wise manner. The relation between \mathcal{A} and the coefficients b, σ in SDE (1) is captured by the following result:

¹ The general time-inhomogeneous case with time-dependent b and σ can be reduced to this form (cf. [24, Chap. 10]).

Lemma 1 ([24]). *Let $\{X_t\}$ be a diffusion process defined by Eq. (1). If $f \in C^{1,2}(\mathbb{R} \times \mathbb{R}^n)$ with compact support, then $f \in \mathcal{D}_A$ and*

$$\mathcal{A}f(t, x) = \frac{\partial f}{\partial t} + \sum_{i=1}^n b_i(x) \frac{\partial f}{\partial x_i} + \frac{1}{2} \sum_{i,j} (\sigma \sigma^\top)_{ij} \frac{\partial^2 f}{\partial x_i \partial x_j}.$$

As a stochastic generalization of the Newton-Leibniz axiom, Dynkin's formula gives the expected value of any adequately smooth function of an Itô diffusion at a stopping time:

Theorem 1 (Dynkin's formula [9]). *Let $\{X_t\}$ be a diffusion process in \mathbb{R}^n . Suppose τ is a stopping time with $E[\tau] < \infty$, and $f \in C^{1,2}(\mathbb{R} \times \mathbb{R}^n)$ with compact support. Then*

$$E^{h,x} [f(\tau, X_\tau)] = f(h, x) + E^{h,x} \left[\int_0^\tau \mathcal{A}f(s, X_s) ds \right].$$

In order to specify the behavior of an Itô diffusion across the domain boundary, we introduce the concept of *stopped process*, which is a stochastic process that is forced to have the same value after a prescribed (possibly random) time.

Definition 2 (Stopped process [12]). *Given a stopping time τ and a stochastic process $\{X_t\}$, the stopped process $\{X_t^\tau\}$ is defined by*

$$X^\tau(t, \omega) \triangleq X_{t \wedge \tau}(\omega) = \begin{cases} X(t, \omega) & \text{if } t \leq \tau(\omega), \\ X(\tau(\omega), \omega) & \text{otherwise.} \end{cases}$$

Remark 1. By definition, a stopped process preserves, among others, continuity and the Markov property, and hence the aforementioned results on a stochastic process apply also to a stopped process.

Now consider a stochastic system modeled by an SDE of the form (1) that evolves "within" a not necessarily bounded set $\mathcal{X} \subseteq \mathbb{R}^n$. Since the solution $\{X_t\}$ of Eq. (1) may escape from \mathcal{X} at any time instant $t > 0$, due to the unbounded nature of Gaussian, we define a stopped process $\tilde{X}_t \triangleq X_{t \wedge \tau_{\mathcal{X}}}$ with $\tau_{\mathcal{X}} \triangleq \inf\{t \mid X_t \notin \mathcal{X}\}$. \tilde{X}_t hence represents the process that will stop at the boundary of \mathcal{X} . Denote the infinitesimal generator of the stopped process as $\tilde{\mathcal{A}}$. One plausible property here is that, for all compactly-supported $f \in C^{1,2}(\mathbb{R} \times \mathbb{R}^n)$,

$$\tilde{\mathcal{A}}f(t, x) = \begin{cases} \mathcal{A}f(t, x) & \text{for } x \in \mathcal{X}^\circ, \\ \frac{\partial f}{\partial t}(t, x) & \text{for } x \in \partial\mathcal{X}. \end{cases} \quad (3)$$

The ∞ -Safety Problem. Given an SDE of the form (1), a (not necessarily bounded²) domain set $\mathcal{X} \subseteq \mathbb{R}^n$, an initial set $\mathcal{X}_0 \subset \mathcal{X}$, and an unsafe set $\mathcal{X}_u \subset \mathcal{X}$. We aim to bound the failure probability

$$P\left(\exists t \in [0, \infty): \tilde{X}_t \in \mathcal{X}_u\right),$$

² In practice, if we can specify \mathcal{X} based on prior knowledge when modeling a physical system, then the larger \mathcal{X} we choose, the greater (bound on) failure probability we will obtain.

for any initial state X_0 whose support lies within \mathcal{X}_0 . Accordingly, the T -safety problem, with $T < \infty$, refers to the problem where one aims to bound the failure probability within the finite time horizon $[0, T]$.

Remark 2. Roughly speaking, if we denote by ϕ the proposition “ \tilde{X}_t evolves within \mathcal{X} ” and by ψ the proposition “ \tilde{X}_t evolves into \mathcal{X}_u ”, then the above ∞ -safety problem asks for a bound on the probability that the LTL formula $\phi\mathcal{U}\psi$ holds.

3 Reducing ∞ -Safety to T -Safety

We dedicate this section to the reduction of the ∞ -safety problem to its bounded counterpart. Observe that for any $0 \leq T < \infty$,

$$P(\exists t \geq 0: \tilde{X}_t \in \mathcal{X}_u) \leq P(\exists t \in [0, T]: \tilde{X}_t \in \mathcal{X}_u) + P(\exists t \geq T: \tilde{X}_t \in \mathcal{X}_u).$$

The key idea behind our approach is to first compute an exponentially decreasing bound on the *tail failure probability* over $[T^*, \infty)$ (the computation of $T^* \geq 0$ will be shown later), and then for any constant $\epsilon > 0$, we can identify (out of the exponentially decreasing bound) a time instant $\tilde{T} \geq T^*$ such that $P(\exists t \geq \tilde{T}: \tilde{X}_t \in \mathcal{X}_u) \leq \epsilon$. The overall bound on the failure probability over $[0, \infty)$ can consequently be obtained by solving the truncated \tilde{T} -safety problem.

3.1 Exponentially Decreasing Bound on the Tail Failure Probability

We first state a result that gives conditions when a linear map keeps vector inequality:

Lemma 2 ([4, Chap. 4]). For a matrix $M \in \mathbb{R}^{n \times n}$,

- $\forall x, y \in \mathbb{R}^n: x \leq y \implies Mx \leq My$ iff M is non-negative, i.e., $M_{ij} \geq 0$ for all $1 \leq i, j \leq n$.
- The matrix e^{Mt} is non-negative for all $t \geq 0$ iff M is essentially non-negative, i.e., $M_{ij} \geq 0$ for $i \neq j$.

