Tolerant Barrier Certificates for Stochastic Systems

Shenghua Feng†, Han Su, Hao Wu, Jie An†, Mingshuai Chen†, and Naijun Zhan†

***Abstract*—We study stochastic verification under tolerant specifications, where unsafe behavior is allowed as long as the total unsafe exposure remains below a prescribed budget with high probability. In discrete time, unsafe exposure is the number of visits to an unsafe set; in continuous time, it is the occupation time spent there. We introduce tolerant barrier certificates, whose key idea is to force the barrier to decrease more aggressively whenever the state is unsafe, so that the barrier accounts for the unsafe exposure accumulated along the trajectory. This construction yields supermartingale-based upper bounds on the probability that the unsafe-exposure budget is exceeded, for both tolerant safety and tolerant reach-avoid properties. We investigate both discrete-time stochastic systems and stochastic differential dynamics, and demonstrate the effectiveness of our approach on a collection of benchmarks.**

***Index Terms*—Stochastic systems, Tolerant safety, Reach-avoid, Barrier certificates**

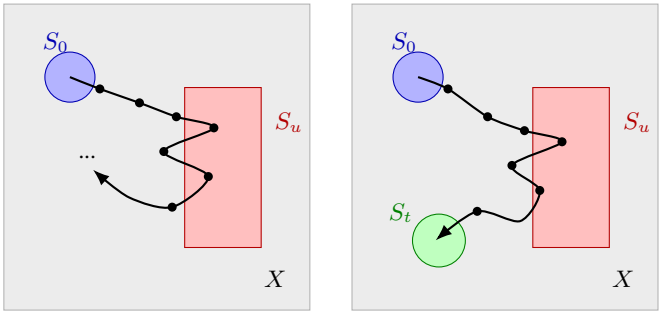

Fig. 1: Conceptual illustration of tolerant properties. Left: tolerant safety allows occasional visits to the unsafe set S_u, provided the total unsafe exposure stays within a budget. Right: tolerant reach-avoid additionally requires the trajectory to reach the target set S_t while keeping the unsafe exposure accumulated before arrival within the prescribed budget.

I. Introduction

STOCHASTIC dynamical systems model the evolution of physical and computational processes under random disturbances. They provide compact mathematical models for real-world dynamics affected by randomness. Such models are particularly relevant to embedded and cyber-physical systems, where sensing noise, actuation errors, and environmental disturbances can induce short unsafe excursions without necessarily indicating mission failure. Prominent instances of stochastic systems include Markov chains [1], Markov decision processes [2], probabilistic graphical models (e.g., Bayesian networks and Markov random fields) [3], random ordinary differential equations [4], stochastic differential equations [5], and stochastic hybrid automata [6]. These system models have witnessed a wide spectrum of applications ranging from control theory [7], physics [8], to finance [9], biology [10], and neuroscience [11].

Manuscript received August 19, 2026; revised August 26, 2026. This work was supported by the National Key R&D Program of China (No. 2022YFA1005101 and 2022YFA1005102), the ZJNSF Major Program (No. LD24F020013), the NSFC (No. 62572427, 62192732, W2511064, and 62502475), the CAS Project for Young Scientists in Basic Research (No. YSBR-123), the ISCAS Basic Research Grant (No. ISCAS-JCZD-202406), and the Fundamental Research Funds for the Central Universities of China (No. 226-2024-00140). Han Su is supported by the ASPIRE grant No. JPMJAP2301, JST.

Shenghua Feng and Jie An are with the National Key Laboratory of Space Integrated Information System, Institute of Software, Chinese Academy of Sciences, Beijing, China (email: {fengshenghua,anjie}@iscas.ac.cn). Han Su is with the National Institute of Informatics (NII), Tokyo, Japan (email: suhan@nii.ac.jp). Hao Wu is with the KLSS, Institute of Software, Chinese Academy of Sciences (email: wuhao@ios.ac.cn). Mingshuai Chen is with Zhejiang University, Hangzhou, China (email: m.chen@zju.edu.cn). Naijun Zhan is with the MoE Key Laboratory of High Confidence Software Technologies, School of Computer Science, Peking University, and the Zhongguancun Laboratory, Beijing, China (email: njzhan@pku.edu.cn).

† Corresponding authors.

Formal verification of stochastic systems has traditionally focused on *strict* specifications, such as safety, reachability, and reach-avoid. Under these specifications, a system must avoid unsafe states altogether, or must reach a target set without ever violating a state constraint. While mathematically clean, such formulations are often too brittle for practical stochastic systems. Random disturbances, modeling errors, sensor noise, and rare environmental perturbations may cause occasional unsafe excursions even when the overall behavior is acceptable. In many applications, the relevant question is therefore not whether unsafe states are visited *at all*, but whether the cumulative unsafe exposure remains sufficiently limited with high probability. For example, consider an autonomous mobile robot navigating through a cluttered workspace under stochastic disturbances. The robot must reach a docking region while avoiding a low-friction area. A classical reach-avoid property would declare the execution unsuccessful even if the robot briefly touches the low-friction area once and still safely reaches the dock. In contrast, a *tolerant* property requires that the robot reach the target while entering the unsafe region at most k times, or while spending at most a prescribed amount of time there.

This observation motivates *tolerant* properties. Instead of requiring zero unsafe exposure, tolerant safety allows a bounded amount of exposure to an unsafe set, and tolerant reach-avoid requires reaching a target while keeping the unsafe exposure below a prescribed budget. In discrete time, this budget can be expressed as the number of visits to the unsafe set; in continuous time, a more natural quantity is the occupation time spent in the unsafe region. Such formulations capture the fact that occasional or short-lived violations may be acceptable, whereas persistent or repeated violations are not.

The central challenge is to certify these tolerant properties for general stochastic systems in a scalable way. For finite-state models, one may in principle compute the exact distribution of unsafe exposure using dynamic programming or probabilistic model checking. However, these approaches do not scale well to nonlinear or continuous-state stochastic systems. Barrier certificates [12] offer an attractive alternative: they provide function-based certificates that avoid explicit state-space discretization and have become a standard tool for safety and reachability analysis of stochastic systems. Existing barrier-based approaches, however, are not designed specifically for tolerant properties. The key difficulty is that the desired guarantee concerns an upper-tail probability for the cumulative unsafe exposure. Rather than only bounding the expected exposure, we must certify that the probability of exceeding a prescribed exposure budget is small. We therefore develop a certificate-based verification method that works directly on the original stochastic system and produces probabilistic upper-tail bounds on cumulative unsafe exposure.

Our starting point is a barrier condition that distinguishes between unsafe and safe regions. Whenever the system lies in the unsafe set, the barrier must decrease sufficiently to account for the additional unsafe exposure incurred at that step or over that time interval. This leads to a family of tolerant barriers that can certify probabilistic bounds on cumulative unsafe exposure. In discrete time, the framework applies to the total number of unsafe visits and to reach-avoid properties where the count is accumulated only until the target is reached. In continuous time, the same idea extends naturally by replacing visit counts with unsafe occupation time and using generator-based conditions for stochastic differential equations.

Intuitively, we want a certificate that "charges" the system every time it spends exposure budget. In discrete time, each visit to the unsafe set should cost one unit; in continuous time, each infinitesimal unsafe time interval should also incur unit cost. The barrier is therefore required to drop more strongly inside the unsafe set than outside it. After adding the accumulated unsafe exposure to the barrier, the resulting process becomes a supermartingale. This allows us to convert barrier inequalities into probabilistic tail bounds on the event that the total unsafe exposure exceeds a budget.

The main contributions of this paper are as follows:

- To the best of our knowledge, we introduce the first formulations of tolerant safety and tolerant reach-avoid properties for stochastic systems by quantifying unsafe exposure over an execution.
- We develop tolerant barrier certificates for discrete-time stochastic systems that certify probabilistic bounds on the tail of cumulative unsafe visits.
- We extend the framework to continuous-time stochastic differential equations by replacing unsafe visit counts with unsafe occupation time and deriving analogous barrier conditions via infinitesimal generators.