The existence of an exponentially decreasing bound on the tail failure probability relies on a witness of a supermartingale of the exponential type:

Theorem 2. Suppose there exists an essentially non-negative matrix $\Lambda \in \mathbb{R}^{m \times m}$, together with an m -dimensional polynomial function (termed exponential stochastic barrier certificate) $V(x) = (V_1(x), V_2(x), \dots, V_m(x))^T$, with $V_i: \mathbb{R}^n \rightarrow \mathbb{R}$ for $1 \leq i \leq m$, satisfying^{3,4}

$$V(x) \geq \mathbf{0} \quad \text{for } x \in \mathcal{X}, \tag{4}$$

$$\mathcal{A}V(x) \leq -\Lambda V(x) \quad \text{for } x \in \mathcal{X}, \tag{5}$$

$$\Lambda V(x) \leq \mathbf{0} \quad \text{for } x \in \partial\mathcal{X}. \tag{6}$$

³ Condition (5) is slightly stronger than the corresponding one used in [27,28], yet will lead to an exponentially decreasing bound on the tail failure probability in return.

⁴ Condition (6) is to ensure that when \tilde{X}_t stops at the boundary of \mathcal{X} , we still have $\tilde{\mathcal{A}}V(x) \leq -\Lambda V(x)$ for $x \in \partial\mathcal{X}$. If $\mathcal{X} = \mathbb{R}^n$, however, this condition can be omitted.

Define a function

$$F(t, x) \triangleq e^{At}V(x),$$

then every component of $F(t, \tilde{X}_t)$ is a supermartingale.

Proof. For cases with a bounded domain \mathcal{X} , one can trivially extend the domain of $F(t, x)$ s.t. F is compactly-supported, and thus Dynkin's formula in Theorem 1 applies immediately. For cases where \mathcal{X} is unbounded, we introduce a stopping time

$$\tau_\delta \triangleq \inf \left\{ t \mid F \left(t, \tilde{X}_t \right) \geq \mathfrak{B}(\mathbf{0}, \delta) \right\},$$

and denote by $X_t^{(\delta)} \triangleq (t \wedge \tau_\delta, \tilde{X}_{t \wedge \tau_\delta})$ the corresponding stopped process involving the timeline, and by $\mathcal{A}^{(\delta)}$ the corresponding infinitesimal generator. Then $X_t^{(\delta)}$ evolves within the δ -closed ball $\mathfrak{B}(\mathbf{0}, \delta)$ and hence boils down to the case with a bounded domain. Moreover, by Eq. (3), we have

$$\begin{aligned} \mathcal{A}^{(\delta)} F \left(X_t^{(\delta)} \right) &= \mathcal{A}^{(\delta)} F \left(t \wedge \tau_\delta, \tilde{X}_{t \wedge \tau_\delta} \right) \\ &= \begin{cases} 0 & \text{if } \tau_\delta(\omega) \leq t, \\ \frac{\partial F}{\partial t}(t, X_t) + e^{At} \mathcal{A}V(X_t) \leq 0 & \text{if } \tau_\delta(\omega) > t \wedge \tau_{\mathcal{X}}(\omega) > t, \\ \frac{\partial F}{\partial t}(t, X_t) \leq 0 & \text{if } \tau_\delta(\omega) > t \wedge \tau_{\mathcal{X}}(\omega) \leq t, \end{cases} \end{aligned}$$

where $\tau_{\mathcal{X}}$ represents the time instant when escaping from the state space \mathcal{X} . Note that the second and the third case hold due to the non-negativity of e^{At} (as A is essentially non-negative), which implies that e^{At} preserves vector inequalities (5) and (6). Hence by Dynkin's formula (in a component-wise manner), for fixed $t, h \in [0, \infty)$, we have

$$\begin{aligned} E \left[F \left((t+h) \wedge \tau_\delta, \tilde{X}_{(t+h) \wedge \tau_\delta} \right) \mid \mathcal{F}_h \right] &= E^{X_h^{(\delta)}} \left[F \left(X_{t+h}^{(\delta)} \right) \right] \\ &= F \left(X_h^{(\delta)} \right) + E^{X_h^{(\delta)}} \left[\int_0^t \mathcal{A}^{(\delta)} F \left(X_s^{(\delta)} \right) ds \right] \\ &\leq F \left(X_h^{(\delta)} \right) \\ &= F \left(h \wedge \tau_\delta, \tilde{X}_{h \wedge \tau_\delta} \right). \end{aligned}$$

Since $F(t, x) > \mathbf{0}$, by Fatou's lemma, we have

$$\begin{aligned} E \left[F \left(t+h, \tilde{X}_{t+h} \right) \mid \mathcal{F}_h \right] &= E \left[\liminf_{\delta \rightarrow \infty} F \left((t+h) \wedge \tau_\delta, \tilde{X}_{(t+h) \wedge \tau_\delta} \right) \mid \mathcal{F}_h \right] \\ &\leq \liminf_{\delta \rightarrow \infty} E \left[F \left((t+h) \wedge \tau_\delta, \tilde{X}_{(t+h) \wedge \tau_\delta} \right) \mid \mathcal{F}_h \right] \\ &\leq \liminf_{\delta \rightarrow \infty} F \left(h \wedge \tau_\delta, \tilde{X}_{h \wedge \tau_\delta} \right) \\ &\leq F \left(h, \tilde{X}_h \right). \end{aligned}$$

It follows consequently that every component of $F(t, \tilde{X}_t)$ is a supermartingale. \square

We will show in Sect. 4 that the synthesis of the exponential stochastic barrier certificate $V(x)$ (and thereby the function $F(t, x)$) boils down to solving a pertinent SDP optimization problem.

In order to further establish the relation between the exponential supermartingale $F(t, \tilde{X}_t)$ (and thereby $V(x)$) and the bound on tail failure probability, we recall Doob's maximal inequality for supermartingales, which gives a bound on the probability that a non-negative supermartingale exceeds some given value over a given time interval:

Lemma 3 (Doob's supermartingale inequality [15]). *Let $\{X_t\}_{t>0}$ be a right continuous non-negative supermartingale adapted to a filtration $\{\mathcal{F}_t \mid t > 0\}$. Then for any $\lambda > 0$,*

$$\lambda P\left(\sup_{t \geq 0} X_t \geq \lambda\right) \leq E[X_0].$$

The following theorem claims an intermediate fact that will later reveal the exponentially decreasing bound on the tail failure probability.