The remainder of the paper is organized as follows. Section II introduces the stochastic models and the tolerant safety and reach-avoid properties studied in this work. Section III presents tolerant barrier certificates for discrete-time stochastic systems. Section IV extends the framework to continuous-time stochastic differential equations.

Related Work. Barrier certificates were introduced by Prajna and Jadbabaie [12] as a Lyapunov-like approach for verifying safety of continuous and hybrid systems without explicit reachable-set computation. A barrier function separates unsafe states from all reachable states; its existence certifies that no trajectory enters the unsafe region. The idea was later extended to control barrier functions (CBFs) for safe controller synthesis [13], [14], making barrier-based methods a cornerstone of cyber-physical systems verification. Since then, significant effort has been devoted to developing more relaxed forms of barrier-certificate condition that still admit efficient synthesis, thereby leading to, e.g., exponential-type barrier certificates [15], Darboux-type barrier certificates [16], vector barrier certificates [17], and invariant barrier certificates [18]. Extensions to stochastic settings replace the strict decrease condition with a supermartingale condition that yields probabilistic safety guarantees [19], [20]. These certificates have become a standard tool for safety and reachability analysis of stochastic systems [21]–[23].

Martingale-based reasoning, which underpins our tolerant barrier certificates, has deep roots in the analysis of probabilistic programs. Chakarov and Sankaranarayanan [24] introduced the use of supermartingales for deriving probabilistic invariants and concentration bounds in probabilistic programs. Chatterjee et al. [25], [26] developed ranking supermartingale techniques for almost-sure and finite-time termination of probabilistic programs, exploiting optional-stopping theorems and tail inequalities. Takisaka et al. [27] unified ranking and repulsing supermartingales for reachability analysis.

An analogous notion of tolerant safety has been investigated in [28] for verifying continuous-time *non-stochastic* systems. The key technique therein is to use *occupation measures* to quantify the duration that a system trajectory stays in a particular set. We generalize this notion of tolerant safety to tolerant reach-avoid properties in the context of stochastic systems and solve the induced verification problems by means of tolerant barrier certificates.

Various alternative methods aim to reason about stochastic dynamics for non-tolerant specifications, including techniques based on sampling [29], [30], recurrence relation [31], [32], dynamic programming [33], [34], Markov abstraction [35], probabilistic model checking [36], [37] (for finite-state models), and various forms of value iteration [38]–[40] for determining reachability probabilities in Markov models.

II. Problem Formulation

We first briefly recall several standard notions used throughout the paper. Detailed treatments can be found in standard textbooks on stochastic processes and stochastic differential equations, e.g. [5], [41].

Notations. Let $\mathbb{N}$ and $\mathbb{R}$ be the set of non-negative integers and real numbers, respectively. A probability space $(\Omega, \mathcal{F}_\Omega, P_\Omega)$ consists of a sample space Ω, a σ-algebra $\mathcal{F}_\Omega$ on Ω, and a probability measure $P_\Omega \colon \mathcal{F}_\Omega \to [0,1]$ on the measurable space $(\Omega, \mathcal{F}_\Omega)$. An n-dimensional *random variable* X defined on the

probability space $(\Omega, \mathcal{F}_\Omega, P_\Omega)$ is an $\mathcal{F}_\Omega$-measurable function $X\colon \Omega \to \mathbb{R}^n$; its *expected value* (w.r.t. P_Ω) is denoted by $E[X]$. For a nonnegative random variable Z, an (upper-)tail bound is a function u satisfying $\mathbb{P}(Z > r) \le u(r)$ for every admissible threshold r.

A (discrete-time) *stochastic process* is a parameterized collection of random variables $\{X_n\}_{n\in\mathbb{N}}$. We occasionally drop the brackets in $\{X_n\}$ when it is clear from the context. A collection $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$ of σ-algebras on sets in $\mathcal{F}$ is a *filtration* if $\mathcal{F}_i \subseteq \mathcal{F}_j$ for all $i \le j$; intuitively, $\mathcal{F}_n$ carries the information known to an observer at step n. A random variable $\tau\colon \Omega \to \mathbb{N}$ is called a *stopping time* w.r.t. some filtration $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$ of $\mathcal{F}$ if $\{\omega \mid \tau(\omega) \le n\} \in \mathcal{F}_n$ for all $n \in \mathbb{N}$. A stochastic process $\{X_n\}$ adapted to a filtration $\{\mathcal{F}_n\}_{n\in\mathbb{N}}$ is called a *supermartingale* (resp. *submartingale*) if $E[X_n] < \infty$ for any $n \ge 0$ and $E[X_m \mid \mathcal{F}_n] \le X_n$ (resp. $E[X_m \mid \mathcal{F}_n] \ge X_n$) for all $0 \le n \le m$. That is, the conditional expected value of any future observation, given all the past observations, is no larger (resp. smaller) than the most recent observation.

Given a stochastic process X_n and a stopping time τ, let X_τ be $X_{\tau(\omega)}(\omega)$, the corresponding *stopped process* $\{X_n^\tau\}_{n\in\mathbb{N}}$ is defined by

$$X_n^\tau(\omega) \triangleq X_{n\wedge\tau}(\omega) = \begin{cases} X_n(\omega) & \text{if } n \le \tau(\omega)\ , \\ X_\tau(\omega) & \text{if } n > \tau(\omega)\ , \end{cases}$$

namely, a process that is forced to keep the same value after a prescribed (random) time τ. Moreover, if process X_n is a supermartingale (resp. submartingale), the stopped process $\{X_n^\tau\}_{n\in\mathbb{N}}$ is also a supermartingale (resp. submartingale). The above notions also have direct continuous-time analogues. In that setting, filtrations are indexed by $t \ge 0$, stopping times take values in $[0,\infty]$, and a process $\{X_t\}_{t\ge0}$ is a supermartingale if $E[X_t \mid \mathcal{F}s] \le X_s$ for all $0 \le s \le t$. The stopped process is defined similarly by $X_t^\tau \triangleq X_{t\wedge\tau}$.

We study stochastic systems evolving on a state space $X \subseteq \mathbb{R}^n$. The state space is equipped with three distinguished subsets: an initial set $S_0 \subseteq X$, an unsafe set $S_u \subseteq X$, and, for reach-avoid specifications, a target set $S_t \subseteq X$. Throughout the paper, we assume that $S_t \cap S_u = \varnothing$. Given an initial condition $x \in X$, we write $\mathbb{P}_x(\cdot)$ and $\mathbb{E}_x[\cdot]$ for probability and expectation conditioned on $X_0 = x$. We consider both discrete-time stochastic systems and continuous-time stochastic systems as follows.

Discrete-time stochastic systems. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and let $\{\xi_k\}_{k\ge0}$ be a sequence of independent random variables taking values in a measurable space Ξ. A discrete-time stochastic system is given by

$$X_{k+1} = f(X_k, \xi_k), \qquad k \in \mathbb{N}, \tag{1}$$

where $f : X \times \Xi \to X$ is measurable. For an initial state $X_0 = x$, the resulting state trajectory is denoted by $\{X_k^x\}_{k\ge0}$, or simply $\{X_k\}_{k\ge0}$ when the dependence on x is clear.

For the unsafe set S_u, the cumulative number of unsafe visits up to time N is

$$D_N := \sum_{k=0}^{N} \mathbf{1}_{S_u}(X_k),$$

where $\mathbf{1}_{S_u}$ denotes the indicator function of S_u. Since D_N is non-decreasing in N, the extended-valued limit $D_\infty := \lim_{N\to\infty} D_N \in \mathbb{N} \cup \{\infty\}$ is well defined. For the target set S_t, we define the first hitting time

$$T := \inf\{k \in \mathbb{N} : X_k \in S_t\},$$

with the convention $T = \infty$ if the target is never reached. When $T < \infty$, the number of unsafe visits accumulated *before* reaching the target is

$$D_{T-1} := \sum_{k=0}^{T-1} \mathbf{1}_{S_u}(X_k),$$

with convention $D_{-1} = 0$.

Continuous-time stochastic systems. We consider a class of continuous-time stochastic dynamics modeled by an Itô stochastic differential equation (SDE) [5] of the form

$$dX_t = b(X_t)\,dt + \sigma(X_t)\,dW_t, \qquad t \ge 0, \tag{2}$$

where W_t is an m-dimensional standard Wiener process (Brownian motion): $W_0 = 0$, its sample paths are continuous, and $W_t - W_s$ is independent of the past and has distribution $\mathcal{N}(0, (t-s)I_m)$ for $0 \le s < t$. The drift and diffusion coefficients are $b : X \to \mathbb{R}^n$ and $\sigma : X \to \mathbb{R}^{n\times m}$, respectively. We assume standard conditions [5] ensuring the existence and uniqueness of a solution for every initial condition $X_0 = x \in X$, and let $\{\mathcal{F}_t\}_{t\ge0}$ denote the natural filtration generated by the process. Intuitively, $\mathcal{F}_t$ is the history available up to time t: it records the initial condition and the randomness observed so far, but excludes future noise.

In continuous time, unsafe exposure is measured by *occupation time* rather than by a visit count. For the unsafe set S_u, the cumulative unsafe occupation time up to horizon t is

$$A_t := \int_0^t \mathbf{1}_{S_u}(X_s)\,ds, \qquad t \ge 0.$$

Again, A_t is non-decreasing in t, so the extended-valued limit

$$A_\infty := \lim_{t\to\infty} A_t \in [0,\infty]$$

is well defined. For the target set S_t, we define the first hitting time $\tau := \inf\{s \ge 0 : X_s \in S_t\}$, with the convention $\tau = \infty$ if the target is never reached. When $\tau < \infty$, the unsafe occupation accumulated before reaching the target is

$$A_\tau := \int_0^\tau \mathbf{1}_{S_u}(X_s)\,ds.$$

As mentioned in Section I, we study two exposure-based trajectory properties: tolerant safety and tolerant reach-avoid.

Tolerant safety. Classical safety requires the system to avoid S_u altogether. This is often too brittle in stochastic settings, where rare or short-lived unsafe excursions may be unavoidable. We therefore quantify safety in terms of a budget on total unsafe exposure.

Definition 1 (Discrete-time tolerant safety)**.** *Let $k \in \mathbb{N}$ and $\lambda \in [0,1]$. The discrete-time stochastic system (1) is said to satisfy (k,λ)-*tolerant safety *with respect to (S_0, S_u) if*

$$\mathbb{P}_x(D_\infty \le k) \ge \lambda, \qquad \forall x \in S_0. \tag{3}$$

Definition 2 (Continuous-time tolerant safety)**.** *Let $t \geq 0$ and $\lambda \in [0,1]$. The SDE (2) is said to satisfy (t,λ)-*tolerant safety *with respect to (S_0, S_u) if*

$$\mathbb{P}_x(A_\infty \leq t) \geq \lambda, \qquad \forall x \in S_0. \tag{4}$$

Therefore, in a discrete-time situation, the tolerance budget is an integer counting unsafe visits, whereas in a continuous-time situation, it is a nonnegative real number that measures unsafe residence time.

Tolerant reach-avoid. Another kind of specification we study is the tolerant reach-avoid property, which requires eventual reachability of a target set while limiting the unsafe exposure incurred before reaching that target.

Definition 3 (Discrete-time tolerant reach-avoid)**.** *Let $k \in \mathbb{N}$ and $\lambda \in [0,1]$. The discrete-time stochastic system (1) is said to satisfy (k,λ)-*tolerant reach-avoid *with respect to (S_0, S_u, S_t) if*

$$\mathbb{P}_x\big(T < \infty,\ D_{T-1} \leq k\big) \geq \lambda, \qquad \forall x \in S_0. \tag{5}$$

where $T < \infty$ in (5) explicitly requires S_t is reached.

Definition 4 (Continuous-time tolerant reach-avoid)**.** *Let $t \geq 0$ and $\lambda \in [0,1]$. The SDE (2) is said to satisfy (t,λ)-*tolerant reach-avoid *with respect to (S_0, S_u, S_t) if*

$$\mathbb{P}_x\big(\tau < \infty,\ A_\tau \leq t\big) \geq \lambda, \qquad \forall x \in S_0. \tag{6}$$

where $\tau < \infty$ in (6) explicitly requires S_t is reached.

Informally, tolerant safety asks whether the trajectory can keep its total unsafe exposure within budget over the whole execution. In discrete time this means "visit the unsafe set at most k times," while in continuous time it means "spend at most t units of time in the unsafe set." Tolerant reach-avoid adds a deadline-like condition: the target must eventually be reached, and only the unsafe exposure accumulated before that first hit counts toward the budget.

Problem formulation. The central goal of this paper is to verify the tolerant safety properties (3) and (4), and also to verify the reach-avoid properties (5) and (6). This amounts to deriving tight bounds on tolerant safety and tolerant reach-avoid probability of stochastic systems.

III. Tolerant Barriers for Discrete-Time Stochastic Systems

In this section, given a discrete-time stochastic system (1), we develop tolerant barrier certificates that yield lower bounds on the tolerant-safety and tolerant reach-avoid probabilities

$$\mathbb{P}_x(D_\infty \leq k) \qquad \text{and} \qquad \mathbb{P}_x(T < \infty,\ D_{T-1} \leq k),$$

for any $k \in \mathbb{N}$.

For a measurable function $B : X \to \mathbb{R}$, define the one-step expectation operator

$$(\mathcal{P}B)(x) := \mathbb{E}\big[B\big(f(x,\xi_0)\big)\big].$$

Equivalently, $(\mathcal{P}B)(x) = \mathbb{E}[B(X_{k+1}) \mid X_k = x]$.

A. Tolerant safety barriers

We begin with tolerant safety. The key idea is to require the barrier to decrease more aggressively whenever the current state lies in the unsafe set. This yields a supermartingale after accounting for the accumulated number of unsafe visits.

Definition 5 (Discrete-time tolerant safety barrier)**.** *Let $q \in (0,1]$ and $\gamma \geq 0$. A measurable function $B : X \to \mathbb{R}_{\geq 0}$ is called a (q,γ)-*tolerant safety barrier *for (1) with respect to (S_0, S_u) if*

$$B(x) \leq \gamma, \qquad \forall x \in S_0, \tag{7}$$
$$(\mathcal{P}B)(x) \leq q\,B(x) - 1, \qquad \forall x \in S_u, \tag{8}$$
$$(\mathcal{P}B)(x) \leq B(x), \qquad \forall x \in X \setminus S_u. \tag{9}$$

The three inequalities in tolerant safety barrier have distinct roles. Condition (7) says that the barrier starts small on the initial set, so γ measures the initial certificate budget. Condition (9) is the usual non-increase condition outside the unsafe set. The new ingredient is (8): whenever the system is currently unsafe, the expected next-step barrier value must not only contract by a factor q, but also decrease by one extra unit. Smaller q enforces a stronger decrease inside the unsafe set and is what eventually produces a geometric tail bound. The extra unit exactly matches the one-unit increase in the unsafe-visit counter. This is what allows the barrier to account for accumulated unsafe exposure.

To state the resulting tail bound compactly, define

$$\beta_q(k;\gamma) := \begin{cases} \dfrac{\gamma}{k+1}, & q = 1, \\ \big(1 + (1-q)\gamma\big)q^{k+1}, & q \in (0,1), \end{cases} \tag{10}$$

for any $k \in \mathbb{N}$ and $\gamma > 0$.