Theorem 3. *Suppose the conditions in Theorem 2 are satisfied. Then for any $T \geq 0$ and any positive vector $\gamma \in \mathbb{R}^m$,*

$$P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq \sup_{t \geq T} (e^{-At}\gamma)\right) \leq E[V_i(X_0)]/\gamma_i \quad (7)$$

holds for all $i \in \{1, \dots, m\}$.

Proof. Observe the following chain of (in-)equalities:

$$\begin{aligned} P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq \sup_{t \geq T} (e^{-At}\gamma)\right) &\leq P\left(\exists t \geq T: V(\tilde{X}_t) \geq e^{-At}\gamma\right) \\ &\leq P\left(\exists t \geq T: e^{At}V(\tilde{X}_t) \geq \gamma\right) && \text{[non-negative } e^{At}] \\ &= P\left(\sup_{t \geq T} F(t, \tilde{X}_t) \geq \gamma\right) \\ &\leq P\left(\sup_{t \geq T} F_i(t, \tilde{X}_t) \geq \gamma_i\right) \\ &\leq E\left[F_i(T, \tilde{X}_T)\right]/\gamma_i && \text{[Lemma 3]} \\ &\leq E[V_i(X_0)]/\gamma_i && \text{[Theorem 2]} \end{aligned}$$

which holds for any $i \in \{1, 2, \dots, m\}$. This completes the proof. \square

Now, we are ready to give the exponentially decreasing bound on the tail failure probability derived from Theorem 3. We start by considering the simple case where the barrier certificate $V(x)$ is a scalar function, i.e., with $m = 1$.

Proposition 1. *Suppose there exists a positive constant $\Lambda \in \mathbb{R}$ and a scalar function $V: \mathbb{R}^n \rightarrow \mathbb{R}$ satisfying Theorem 2. Then,*

$$P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq \gamma\right) \leq \frac{E[V(X_0)]}{e^{\Lambda T}\gamma} \quad (8)$$

holds for any $\gamma > 0$ and $T \geq 0$. Moreover, if there exists $l > 0$ such that

$$V(x) \geq l \quad \text{for all } x \in \mathcal{X}_u,$$

then

$$P\left(\exists t \geq T: \tilde{X}_t \in \mathcal{X}_u\right) \leq \frac{E[V(X_0)]}{e^{AT}l} \quad (9)$$

holds for any $T \geq 0$.

Proof. Eq. (8) holds since

$$\begin{aligned} P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq \gamma\right) &= P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq e^{-AT}(e^{AT}\gamma)\right) \\ &\leq P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq \sup_{t \geq T}(e^{-At}(e^{AT}\gamma))\right) \quad [\text{monotonicity on } t] \\ &\leq \frac{E[V(X_0)]}{e^{AT}\gamma}. \quad [\text{Theorem 3}] \end{aligned}$$

For Eq. (9), it is immediately obvious that

$$P\left(\exists t \geq T: \tilde{X}_t \in \mathcal{X}_u\right) \leq P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq l\right) \leq \frac{E[V(X_0)]}{e^{AT}l}.$$

This completes the proof. \square

Now we lift the results to the slightly more involved case with $m > 1$.

Proposition 2. *Suppose there exists an essentially non-negative matrix $\Lambda \in \mathbb{R}^{m \times m}$ and an m -dimensional polynomial function $V: \mathbb{R}^n \rightarrow \mathbb{R}^m$ satisfying Theorem 2. If all of the eigenvalues of Λ have positive real parts, i.e.,*

$$\min_{1 \leq i \leq m} \{\Re(\lambda_i) \mid \lambda_i \text{ is an eigenvalue of } \Lambda\} > 0,$$

then for any positive vector $\gamma \in \mathbb{R}^m$, there exists $T^* = T^*(\gamma, M, \Lambda) \in \mathbb{R}$ such that for any $T \geq T^*$,

$$P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq \gamma\right) \leq \frac{E[V_i(X_0)]}{(e^{MT}\gamma)_i} \quad (10)$$

holds for all $i \in \{1, \dots, m\}$. Here, M is an essentially non-negative matrix s.t. all of the eigenvalues of $\Lambda - M$ have positive real parts⁵. Moreover, if there exists a positive vector $l \in \mathbb{R}^m$ such that

$$V(x) \geq l \quad \text{for all } x \in \mathcal{X}_u,$$

then for any $T \geq T^*$,

$$P\left(\exists t \geq T: \tilde{X}_t \in \mathcal{X}_u\right) \leq \frac{E[V_i(X_0)]}{(e^{MT}l)_i} \quad (11)$$

holds for all $i \in \{1, \dots, m\}$.

⁵ Such matrix M always exists, for instance, $M \hat{=} \Lambda/2$.

Proof. By substituting γ in Eq. (7) with $e^{MT}\gamma$, we have that for all $T \geq 0$,

$$\begin{aligned} \frac{E[V_i(X_0)]}{(e^{MT}\gamma)_i} &\geq P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq \sup_{t \geq T} (e^{-\Lambda t} e^{MT}\gamma)\right) \\ &= P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq \sup_{t \geq T} (e^{-\Lambda(t-T)} e^{-(\Lambda-M)T}\gamma)\right) \end{aligned} \quad (12)$$

holds for any $\gamma \in \mathbb{R}^m$ with $\gamma > \mathbf{0}$. Observe that

$$\begin{aligned} \left| \sup_{t \geq T} (e^{-\Lambda(t-T)} e^{-(\Lambda-M)T}\gamma) \right|_{\infty} &= \left| \sup_{t \geq 0} (e^{-\Lambda t} e^{-(\Lambda-M)T}\gamma) \right|_{\infty} \\ &\leq \left| \sup_{t \geq 0} (e^{-\Lambda t}) \right|_{\infty} \left| e^{-(\Lambda-M)T}\gamma \right|_{\infty}, \end{aligned}$$

where $|\cdot|_{\infty}$ denotes the infinity norm. Moreover, since all of the eigenvalues of $\Lambda - M$ have positive real parts, then by the Lyapunov stability established in the theory of ODEs, we have

$$\lim_{T \rightarrow \infty} e^{-(\Lambda-M)T}\gamma = \mathbf{0}.$$

There hence exists T^* s.t. for all $T \geq T^*$,

$$\sup_{t \geq T} (e^{-\Lambda(t-T)} e^{-(\Lambda-M)T}\gamma) \leq \gamma. \quad (13)$$

By Combining Eq. (13) and Eq. (12), we obtain Eq. (10). For Eq. (11), it follows immediately that

$$P\left(\exists t \geq T: \tilde{X}_t \in \mathcal{X}_u\right) \leq P\left(\sup_{t \geq T} V(\tilde{X}_t) \geq l\right) \leq \frac{E[V_i(X_0)]}{(e^{MT}l)_i}.$$

This completes the proof. \square

Remark 3. Proposition 2 argues the existence of T^* that suffices to “split off” the tail failure probability. From a computational perspective, this is algorithmically tractable as the matrix exponential involved in Eq. (13) is symbolically computable (cf., e.g., [23]).