Theorem 1 (Tolerant safety)**.** *Suppose there exists a (q,γ)-tolerant safety barrier B in the sense of Definition 5. Then, for every $x \in S_0$ and every $k \in \mathbb{N}$,*

$$\mathbb{P}_x(D_\infty > k) \leq \beta_q(k;\gamma). \tag{11}$$

Equivalently, the system (1) satisfies $(k, 1-\beta_q(k;\gamma))$-tolerant safety for any $x \in S_0$, i.e.

$$\mathbb{P}_x(D_\infty \leq k) \geq 1 - \beta_q(k;\gamma).$$

Proof. We distinguish the cases $q = 1$ and $q \in (0,1)$.
For Case 1: $q = 1$**.** Define $Y_n := B(X_n) + D_{n-1}$ for $n \in \mathbb{N}$. We show that $\{Y_n\}_{n\geq 0}$ is a supermartingale with respect to the natural filtration. Indeed,

$$Y_{n+1} = B(X_{n+1}) + D_n.$$

If $X_n \in S_u$, then $D_n = D_{n-1} + 1$, and by (8), we have

$$\begin{aligned} \mathbb{E}[Y_{n+1} \mid \mathcal{F}_n] &= \mathbb{E}[B(X_{n+1}) \mid \mathcal{F}_n] + D_{n-1} + 1 \\ &\leq B(X_n) - 1 + D_{n-1} + 1 = Y_n. \end{aligned}$$

If $X_n \notin S_u$, then $D_n = D_{n-1}$, and by (9),

$$\begin{aligned} \mathbb{E}[Y_{n+1} \mid \mathcal{F}_n] &= \mathbb{E}[B(X_{n+1}) \mid \mathcal{F}_n] + D_{n-1} \\ &\leq B(X_n) + D_{n-1} = Y_n. \end{aligned}$$

Hence $\{Y_n\}_{n\geq 0}$ is a supermartingale. Since $x \in S_0$ and $B(x) \leq \gamma$, we have

$$\mathbb{E}_x[Y_n] \;\leq\; \mathbb{E}_x[Y_0] = B(x) \;\leq\; \gamma \qquad \forall n \geq 0.$$

Moreover, since $B(x) \geq 0$ on X, we have $D_{n-1} \leq Y_n$. Therefore, by Markov's inequality,

$$\mathbb{P}_x(D_{n-1} > k) \;\leq\; \mathbb{P}_x(Y_n \geq k+1) \;\leq\; \frac{\mathbb{E}_x[Y_n]}{k+1} \;\leq\; \frac{\gamma}{k+1}.$$

Since $D_n \uparrow D_\infty$ almost surely, the events $\{D_n > k\}$ increase to $\{D_\infty > k\}$. Thus,

$$\mathbb{P}_x(D_\infty > k) \;=\; \lim_{n\to\infty} \mathbb{P}_x(D_n > k) \;\leq\; \frac{\gamma}{k+1}.$$

For case 2: $q \in (0,1)$**.** Let

$$c_q := \frac{1}{1-q}, \qquad Z_n := q^{-D_{n-1}}\big(B(X_n) + c_q\big),$$

for $n \in \mathbb{N}$. We show that $\{Z_n\}_{n\geq 0}$ is a supermartingale. In fact, if $X_n \notin S_u$, then $D_n = D_{n-1}$ and (9) gives

$$\begin{aligned}\mathbb{E}[Z_{n+1} \mid \mathcal{F}_n] \;&=\; q^{-D_n}\,\mathbb{E}[B(X_{n+1}) + c_q \mid \mathcal{F}_n] \\ &\leq\; q^{-D_{n-1}}\big(B(X_n) + c_q\big) \;=\; Z_n.\end{aligned}$$

If $X_n \in S_u$, then $D_n = D_{n-1} + 1$ and (8) yields

$$\begin{aligned}\mathbb{E}[Z_{n+1} \mid \mathcal{F}_n] \;&=\; q^{-D_n}\,\mathbb{E}[B(X_{n+1}) + c_q \mid \mathcal{F}_n] \\ &\leq\; q^{-D_n}\big(qB(X_n) - 1 + c_q\big).\end{aligned}$$

Since $c_q(1-q) = 1$, we have $-1 + c_q = qc_q$, hence

$$qB(X_n) - 1 + c_q = q\big(B(X_n) + c_q\big).$$

Therefore,

$$\begin{aligned}\mathbb{E}[Z_{n+1} \mid \mathcal{F}_n] \;&\leq\; q^{-D_n}\,q\big(B(X_n) + c_q\big) \\ &=\; q^{-D_{n-1}}\big(B(X_n) + c_q\big) \;=\; Z_n.\end{aligned}$$

Hence $\{Z_n\}_{n\geq 0}$ is a supermartingale. Moreover, since $x \in S_0$ and $B(x) \leq \gamma$, then

$$\mathbb{E}_x[Z_n] \;\leq\; \mathbb{E}_x[Z_0] \;=\; B(x) + c_q \;\leq\; \gamma + c_q.$$

Also, since $B(x) \geq 0$, we have $Z_n \geq c_q\, q^{-D_{n-1}}$, thus

$$\mathbb{E}_x[q^{-D_{n-1}}] \leq \frac{\mathbb{E}_x[Z_n]}{c_q} \leq \frac{\gamma + c_q}{c_q} = 1 + (1-q)\gamma.$$

Applying Markov's inequality to the nonnegative random variable $q^{-D_{n-1}}$ gives

$$\mathbb{P}_x(D_{n-1} > k) = \mathbb{P}\big(q^{-D_{n-1}} \geq q^{-(k+1)}\big) \leq \big(1+(1-q)\gamma\big)q^{k+1}.$$

Again using $D_n \uparrow D_\infty$ almost surely, we obtain

$$\mathbb{P}_x(D_\infty > k) \;\leq\; \big(1 + (1-q)\gamma\big)q^{k+1}$$

Combining the above two cases proves (11). □

Theorem 1 shows that the same barrier template yields two distinct kinds of guarantees. When $q = 1$, the certificate controls the first moment of total unsafe exposure, which gives a reciprocal tail bound. When $q < 1$, the unsafe-set drift is stronger, and the certificate controls an exponential transform of the exposure count, which yields a geometric tail. Thus, the parameter q is not just a technical parameter; it directly controls the sharpness of the certified probability bound.

B. Tolerant reach-avoid barriers

We now turn to the tolerant reach-avoid property. As in the tolerant-safety case, we can use a barrier to quantify the tolerant reach-avoid probability.

Definition 6 (Discrete-time tolerant reach-avoid barrier)**.** *Let $q \in (0,1]$ and $\gamma \geq 0$. A measurable function $B : X \to \mathbb{R}_{\geq 0}$ is called a* (q,γ)-tolerant reach-avoid barrier *for* (1) *with respect to (S_0, S_u, S_t) if*

$$B(x) \leq \gamma, \qquad \forall x \in S_0, \tag{12}$$
$$(\mathcal{P}B)(x) \leq q\,B(x) - 1, \qquad \forall x \in S_u, \tag{13}$$
$$(\mathcal{P}B)(x) \leq B(x), \qquad \forall x \in X \setminus (S_u \cup S_t). \tag{14}$$

No condition is imposed on S_t, because the process will be stopped upon first hitting the target.