The following theorem states the main result of this section, that is, for any given constant ϵ , there exists $\tilde{T} \geq 0$ such that the truncated \tilde{T} -tail failure probability is bounded by ϵ :

Theorem 4. *Suppose the conditions in Proposition 1 and 2 are satisfied. If there exists $\alpha > 0$, s.t. $\forall x \in \mathcal{X}_0: V_i(x) \leq \alpha$ holds for some $i \in \{1, \dots, m\}$. Then for any $\epsilon > 0$, there exists $\tilde{T} \geq 0$ such that*

$$P\left(\exists t \geq \tilde{T}: \tilde{X}_t \in \mathcal{X}_u\right) \leq \epsilon.$$

Proof. Observe that for Eq. (11) in Proposition 2, the assumption $\forall x \in \mathcal{X}_0: V_i(x) \leq \alpha$ guarantees an upper bound on the numerator $E[V_i(X_0)]$, while the essential non-negativity of M (with all its eigenvalues having positive real parts) ensures that the denominator $(e^{MT}l)_i \rightarrow +\infty$ as $T \rightarrow \infty$. An analogous argument applies to Eq. (9) in Proposition 1. The claim in this theorem then follows immediately. \square

3.2 Bounding the Failure Probability over $[0, T]$

The reduced T -safety problem can be solved by existing methods tailored for bounded verification of SDEs, e.g., [35,32]. In what follows, we propose an alternative method leveraging time-dependent polynomial stochastic barrier certificates. Our method requires constraints (on the barrier certificates) of simpler form compared to [35]; meanwhile, it yields strictly more expressive form of barrier certificates, against the approach on unbounded verification as in [27,28], thus leading to theoretically non-looser (usually tighter) failure bound. A detailed argument will be given at the end of this section.

The following theorem states a sufficient condition, i.e., a collection of constraints on the time-dependent polynomial stochastic barrier certificates $H(t, x)$, under which the failure probability of a stochastic system over a finite time horizon can be explicitly bounded from above.

Theorem 5. *Suppose there exists a constant $\eta > 0$ and a polynomial function (termed time-dependent stochastic barrier certificate) $H(t, x): \mathbb{R} \times \mathbb{R}^n \rightarrow \mathbb{R}$, satisfying⁶*

$$H(t, x) \geq 0 \quad \text{for } (t, x) \in [0, T] \times \mathcal{X}, \quad (14)$$

$$\mathcal{A}H(t, x) \leq 0 \quad \text{for } (t, x) \in [0, T] \times (\mathcal{X} \setminus \mathcal{X}_u), \quad (15)$$

$$\frac{\partial H}{\partial t} \leq 0 \quad \text{for } (t, x) \in [0, T] \times \partial\mathcal{X}, \quad (16)$$

$$H(t, x) \geq \eta \quad \text{for } (t, x) \in [0, T] \times \mathcal{X}_u. \quad (17)$$

Then,

$$P\left(\exists t \in [0, T]: \tilde{X}_t \in \mathcal{X}_u\right) \leq \frac{E[H(0, X_0)]}{\eta}. \quad (18)$$

Proof. Assume in the following that the system evolves within a bounded domain \mathcal{X} ⁷. Define a stopping time

$$\tau_u \triangleq \inf \left\{ t \mid \tilde{X}_t \notin \mathcal{X} \setminus \mathcal{X}_u \right\},$$

and denote by $X_t^{(u)} \triangleq (t \wedge \tau_u \wedge T, \tilde{X}_{t \wedge \tau_u \wedge T})$ the corresponding stopped process, and by $\mathcal{A}^{(u)}$ the corresponding infinitesimal generator. By Eq. (3), we have

$$\begin{aligned} \mathcal{A}^{(u)} H \left(X_t^{(u)} \right) &= \mathcal{A}^{(u)} H \left(t \wedge \tau_u \wedge T, \tilde{X}_{t \wedge \tau_u \wedge T} \right) \\ &= \begin{cases} 0 & \text{if } t \geq T \vee t \geq \tau_u(\omega), \\ \mathcal{A}H(t, X_t) \leq 0 & \text{if } t < \min\{T, \tau_u(\omega), \tau_{\mathcal{X}}(\omega)\}, \\ \frac{\partial H}{\partial t}(t, X_t) \leq 0 & \text{if } t < \min\{T, \tau_u(\omega)\} \wedge t \geq \tau_{\mathcal{X}}(\omega). \end{cases} \end{aligned}$$

⁶ Condition (16) is to ensure that when \tilde{X}_t stops at the boundary of \mathcal{X} , we still have $\tilde{A}H(t, x) \leq 0$ for $x \in \partial\mathcal{X}$. If $\mathcal{X} = \mathbb{R}^n$, however, this condition can be dropped.

⁷ For cases with an unbounded \mathcal{X} , the same proof technique of introducing a δ -closed ball as in the proof of Theorem 2 applies.

By Dynkin's formula, for fixed $t, h \in [0, T]$, we have

$$\begin{aligned} E \left[H \left(X_{t+h}^{(u)} \right) \mid \mathcal{F}_t \right] &= E^{X_t^{(u)}} \left[H \left(X_{t+h}^{(u)} \right) \right] \\ &= E \left[H \left(X_h^{(u)} \right) \right] + E^{X_t^{(u)}} \left[\int_0^t \mathcal{A}^{(u)} H \left(X_s^{(u)} \right) ds \right] \\ &\leq E \left[H \left(X_h^{(u)} \right) \right]. \end{aligned}$$

Thus $H(X_t^{(u)})$ is a non-negative supermartingale. Then by Doob's maximal inequality in Lemma 3, we have

$$\begin{aligned} P \left(\exists t \in [0, T]: \tilde{X}_t \in \mathcal{X}_u \right) &= P \left(\exists t \geq 0: \tilde{X}_{t \wedge \tau_u \wedge T} \in \mathcal{X}_u \right) \\ &\leq P \left(\exists t \geq 0: H \left(X_t^{(u)} \right) \geq \eta \right) \\ &\leq \frac{E[H(0, X_0)]}{\eta}. \end{aligned}$$

This completes the proof. \square

The following fact is then immediately obvious:

Corollary 1. *Suppose the conditions in Theorem 5 hold, and there exists $\beta > 0$, s.t. $H(0, x) \leq \beta$ for $x \in \mathcal{X}_0$. Then,*

$$P \left(\exists t \in [0, T]: \tilde{X}_t \in \mathcal{X}_u \right) \leq \frac{\beta}{\eta}.$$

Proof. This is a direct consequence of Theorem 5. \square

Remarks on Potentially Tighter Bound. There exists already in the literature a barrier certificate-based method proposed in [27,28] that can deal with the ∞ -safety problem. It is worth highlighting, however, that our bound on the overall failure probability derived from Proposition 1, 2 and Theorem 5 (with appropriate \tilde{T} chosen) is at least as tight as (and usually tighter than, as can be seen later in the experiments) that in [27,28]. The reasons are twofold: (1) the reduction to a finite-time horizon \tilde{T} -safety problem substantially "trims off" verification efforts pertaining to $t > \tilde{T}$; (2) our method for the reduced \tilde{T} -safety problem admits time-dependent barrier certificates, which are strictly more expressive than those time-independent ones exploited in [27,28], in the sense that any feasible solution thereof shall also be a feasible solution satisfying Theorem 5.

Remark 4. Roughly speaking, by setting the diffusion coefficients σ in SDEs to zero, our method applies trivially to ODE dynamics with either a known or an unknown probability distribution over the initial set of states. For the former, we can even obtain a tighter bound on the failure probability, since in this case we do not need to compute a bound on the barrier certificate over all possible initial distributions.

4 Synthesizing Stochastic Barrier Certificates Using SDP

In this section, we encode the synthesis of the aforementioned exponential and time-dependent stochastic barrier certificates into semidefinite programming [38] optimizations, and thus a solution thereof yields an upper bound on the failure probability over the infinite-time horizon. Specifically, an SDP problem is formulated, for each of the two barrier certificates, to encode the constraints for “being an exponential/time-dependent stochastic barrier certificate”, while in the meantime optimizing the tightness of the failure probability bound.

It is worth noting that SDP is a generalization of the standard linear programming in which the element-wise non-negativity constraints are replaced by a generalized inequality w.r.t. the cone of positive semidefinite matrices. The generalization preserves *convexity*, leading to the fact that SDP admits polynomial-time algorithms, say the well-known *interior-point methods*, that can efficiently solve the synthesis problem, albeit numerically. We remark that the numerical computation employed in off-the-shelf SDP solvers and the use of interior-point algorithms may potentially lead to erroneous results and thereby unsoundness in the verification/synthesis results. There have been numerous attempts to validate the results from the solver through a-posteriori numerical verification of the solution. For more details, we refer the readers to [30] and the references therein.

Exponential Stochastic Barrier Certificate $V(x)$. To encode the synthesis problem into an SDP optimization, we first fix the dimension m together with Λ satisfying Proposition 1 or 2 (depending on m), and then assume a polynomial template $V^a(x)$ of certain degree k with unknown parameters a , as the barrier certificate to be discovered. It then suffices to solve the following SDP problem⁸:

$$\underset{a, \alpha}{\text{minimize}} \quad \alpha \tag{19}$$

$$\text{subject to} \quad V^a(x) \geq \mathbf{0} \quad \text{for } x \in \mathcal{X} \tag{20}$$

$$\mathcal{A}V^a(x) \leq -\Lambda V^a(x) \quad \text{for } x \in \mathcal{X} \tag{21}$$

$$\Lambda V^a(x) \leq \mathbf{0} \quad \text{for } x \in \partial\mathcal{X} \tag{22}$$

$$V^a(x) \geq \mathbf{1} \quad \text{for } x \in \mathcal{X}_u \tag{23}$$

$$V^a(x) \leq \alpha \mathbf{1} \quad \text{for } x \in \mathcal{X}_0 \tag{24}$$

Here, the constraints (20)–(22) encode the definition of an exponential stochastic barrier certificate (cf. Theorem 2), while constraint (23) (resp., (24)) corresponds to the lower (resp., upper) bound of $V(x)$ as in Proposition 1 and 2 (resp., Theorem 4)⁹. Hence, minimizing the upper bound α of (each component of) $V^a(x)$ gives a tight exponentially decreasing bound on the tail failure probability, as claimed in Proposition 1 and 2.

⁸ SDP problems in this paper refer to those that can be readily translated into the standard form of SDP, through, e.g., Stengle’s Positivstellensatz [36] and sum-of-squares decomposition [26].

⁹ The lower bound l of $V(x)$ in Proposition 1 and 2 is normalized to a vector with all its components no less than 1, based on the observation that, for any $c > 0$, $V^a(x)$ is a feasible solution implies $cV^a(x)$ is also a feasible solution.

Remark 5. If Λ is chosen as a non-negative matrix, the combination of condition (20) and (22) will force $V^a(x) = \mathbf{0}$ for $x \in \partial\mathcal{X}$, whereof the strict equality may be violated due to numerical computations in SDP. In practice, however, this issue can be well addressed by looking for a barrier certificate of the form $g(x)V(x)$, where $g(x)$ satisfies $\partial\mathcal{X} \subseteq \{x \mid g(x) = 0\}$, namely, an overapproximation of the boundary of \mathcal{X} .

Remark 6. The choice of m is arbitrary, while the choices of Λ and k can be heuristic: If Λ_1 admits no feasible solution, neither will $\Lambda_2 \geq \Lambda_1$ (point-wise, with all the rest parameters fixed); similarly, if k_1 admits no feasible solution, neither will $k_2 \leq k_1$ (with all the rest parameters fixed). Therefore, one may decrease Λ (say, by a half) or increase k (say, by one) whenever a valid barrier certificate was not found.