Theorem 2 (Tolerant reach-avoid)**.** *Suppose there exists a (q,γ)-tolerant reach-avoid barrier B in the sense of Definition 6. Then, for every $x \in S_0$ and every $k \in \mathbb{N}$,*

$$\mathbb{P}_x\big(T < \infty,\; D_{T-1} > k\big) \leq \beta_q(k;\gamma),$$

where $\beta_q(k;\gamma)$ is defined in (10)*.*

Proof. We distinguish the cases $q = 1$ and $q \in (0,1)$.
Case 1: $q = 1$**.** Define the stopped process

$$Y_n^T := B(X_{n\wedge T}) + D_{(n\wedge T)-1}, \qquad n \in \mathbb{N},$$

where $n \wedge T := \min\{n, T\}$. We claim that $\{Y_n^T\}_{n\geq 0}$ is a supermartingale. If $T \leq n$, then $Y_{n+1}^T = Y_n^T$, so the supermartingale property is immediate. Suppose instead that $T > n$, then $X_n \notin S_t$. If $X_n \in S_u$, then (13) gives

$$\begin{aligned}\mathbb{E}[Y_{n+1}^T \mid \mathcal{F}_n] \;&=\; \mathbb{E}[B(X_{n+1}) \mid \mathcal{F}_n] + D_{n-1} + 1 \\ &\leq B(X_n) - 1 + D_{n-1} + 1 \;=\; Y_n^T.\end{aligned}$$

If $X_n \in X \setminus (S_u \cup S_t)$, then $D_n = D_{n-1}$ and (14) yields

$$\begin{aligned}\mathbb{E}[Y_{n+1}^T \mid \mathcal{F}_n] \;&=\; \mathbb{E}[B(X_{n+1}) \mid \mathcal{F}_n] + D_{n-1} \\ &\leq\; B(X_n) + D_{n-1} \;=\; Y_n^T.\end{aligned}$$

Hence $\{Y_n^T\}_{n\geq 0}$ is a supermartingale. Since $B(x) \leq \gamma$ for $x \in S_0$, we have $\mathbb{E}_x[Y_n^T] \leq \mathbb{E}_x[Y_0^T] = B(x) \leq \gamma$. Now define the increasing events $E_n := \{T \leq n,\; D_{T-1} > k\}$. On E_n, we have $Y_n^T = B(X_T) + D_{T-1} \geq k+1$ because $B(X_T) \geq 0$. Therefore $\mathbb{P}_x(E_n) \leq \mathbb{E}_x[Y_n^T]/k+1 \leq \gamma/k+1$. Since $E_n \uparrow \{T < \infty,\; D_{T-1} > k\}$ as $n \to \infty$, continuity from below gives $\mathbb{P}_x(T < \infty,\; D_{T-1} > k) \;\leq\; \gamma/k+1$.

Case 2: $q \in (0,1)$**.** Let

$$c_q := \frac{1}{1-q}, \qquad Z_n^T := q^{-D_{(n\wedge T)-1}}\big(B(X_{n\wedge T}) + c_q\big).$$

Using the same case split as above, one checks that $\{Z_n^T\}_{n\geq 0}$ is a supermartingale. Indeed, the process is constant after time T, and before T the proof is identical to that of Thm 1. Hence

$$\mathbb{E}_x[Z_n^T] \leq \mathbb{E}_x[Z_0^T] = B(x) + c_q \leq \gamma + c_q.$$

On the event $E_n = \{T \leq n,\; D_{T-1} > k\}$, we have

$$Z_n^T \;=\; q^{-D_{T-1}}\big(B(X_T) + c_q\big) \;\geq\; c_q\, q^{-(k+1)}.$$

Therefore,

$$\mathbb{P}_x(E_n) \le \frac{\mathbb{E}_x[Z_n^T]}{c_q\, q^{-(k+1)}} \le \frac{\gamma + c_q}{c_q}\, q^{k+1} = \big(1 + (1-q)\gamma\big) q^{k+1}$$

Letting $n \to \infty$ yields

$$\mathbb{P}_x(T < \infty,\ D_{T-1} > k) \le \big(1 + (1-q)\gamma\big) q^{k+1}.$$

This completes the proof. □

Theorem 2 only controls the bad event that the target is reached after exceeding the unsafe-visit budget. It does not by itself guarantee that the target is reached with sufficiently high probability. That second ingredient is captured by an independent lower bound α on reachability. The following corollary is obtained by combining these two pieces.

Corollary 1 (Tolerant reach-avoid guarantee)**.** *Assume that there exists a constant $\alpha \in [0,1]$ such that*

$$\mathbb{P}_x(T < \infty) \ge \alpha, \qquad \forall x \in S_0. \tag{15}$$

If there also exists a (q,γ)-tolerant reach-avoid barrier, then for every $x \in S_0$ and every $k \in \mathbb{N}$,

$$\mathbb{P}_x(T < \infty,\ D_{T-1} \le k) \ge \alpha - \beta_q(k;\gamma).$$

In particular, whenever $\alpha - \beta_q(k;\gamma) \ge \lambda$, the system satisfies (k,λ)-tolerant reach-avoid from every $x \in S_0$.

Proof. For every $x \in S_0$,

$$\begin{aligned}&\mathbb{P}_x(T < \infty,\ D_{T-1} \le k)\\ =&\mathbb{P}_x(T < \infty) - \mathbb{P}_x(T < \infty,\ D_{T-1} > k) \ge \alpha - \beta_q(k;\gamma),\end{aligned}$$

where the first term is bounded below by (15) and the second term is bounded above by Theorem 2. □

Remark 1. *The lower bound α in Corollary 1 can be obtained independently using a variety of existing reachability-analysis techniques, see e.g., [24], [42]–[44]. In this paper, we assume that such a lower bound is already given.*

C. Counter-augmentation approach

A natural alternative to the direct tolerant-barrier construction is to encode the accumulated unsafe exposure as an additional state component and then apply a standard non-tolerant safety or reach-avoid certificate to the augmented system. We briefly describe this reduction and compare it with the tolerant-barrier method.

For a fixed tolerance budget $K \in \mathbb{N}$, introduce a counter

$$C_0 = 0, \qquad C_{n+1} = \mathrm{sat}_{K+1}\big(C_n + \mathbf{1}_{S_u}(X_n)\big),$$

where $\mathrm{sat}_{K+1}(r) := \min\{r, K+1\}$. This counter records the accumulated number of unsafe visits, saturated at $K+1$. We then augment the system state with this counter:

$$\widehat{X}_n^K := (X_n, C_n) \in \widehat{X}_K, \qquad \widehat{X}_K := X \times \{0,1,\dots,K+1\},$$

The corresponding augmented transition map is

$$\widehat{f}_K((x,c),\xi) := \Big(f(x,\xi), \mathrm{sat}_{K+1}\big(c + \mathbf{1}_{S_u}(x)\big)\Big).$$

The augmented initial set is $\widehat{S}_0^K := S_0 \times \{0\}$, and the augmented unsafe set is defined as

$$\widehat{S}_{\mathrm{bad}}^K := \Big\{(x,c) \in \widehat{X}_K : c + \mathbf{1}_{S_u}(x) > K\Big\}.$$

Then,

$$\{D_\infty > K\} = \Big\{\widehat{X}_n^K \in \widehat{S}_{\mathrm{bad}}^K \text{ for some } n \ge 0\Big\}.$$

Hence, for a fixed K, the tolerant-safety property is equivalent to standard safety of the augmented system with respect to $\widehat{S}_{\mathrm{bad}}^K$. Applying standard barrier-certificate methods [12], [19] for safety verification, one obtains

$$\mathbb{P}_x(D_\infty > K) \le \rho, \qquad \forall\, x \in S_0,$$

provided that there exist $\rho \ge 0$ and a family of barrier functions $V_c(x)$ for $c = 0,\dots,K+1$, such that

$$\begin{aligned}
&\rho \ge 0, \quad V_c(x) \ge 0, && x \in X,\ c = 0,\dots,K+1,\\
&V_0(x) \le \rho, && x \in S_0,\\
&V_K(x) \ge 1, && x \in S_u,\\
&V_{K+1}(x) \ge 1, && x \in X,\\
&\mathbb{E}[V_{c+1}(f(x,\xi_0))] \le V_c(x), && x \in S_u,\ c = 0,\dots,K-1,\\
&\mathbb{E}[V_c(f(x,\xi_0))] \le V_c(x), && x \in X \setminus S_u,\ c = 0,\dots,K.
\end{aligned}$$

The reach-avoid case can be handled analogously by defining the augmented target set

$$\widehat{S}_t^K := S_t \times \{0,1,\dots,K\}.$$

Then reaching $\widehat{S}_t^K$ before $\widehat{S}_{\mathrm{bad}}^K$ is equivalent to the original event $\{T < \infty,\ D_{T-1} \le K\}$.