Time-Dependent Stochastic Barrier Certificate $H(t, x)$. Given the results established in Sect. 3, the corresponding synthesis problem can be analogously encoded as the following SDP problem:

$$\underset{b, \beta}{\text{minimize}} \quad \beta \quad (25)$$

$$\text{subject to} \quad H^b(t, x) \geq 0 \quad \text{for } (t, x) \in [0, T] \times \mathcal{X} \quad (26)$$

$$\mathcal{A}H^b(t, x) \leq 0 \quad \text{for } (t, x) \in [0, T] \times (\mathcal{X} \setminus \mathcal{X}_u) \quad (27)$$

$$\frac{\partial H^b}{\partial t} \leq 0 \quad \text{for } (t, x) \in [0, T] \times \partial\mathcal{X} \quad (28)$$

$$H^b(t, x) \geq 1 \quad \text{for } (t, x) \in [0, T] \times \mathcal{X}_u \quad (29)$$

$$H^b(0, x) \leq \beta \quad \text{for } x \in \mathcal{X}_0 \quad (30)$$

Similarly, the constraints (26)–(29) encode the definition of a time-dependent stochastic barrier certificate (cf. Theorem 5), while constraint (30) corresponds to the upper bound of $H(t, x)$ as in Corollary 1 (with η being normalized to 1, as in constraint (29)). Consequently, minimizing the upper bound β of $H^b(t, x)$ produces a tight bound on the failure probability over the reduced finite-time horizon, as stated in Corollary 1.

Remark 7. The state-of-the-art interior-point methods solve an SDP problem up to an error ε in time that is polynomial in the program description size (number of variables) and $\log(1/\varepsilon)$. The former is exponential in the degree of V^a and H^b , as it corresponds to the number of monomials in the template polynomials.

5 Implementation and Experimental Results

To further demonstrate the practical performance of our approach, we have carried out a prototypical implementation in MATLAB R2019b, with the toolbox YALMIP [21] and MOSEK [2] equipped for formulating and solving the underlying SDP problems. Given an ∞ -safety problem as input, our implementation works toward an upper bound on the failure probability over the infinite time horizon, leveraging the reduction to a T -safety problem based on a computed exponentially decreasing bound on the tail failure probability. A collection of benchmark examples from the literature has been evaluated on a 1.8GHz Intel Core-i7 processor with 8GB RAM running 64-bit Windows 10.

Each of the examples has been successfully tackled within 30 seconds. In what follows, we demonstrate the applicability of our techniques to SDEs featuring different dimensionalities and nonlinear dynamics, and show particularly that our approach usually produces tighter bounds compared to existing methods.

Example 1 (Population growth [25]). Consider the stochastic system

$$dX_t = b(X_t) dt + \sigma(X_t) dW_t,$$

which is a stochastic model of population dynamics subject to random fluctuations that, possibly, can be attributed to extraneous or chance factors such as the weather, location, and the general environment. Suppose that the state space is restricted within $\mathcal{X} = \{x \mid x \geq 0\}$ with $b(X_t) = -X_t$ and $\sigma(X_t) = \sqrt{2}/2X_t$. We instantiate the ∞ -safety problem as $\mathcal{X}_0 = \{x \mid x = 1\}$ and $\mathcal{X}_u = \{x \mid x \geq 2\}$, namely, we expect that the population does not diverge beyond 2.

Let $\Lambda = 1$ (with $m = 1$) and set the polynomial template degree of the exponential stochastic barrier certificate $V^a(x)$ to 4, the SDP solver gives

$$V^a(x) = 0.000001474596322 - 0.000044643990040x \\ + 0.125023372121222x^2 + 0.000000001430428x^3,$$

which satisfies

$$V^a(x) \geq 1 \text{ for } x \in \mathcal{X}_u \quad \text{and} \quad V^a(x) \leq 0.12498 \text{ for } x \in \mathcal{X}_0.$$

Thus by Proposition 1, we obtain the exponentially decreasing bound

$$P\left(\exists t \geq T: \tilde{X}_t \in \mathcal{X}_u\right) \leq \frac{0.12498}{e^T} \quad \text{for all } T > 0.$$

The user then may choose any $T > 0$ and solve the reduced T -safety problem. As depicted in the left of Fig. 1, different choices lead to different bounds on the failure probability. Nevertheless, one may surely select an appropriate T that yields a way tighter overall bound on the failure probability than that produced by the method in [27,28].

Example 2 (Harmonic oscillator [13]). Consider a two-dimensional harmonic oscillator with noisy damping:

$$dX_t = \begin{pmatrix} 0 & \omega \\ -\omega & -k \end{pmatrix} X_t dt + \begin{pmatrix} 0 & 0 \\ 0 & -\sigma \end{pmatrix} X_t dW_t,$$

with constants $\omega = 1, k = 7$ and $\sigma = 2$. We instantiate the ∞ -safety problem as $\mathcal{X} = \mathbb{R}^n$, $\mathcal{X}_0 = \{(x_1, x_2) \mid -1.2 \leq x_1 \leq 0.8, -0.6 \leq x_2 \leq 0.4\}$ and $\mathcal{X}_u = \{(x_1, x_2) \mid |x_1| \geq 2\}$.

Let $\Lambda = \begin{pmatrix} 0.45 & 0.1 \\ 0.1 & 0.45 \end{pmatrix}$ and set the polynomial template degree of the exponential stochastic barrier certificate $V^a(x)$ to 4, the SDP solver produces a two-dimensional $V^a(x)$ (abbreviated for clear presentation) satisfying

$$V^a(x) \leq \begin{pmatrix} 0.19946 \\ 0.19946 \end{pmatrix} \text{ for } x \in \mathcal{X}_0 \quad \text{and} \quad V^a(x) \geq l = \begin{pmatrix} 1.000237 \\ 1.000236 \end{pmatrix} \text{ for } x \in \mathcal{X}_u.$$

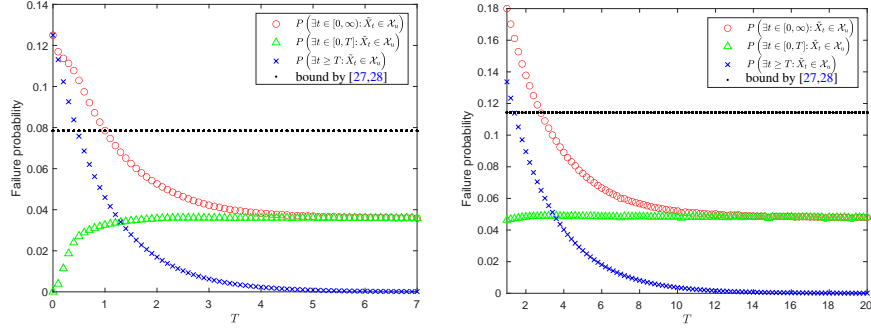


Fig. 1: Different choices of T lead to different bounds on the failure probability (with the time-dependent stochastic barrier certificates of degree 4). Note that ‘ \circ ’ = ‘ \times ’ + ‘ \triangle ’ and ‘ \bullet ’ depicts the overall bound on the failure probability produced by the method in [27,28].

According to the proof of Proposition 2, we set $M = \begin{pmatrix} 0.3 & 0.1 \\ 0.1 & 0.3 \end{pmatrix}$ and aim to find $T^* \geq 0$ such that for all $T \geq T^*$,

$$\sup_{t \geq 0} \left(e^{-\Lambda t} e^{-(\Lambda - M)T} \begin{pmatrix} 1.000237 \\ 1.000236 \end{pmatrix} \right) \leq \begin{pmatrix} 1.000237 \\ 1.000236 \end{pmatrix}. \quad (31)$$

Symbolic computation on the matrix exponential gives

$$\begin{aligned} \sup_{t \geq 0} \left(e^{-\Lambda t} e^{-(\Lambda - M)T} \begin{pmatrix} 1.000237 \\ 1.000236 \end{pmatrix} \right) &= \sup_{t \geq 0} \left(\begin{matrix} e^{-0.15T} (1.0002365e^{-0.55t} + 0.0000005e^{-0.35t}) \\ e^{-0.15T} (1.0002365e^{-0.55t} - 0.0000005e^{-0.35t}) \end{matrix} \right) \\ &\leq \begin{pmatrix} 1.0002365e^{-0.15T} \\ 1.0002365e^{-0.15T} \end{pmatrix}. \end{aligned}$$

Therefore, $T^* = 1$ satisfies condition (31). Further by Corollary 2, for any $T \geq T^* = 1$, we have

$$P(\exists t \geq T: \tilde{X}_t \in \mathcal{X}_u) \leq \frac{E[V_1(X_0)]}{(e^{MT}l)_1} \leq \frac{0.19946}{0.0000005e^{0.2T} + 1.00024e^{0.4T}}.$$

Analogously, a comparison with existing methods concerning the tightness of the synthesized failure probability bound (under different choices of T) is shown in the right of Fig. 1.

Example 3 (Nonlinear drift [27]). We consider in this example a stochastic system involving nonlinear dynamics in its drift coefficient:

$$\begin{aligned} dx_1(t) &= x_2(t) dt \\ dx_2(t) &= -x_1(t) - x_2(t) - 0.5x_1^3(t) dt + 0.1 dW_t. \end{aligned}$$

As in [27], let $\mathcal{X} = \{(x_1, x_2) \mid |x_1| \leq 3, |x_2| \leq 3, x_1^2 + x_2^2 \geq 0.5^2\}$, $\mathcal{X}_0 = \{(x_1, x_2) \mid (x_1 + 2)^2 + x_2^2 \leq 0.1^2\}$ and $\mathcal{X}_u = \{(x_1, x_2) \in \mathcal{X} \mid x_2 \geq 2.25\}$. With $\Lambda = 1.5$ ($m = 1$), we obtain an exponential stochastic barrier certificate $V^a(x)$ of degree 8 satisfying

$$V^a(x) \leq 4.00014 \text{ for } x \in \mathcal{X}_0 \quad \text{and} \quad V^a(x) \geq 1.05248 \text{ for } x \in \mathcal{X}_u.$$

Thus by Corollary 1, we have for any $T \geq 0$,

$$P\left(\exists t \geq T: \tilde{X}_t \in \mathcal{X}_u\right) \leq \frac{3.80070}{e^{1.5T}}.$$

Setting, for instance, $T = 6$, we have

$$P\left(\exists t \geq 0: \tilde{X}_t \in \mathcal{X}_u\right) \leq P\left(\exists t \in [0, 6]: \tilde{X}_t \in \mathcal{X}_u\right) + \frac{3.80070}{e^9}.$$

For the reduced T -safety problem with $T = 6$, a time-dependent stochastic barrier certificate of degree 8 is synthesized, thereby yielding $P\left(\exists t \in [0, 6]: \tilde{X}_t \in \mathcal{X}_u\right) \leq 0.196124$, thus together we get

$$P\left(\exists t \geq 0: \tilde{X}_t \in \mathcal{X}_u\right) \leq 0.196593,$$

which is tighter than 0.265388 produced (on the same machine) by the method in [27] under the same template degree.

6 Conclusion

We proposed a constructive method, based on the synthesis of stochastic barrier certificates, for computing an exponentially decreasing upper bound, if existent, on the tail probability that an SDE system violates a given safety specification. We showed that such an upper bound facilitates a reduction of the verification problem over an unbounded temporal horizon to that over a bounded one. Preliminary experimental results on a set of interesting examples from the literature demonstrated the effectiveness of the reduction and that our method often produces tighter bounds on the failure probability.

For future work, we plan to investigate a possible convergence result in the sense that the derived failure probability bound may converge to the exact one as increasing the degree of the barrier certificates. Extending our technique to tackle SDEs with control inputs will also be of interest. Moreover, checking whether a given parametric (polynomial) formula keeps probabilistic invariance plays a central role in the verification of SDEs. Several kinds of sufficient conditions on probabilistic barrier certificates were proposed, including the ones given in this paper. It consequently deserves to investigate a necessary and sufficient condition for checking the probabilistic invariance of a given template, like for ODEs in [19]. Apart from that, we are interested in carrying our results to the verification of probabilistic programs without conditioning, which can be viewed as discrete-time stochastic dynamics.

References

1. A. Abate, M. Prandini, J. Lygeros, and S. Sastry. Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems. *Automatica*, 44(11):2724–2734, 2008.
2. E. D. Andersen, C. Roos, and T. Terlaky. On implementing a primal-dual interior-point method for conic quadratic optimization. *Math. Program.*, 95(2):249–277, 2003.