Compared with tolerant barriers, the counter-augmentation approach applies to a fixed tolerance budget K. In contrast, tolerant barriers yield certified tail bounds parameterized by the budget k, thereby providing a single flexible certificate that can be evaluated at multiple tolerance levels. Moreover, the reduced counter-augmentation synthesis problem involves $K+2$ coupled component functions $V_0,\dots,V_{K+1}$, so both the number of barrier unknowns and the number of constraints grow with the tolerance budget. The tolerant-barrier conditions avoid this budget-dependent state augmentation by encoding the compensation for unsafe exposure directly on the original state space.

IV. Tolerant Barriers in Continuous Time

We now extend the discrete-time tolerant-barrier framework to stochastic differential equations. The underlying principle is unchanged: outside the unsafe set, the barrier should not increase, while inside the unsafe set, it should decrease sufficiently fast to offset the accumulation of unsafe exposure. The key difference is that, in continuous time, unsafe exposure is measured by occupation time rather than by the number of unsafe visits. For the SDE (2), our objective is to establish lower bounds on

$$\mathbb{P}_x(A_\infty \le \ell) \qquad \text{and} \qquad \mathbb{P}_x(\tau < \infty,\ A_\tau \le \ell),$$

where A_t, A_∞, and τ are defined in Section II.

To this end, we replace the discrete-time one-step operator P by the infinitesimal generator $\mathcal{L}$ (cf. [5, Chap. 7]), which

describes the instantaneous expected change of a smooth function along the SDE. Specifically, for any $B \in C^2(X;\mathbb{R})$,

$$(\mathcal{L}B)(x) := \nabla B(x)^\top b(x) + \frac{1}{2}\mathrm{Tr}\big(\sigma(x)\sigma(x)^\top \nabla^2 B(x)\big).$$

Using $\mathcal{L}$, we can formulate continuous-time tolerant barrier conditions that mirror those in the discrete-time setting.

A. Continuous-time tolerant safety barriers

We first consider tolerant safety in continuous time. As in the discrete-time case, the key idea is to require a stronger barrier decrease whenever the state lies in the unsafe set. The unsafe-visit count is now replaced by unsafe occupation time.

Definition 7 (Continuous-time tolerant safety barrier)**.** *Let $\eta \geq 0$ and $\gamma \geq 0$. A function $B \in C^2(X;\mathbb{R}_{\geq 0})$ is called an (η,γ)-*tolerant safety barrier *for (2) with respect to (S_0, S_u) if*

$$B(x) \leq \gamma, \qquad \forall x \in S_0, \tag{16}$$
$$(\mathcal{L}B)(x) \leq -1 - \eta B(x), \qquad \forall x \in S_u, \tag{17}$$
$$(\mathcal{L}B)(x) \leq 0, \qquad \forall x \in X \setminus S_u. \tag{18}$$

This is exactly the continuous-time analogue of Definition 5. Condition (17) is the continuous-time analogue of the unsafe drift condition in the discrete-time case. The term -1 accounts for the rate at which unsafe occupation accumulates, while the additional term $-\eta B(x)$ strengthens the decay within S_u. When $\eta = 0$, it reduces to the additive drift inequality

$$(\mathcal{L}B)(x) \leq -1, \qquad \forall x \in S_u,$$

which yields a first-moment bound on A_∞. When $\eta > 0$, it yields an exponential tail bound.

To state the resulting bound compactly, define

$$\beta_\eta(\ell;\gamma) := \begin{cases} \dfrac{\gamma}{\ell}, & \eta = 0, \\ (1+\eta\gamma)e^{-\eta\ell}, & \eta > 0, \end{cases} \qquad \ell > 0. \tag{19}$$

Theorem 3 (Tail bound for tolerant safety)**.** *Suppose there exists a (η,γ)-tolerant safety barrier B in the sense of Definition 7. Then, for every $x \in S_0$ and every $\ell > 0$,*

$$\mathbb{P}_x(A_\infty > \ell) \leq \beta_\eta(\ell;\gamma). \tag{20}$$

Equivalently,

$$\mathbb{P}_x(A_\infty \leq \ell) \geq 1 - \beta_\eta(\ell;\gamma). \tag{21}$$

Proof. We distinguish the cases $\eta = 0$ and $\eta > 0$.

Case 1: $\eta = 0$**.** Define $Y_t := B(X_t) + A_t$, for $t \geq 0$. Applying Itô's formula [5] gives

$$dY_t = \Big((\mathcal{L}B)(X_t) + \mathbf{1}_{S_u}(X_t)\Big)\,dt + \nabla B(X_t)^\top \sigma(X_t)\,dW_t.$$

If $X_t \in S_u$, then by (17), $(\mathcal{L}B)(X_t) + 1 \leq 0$. If $X_t \notin S_u$, then by (18), $(\mathcal{L}B)(X_t) + 0 \leq 0$. Hence the drift of Y_t is nonpositive, so $\{Y_t\}_{t\geq 0}$ is a supermartingale (Strictly speaking, a localization argument is needed to justify the supermartingale property. This step is somewhat technical but standard in the literature; see, e.g., [45], [46]). Since $x \in S_0$ and $B(x) \leq \gamma$,

$$\mathbb{E}_x[Y_t] \leq \mathbb{E}_x[Y_0] = B(x) \leq \gamma \qquad \forall t \geq 0.$$

Since $B(x) \geq 0$, we have $A_t \leq Y_t$, and therefore $\mathbb{E}_x[A_t] \leq \gamma$ for all $t \geq 0$. Since $A_t \uparrow A_\infty$ almost surely, the monotone convergence theorem [41] yields $\mathbb{E}_x[A_\infty] \leq \gamma$. Applying Markov's inequality gives $\mathbb{P}_x(A_\infty > \ell) \leq \gamma/\ell$.

Case 2: $\eta > 0$**.** Let $c_\eta := 1/\eta$ and $Z_t := e^{\eta A_t}\big(B(X_t) + c_\eta\big)$. Using Itô's formula together with

$$d(e^{\eta A_t}) = \eta \mathbf{1}_{S_u}(X_t) e^{\eta A_t}\,dt,$$

we obtain

$$dZ_t = e^{\eta A_t}\Big((\mathcal{L}B)(X_t) + \eta \mathbf{1}_{S_u}(X_t)\big(B(X_t) + c_\eta\big)\Big)\,dt + e^{\eta A_t}\nabla B(X_t)^\top \sigma(X_t)\,dW_t.$$

If $X_t \in S_u$, then by (17),

$$\begin{aligned}&(\mathcal{L}B)(X_t) + \eta\big(B(X_t) + c_\eta\big) \\ &\leq -1 - \eta B(X_t) + \eta B(X_t) + \eta c_\eta = -1 + 1 = 0.\end{aligned}$$

If $X_t \notin S_u$, then by (18), $(\mathcal{L}B)(X_t) + 0 \leq 0$. Hence the drift of Z_t is nonpositive, so $\{Z_t\}_{t\geq 0}$ is a supermartingale. Therefore, $\mathbb{E}_x[Z_t] \leq \mathbb{E}_x[Z_0] = B(x) + c_\eta \leq \gamma + c_\eta$. Since $B(x) \geq 0$, we have $c_\eta e^{\eta A_t} \leq Z_t$. Thus

$$\mathbb{E}_x[e^{\eta A_t}] \leq (\mathbb{E}_x[Z_t])/c_\eta \leq (\gamma + c_\eta)/c_\eta = 1 + \eta\gamma.$$

Letting $t \to \infty$ and using monotone convergence yields $\mathbb{E}_x[e^{\eta A_\infty}] \leq 1 + \eta\gamma$. Applying Markov's inequality gives

$$\mathbb{P}_x(A_\infty > \ell) = \mathbb{P}_x\big(e^{\eta A_\infty} \geq e^{\eta\ell}\big) \leq (1+\eta\gamma)e^{-\eta\ell}.$$

Combining the two cases proves (20). □

Theorem 3 is the continuous-time counterpart of Theorem 1. The case $\eta = 0$ yields a tail bound through a first-moment estimate, whereas every $\eta > 0$ yields a genuine exponential tail for the unsafe occupation time.

B. Continuous-time tolerant reach-avoid barriers

We now turn to tolerant reach-avoid in continuous time.

Definition 8 (Continuous-time tolerant reach-avoid barrier)**.** *Let $\eta \geq 0$ and $\gamma \geq 0$. A function $B \in C^2(X;\mathbb{R}_{\geq 0})$ is called an (η,γ)-*tolerant reach-avoid barrier *for (2) with respect to (S_0, S_u, S_t) if*

$$B(x) \leq \gamma, \qquad \forall x \in S_0, \tag{22}$$
$$(\mathcal{L}B)(x) \leq -1 - \eta B(x), \qquad \forall x \in S_u, \tag{23}$$
$$(\mathcal{L}B)(x) \leq 0, \qquad \forall x \in X \setminus (S_u \cup S_t). \tag{24}$$

No condition is imposed on S_t, since the process will be stopped at the first hitting time of the target.

Theorem 4 (Tail bound for the reach-avoid bad event)**.** *Suppose there exists a (η,γ)-tolerant reach-avoid barrier B in the sense of Definition 8. Then, for every $x \in S_0$ and $\ell > 0$,*

$$\mathbb{P}_x(\tau < \infty,\ A_\tau > \ell) \leq \beta_\eta(\ell;\gamma), \tag{25}$$

where $\beta_\eta(\ell;\gamma)$ is defined in (19).

Proof. We again distinguish the cases $\eta = 0$ and $\eta > 0$.

Case 1: $\eta = 0$**.** Define the stopped process $Y_t^\tau := B(X_{t\wedge\tau}) + A_{t\wedge\tau}$ for $t \geq 0$. If $\tau \leq t$, then the process is constant after

time τ. Before time τ, the state lies outside S_t. If moreover $X_t \in S_u$, then by (23), $(\mathcal{L}B)(X_t) + 1 \le 0$. If $X_t \in X \setminus (S_u \cup S_t)$, then by (24), $(\mathcal{L}B)(X_t) + 0 \le 0$. Hence $\{Y_t^\tau\}_{t\ge 0}$ is a supermartingale. Since $x \in S_0$ and $B(x) \le \gamma$, we have $\mathbb{E}_x[Y_t^\tau] \le \mathbb{E}_x[Y_0^\tau] = B(x) \le \gamma$. Moreover, let $E_t := \{\tau \le t,\ A_\tau > \ell\}$, one has

$$Y_t^\tau = B(X_\tau) + A_\tau \ge \ell,$$

over E_t since $B(x) \ge 0$. Therefore,

$$\mathbb{P}_x(E_t) \le \frac{\mathbb{E}_x[Y_t^\tau]}{\ell} \le \frac{\gamma}{\ell}.$$

As $t \to \infty$, the events E_t increase to $\{\tau < \infty,\ A_\tau > \ell\}$. Thus, $\mathbb{P}_x(\tau < \infty,\ A_\tau > \ell) \le \gamma/\ell$.

Case 2: $\eta > 0$**.** Let

$$c_\eta := \frac{1}{\eta}, \qquad Z_t^\tau := e^{\eta A_{t\wedge\tau}}\big(B(X_{t\wedge\tau}) + c_\eta\big).$$

Using the same calculation as in the proof of Theorem 3, now applied up to the stopping time τ, one checks that $\{Z_t^\tau\}_{t\ge 0}$ is a supermartingale. Hence

$$\mathbb{E}_x[Z_t^\tau] \le \mathbb{E}_x[Z_0^\tau] = B(x) + c_\eta \le \gamma + c_\eta.$$

On the event $E_t = \{\tau \le t,\ A_\tau > \ell\}$, we have

$$Z_t^\tau = e^{\eta A_\tau}\big(B(X_\tau) + c_\eta\big) \ge c_\eta e^{\eta\ell}.$$

Therefore,

$$\mathbb{P}_x(E_t) \le \frac{\mathbb{E}_x[Z_t^\tau]}{c_\eta e^{\eta\ell}} \le \frac{\gamma + c_\eta}{c_\eta} e^{-\eta\ell} = (1 + \eta\gamma)e^{-\eta\ell}.$$

Letting $t \to \infty$ gives $\mathbb{P}_x(\tau < \infty,\ A_\tau > \ell) \le (1 + \eta\gamma)e^{-\eta\ell}$. This proves (25). □

Theorem 4 upper-bounds the probability of the bad event that the target is eventually reached only after the unsafe occupation exceeds ℓ. As in the discrete-time setting, this bound can be combined with any lower bound on reachability.

Corollary 2 (Tolerant reach-avoid guarantee)**.** *Assume that there exists a constant $\alpha \in [0, 1]$ such that*

$$\mathbb{P}_x(\tau < \infty) \ge \alpha, \qquad \forall x \in S_0. \tag{26}$$

If there also exists a (η, γ)-tolerant reach-avoid barrier, then for every $x \in S_0$ and every $\ell > 0$,

$$\mathbb{P}_x(\tau < \infty,\ A_\tau \le \ell) \ge \alpha - \beta_\eta(\ell; \gamma).$$

In particular, whenever $\alpha - \beta_\eta(\ell; \gamma) \ge \lambda$, the system satisfies (ℓ, λ)-tolerant reach-avoid from every $x \in S_0$.

Proof. For every $x \in S_0$,

$$\begin{aligned}\mathbb{P}_x(\tau < \infty,\ A_\tau \le \ell) &= \mathbb{P}_x(\tau < \infty) - \mathbb{P}_x(\tau < \infty,\ A_\tau > \ell)\\ &\ge \alpha - \beta_\eta(\ell; \gamma),\end{aligned}$$

where the first term is bounded below by (26) and the second term is bounded above by Theorem 4. □

Remark 2. *Similar to the discrete-time case, the lower bound α in Corollary 2 can be obtained independently using existing reachability-analysis techniques for stochastic differential equations, see e.g., [45], [47].*

The barrier conditions in Def. 7 and 8 have the same structure as their discrete-time counterparts: the barrier condition depends only on whether the current state lies in the unsafe set, outside it, or, in the reach-avoid setting, outside both the unsafe and target sets. This makes them natural continuous-time analogues of the discrete-time tolerant barriers.

Moreover, the parameter η plays the same conceptual role as the parameter q in Section III. When $\eta = 0$, the barrier controls the expected unsafe occupation time and implies a reciprocal tail bound. When $\eta > 0$, the strengthened unsafe-set drift yields an exponential tail in the occupation budget.

Remark 3 (Graded tolerance)**.** *The same idea can be extended beyond binary unsafe exposure. Instead of using the indicator $\mathbf{1}_{S_u}$, one may introduce a nonnegative severity function $r : X \to \mathbb{R}_{\ge 0}$ and measure weighted exposure by $\sum_{k=0}^{N} r(X_k)$ or $\int_0^t r(X_s)\,ds$. This would allow different regions to contribute different amounts of exposure. The present paper focuses on the binary case $r = \mathbf{1}_{S_u}$ and leaves a systematic treatment of graded tolerance to future work.*

Continuous-time counter augmentation. One may be interested in whether the continuous-time occupation budget can be handled analogously by adding one state variable. However, the counter-augmentation idea does not carry over as directly to the continuous-time setting. Formally, for a fixed budget $\ell > 0$, one could introduce $Y_t := A_t = \int_0^t \mathbf{1}_{S_u}(X_s)\,ds$, then

$$dX_t = b(X_t)\,dt + \sigma(X_t)\,dW_t, \qquad dY_t = \mathbf{1}_{S_u}(X_t)\,dt.$$

Then exceeding the occupation-time budget corresponds to hitting the augmented boundary $Y_t = \ell$.