3. C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
4. E. F. Beckenbach and R. E. Bellman. *Inequalities, Ergeb. Math. Grenzgeb.*, volume 30. Springer-Verlag, Berlin–Göttingen–Heidelberg, xii, 1961.
5. F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
6. H. Blom, G. Bakker, and J. Krystul. Probabilistic reachability analysis for large scale stochastic hybrid systems. In *CDC'07*, pages 3182–3189, 2007.
7. M. L. Bujorianu. Extended stochastic hybrid systems and their reachability problem. In *HSCC'04*, volume 2993 of *LNCS*, pages 234–249, 2004.
8. J. Deshmukh and S. Sankaranarayanan. Formal techniques for verification and testing of cyber-physical systems. In *Design Automation of Cyber-Physical Systems (Arquimedes Canedo and Mohammad Al Faruque)*, pages 69–105. Springer-Verlag, 2019.
9. E. B. Dynkin. *Markov Processes, 2 vols.* Springer Verlag, New York, Berlin, Heidelberg, 1965.
10. A. Einstein. On the theory of Brownian motion. *Annalen der Physik*, pages 371–381, 1906.
11. S. Feng, M. Chen, N. Zhan, M. Fränzle, and B. Xue. Taming delays in dynamical systems - unbounded verification of delay differential equations. In *CAV'19*, volume 11561 of *LNCS*, pages 650–669. Springer, 2019.
12. R. G. Gallager. *Stochastic Processes: Theory for Applications*. Cambridge University Press, 2013.
13. S. Hafstein, S. Gudmundsson, P. Giesl, and E. Scalas. Lyapunov function computation for autonomous linear stochastic differential equations using sum-of-squares programming. *Discrete and Continuous Dynamical Systems-Series B*, 23(2):939–956, 2018.
14. S. Hoogendoorn and P. Bovy. Pedestrian route-choice and activity scheduling theory and models. *Transportation Research Part B: Methodological*, 38(2):169–190, 2004.
15. I. Karatzas and S. Shreve. *Brownian Motion and Stochastic Calculus*. Graduate Texts in Mathematics. Springer New York, 2014.
16. X. D. Koutsoukos and D. Riley. Computational methods for verification of stochastic hybrid systems. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 38(2):385–396, 2008.
17. H. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer, 2001.
18. A. Lecchini-Visintini, J. Lygeros, and J. Maciejowski. Stochastic optimization on continuous domains with finite-time guarantees by Markov chain Monte Carlo methods. *IEEE Trans. Automat. Contr.*, 55(12):2858–2863, 2010.
19. J. Liu, N. Zhan, and H. Zhao. Computing semi-algebraic invariants for polynomial dynamical systems. In *EMSOFT'11*, pages 97–106. ACM, 2011.
20. K. Liu, M. Li, and Z. She. Reachability estimation of stochastic dynamical systems by semi-definite programming. In *CDC'19*, pages 7727–7732. IEEE, 2019.
21. J. Löfberg. YALMIP: A toolbox for modeling and optimization in MATLAB. In *CACSD'04*, pages 284–289, 2004.
22. I. Mitchell and J. Templeton. A toolbox of Hamilton-Jacobi solvers for analysis of non-deterministic continuous and hybrid systems. In *HSCC'05*, volume 3414 of *LNCS*, pages 573–589, 2005.
23. C. Moler and C. Van Loan. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM review*, 45(1):3–49, 2003.
24. B. Øksendal. *Stochastic Differential Equations: An introduction with applications*. Springer Science & Business Media, 2013.
25. M. Panik. *Stochastic Differential Equations: An Introduction with Applications in Population Dynamics Modeling*. Wiley, 2017.

26. P. A. Parillo. Semidefinite programming relaxation for semialgebraic problems. *Mathematical Programming Ser. B*, 96(2):293–320, 2003.
27. S. Prajna, A. Jadbabaie, and G. J. Pappas. Stochastic safety verification using barrier certificates. In *CDC'04*, volume 1, pages 929–934. IEEE, 2004.
28. S. Prajna, A. Jadbabaie, and G. J. Pappas. A framework for worst-case and stochastic safety verification using barrier certificates. *IEEE Trans. Automat. Contr.*, 52(8):1415–1428, 2007.
29. R. Rajkumar, I. Lee, L. Sha, and J. Stankovic. Cyber-physical systems: The next computing revolution. In *DAC'10*, pages 731–736. ACM, 2010.
30. P. Roux, Y. Voronin, and S. Sankaranarayanan. Validating numerical semidefinite programming solvers for polynomial invariants. *Formal Methods Syst. Des.*, 53(2):286–312, 2018.
31. S. Sankaranarayanan, A. Chakarov, and S. Gulwani. Static analysis for probabilistic programs: inferring whole program properties from finitely many paths. In *PLDI'13*, pages 447–458, 2013.
32. C. Santoyo, M. Dutreix, and S. Coogan. Verification and control for finite-time safety of stochastic systems via barrier functions. In *CCTA'19*, pages 712–717. IEEE, 2019.
33. C. Sloth and R. Wisniewski. Safety analysis of stochastic dynamical systems. In *ADHS'15*, pages 62–67, 2015.
34. A. Sogokon, K. Ghorbal, Y. K. Tan, and A. Platzer. Vector barrier certificates and comparison systems. In *FM'18*, volume 10951 of *Lecture Notes in Computer Science*, pages 418–437. Springer, 2018.
35. J. Steinhardt and R. Tedrake. Finite-time regional verification of stochastic non-linear systems. *The International Journal of Robotics Research*, 31(7):901–923, 2012.
36. G. Stengle. A nullstellensatz and a positivstellensatz in semialgebraic geometry. *Math. Ann.*, 207(2):87–97, 1974.
37. X. Wang, H. Chiang, J. Wang, H. Liu, and T. Wang. Long-term stability analysis of power systems with wind power based on stochastic differential equations: Model development and foundations. *IEEE Transactions on Sustainable Energy*, 6(4):1534–1542, Oct 2015.
38. H. Wolkowicz, R. Saigal, and L. Vandenberghe. *Handbook of Semidefinite Programming: Theory, Algorithms, and Applications*. International Series in Operations Research & Management Science. Springer US, 2012.
39. H. L. S. Younes and R. G. Simmons. Probabilistic verification of discrete event systems using acceptance sampling. In *CAV'02*, volume 2404 of *LNCS*, pages 223–235. Springer-Verlag, 2002.