This formulation is less direct than discrete-time counter augmentation for several reasons. First, the added drift $\mathbf{1}_{S_u}(x)$ is discontinuous at the boundary of S_u. Thus, even if the original coefficients b and σ are polynomial, the augmented dynamics need not define a polynomial SDE. This is inconvenient both for formulating barrier conditions and for applying the synthesis procedure discussed in Section V. Second, the augmented barrier would depend on an additional variable y, and the corresponding constraints would be imposed on $X \times [0, \ell]$ rather than on X. Hence the synthesis complexity increases, and the certificate is tied to the particular budget ℓ.

V. Synthesizing Tolerant Barrier Certificates

This section briefly discusses how the barrier conditions of Sections III and IV can be synthesized. Our focus is on polynomial stochastic systems and semialgebraic state sets, where the barrier constraints can be relaxed to sum-of-squares (SOS) programs [48], which can be further translated to semi-definite programming (SDP) problems [49] that admit polynomial-time algorithms implemented by many off-the-shelf SDP solvers.

We restrict attention to polynomial barrier templates of a fixed degree. For instance, in the discrete-time setting, choose

$$B_a(x) = \sum_{|\alpha| \le d} a_\alpha x^\alpha,$$

where $a := \{a_\alpha\}$ is the vector of unknown coefficients. We assume that the sets X, S_0, S_u, and S_t are semialgebraic, i.e.,

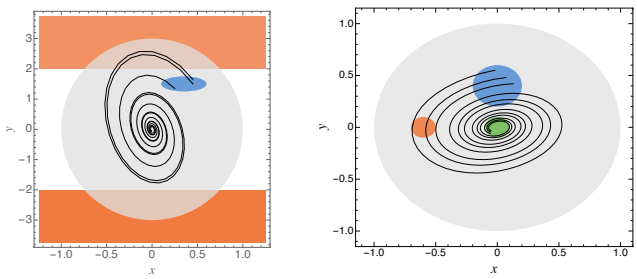

(a) circuit (safety) (b) vanderpol (reach-avoid)

Fig. 2: Portraits of two selected examples. The solid curves (connecting states between consecutive time steps) depict randomly sampled trajectories.

described by finitely many polynomial inequalities. Under this assumption, pointwise nonnegativity constraints on these sets can be enforced by SOS relaxations.

Discrete-time tolerant barriers. Consider first the discrete-time tolerant safety barrier conditions (7)–(9). For a polynomial template B_a, synthesis amounts to finding coefficients a together with parameters $q \in (0,1]$ and $\gamma \geq 0$ such that

$$\begin{aligned} B_a(x) - \gamma &\leq 0, && x \in S_0, \\ \mathbb{E}[B_a(f(x,\xi))] - qB_a(x) + 1 &\leq 0, && x \in S_u, \\ \mathbb{E}[B_a(f(x,\xi))] - B_a(x) &\leq 0, && x \in X \setminus S_u, \\ -B_a(x) &\leq 0, && x \in X. \end{aligned}$$

When the moments of ξ are known and f is polynomial in (x,ξ), the expressions $\mathbb{E}[B_a(f(x,\xi))]$ are again polynomial in x. Hence the above constraints can be relaxed to an SOS feasibility problem.

The same idea applies to the discrete-time reach-avoid barrier conditions (12)–(14), replacing $X \setminus S_u$ by $X \setminus (S_u \cup S_t)$. In practice, one may either fix q and minimize γ, or jointly search over a grid of q values and select the pair (q,γ) that minimizes the certified tail bound $\beta_q(k;\gamma)$ for a target exposure level k.

Continuous-time tolerant barriers. For the SDE model (2), synthesis proceeds analogously using the generator inequalities in (16)–(18) and (22)–(24). Given a polynomial template B_a, the constraints take the form

$$\begin{aligned} B_a(x) - \gamma &\leq 0, && x \in S_0, \\ (\mathcal{L}B_a)(x) + 1 + \eta B_a(x) &\leq 0, && x \in S_u, \\ (\mathcal{L}B_a)(x) &\leq 0, && x \in X \setminus S_u, \end{aligned}$$

for safety, and with $X \setminus S_u$ replaced by $X \setminus (S_u \cup S_t)$ for reach-avoid. When the drift b and diffusion σ are polynomial, the generator $\mathcal{L}B_a$ is again polynomial, and the constraints can be relaxed to SOS conditions. As in discrete time, one may fix η and minimize γ, or search over $\eta \geq 0$ to optimize the certified continuous-time tail bound $\beta_\eta(\ell;\gamma)$ for a target unsafe-occupation budget ℓ.

Observe that all the constraints above share a common form

$$V^a(x) \geq 0 \quad \text{for} \quad x \in \left\{ x \;\middle|\; \bigvee\nolimits_{i=0}^{m} \bigwedge\nolimits_{j=0}^{l} P_{ij}(x) \triangleright 0 \right\}, \quad (27)$$

i.e., $V^a(\cdot)$ is a parametrized polynomial that is non-negative over a semi-algebraic set. Based on the well-known Putinar's Positivstellensatz [57], constraint (27) can be reformulated into a group of SOS constraints:

$$\begin{aligned} & V^a(x) + \sum_{j=0}^{l} s_{ij}(x) \cdot P_{ij}(x) \;\in\; \mathrm{sos}[x], \quad \text{for } 0 \leq i \leq m\,, \\ & s_{ij} \;\in\; \mathrm{sos}[x], \quad \text{for } 0 \leq i \leq m,\, 0 \leq j \leq l \end{aligned}$$

where $\mathrm{sos}[x] \triangleq \{g(x) \in \mathbb{R}[x] \mid g = h_1^2 + h_2^2 + \cdots + h_k^2\}$ denotes the set of all sum-of-squares polynomials in x (over the reals). Consequently, the constraints for tolerant barrier certificates can be encoded as SOS programming problems, which can be further solved by an off-the-shelf SOS/SDP solver.

Interpretation of synthesis failure. The barrier conditions are sufficient but not necessary for the tolerant property. Moreover, the implemented search is restricted to a chosen polynomial degree and a particular SOS relaxation. Therefore, failure to synthesize a barrier does not imply that the system violates tolerant safety or tolerant reach-avoid. It may instead mean that no certificate exists in the selected template class, that the SOS relaxation is too conservative, or that the numerical solver cannot resolve the instance reliably. Such a failure is therefore inconclusive; only a successfully verified certificate yields a sound probabilistic guarantee.

VI. Implementation and Experimental Results

To demonstrate the effectiveness of our SOS-based synthesis approach for tolerant barrier certificates, we developed a prototype implementation in Julia, interfaced with Mosek [58] to solve the underlying SOS programs. We evaluated our approach on a collection of benchmark examples for both continuous-time and discrete-time stochastic systems drawn from the literature. All experiments were conducted on an Apple MacBook equipped with an M4 chip and 24 GB of memory, running macOS.

Design of Experiments. All benchmarks are *infinite-state* stochastic systems that may enter the unsafe region with non-negligible probability; representative trajectories are shown in Fig. 2. We consider four verification settings: discrete-time safety, discrete-time reach-avoid, continuous-time safety, and continuous-time reach-avoid. Since the evaluation procedure is similar across all four cases, we explain the discrete-time safety setting in detail and summarize the others analogously.

For each discrete-time safety benchmark, we instantiate the tolerant barrier conditions using several decay parameter q, and solve the corresponding SOS programs. For each q, we search for a polynomial tolerant barrier template of degree ranging from 2 to $d_{\max}$. To balance effectiveness and computational cost, we set $d_{\max} = 16, 12, 10, 6$ for systems of dimension $1, 2, 3, 4$, respectively. For each selected barrier, we compute certified lower bounds on $\mathbb{P}_x(D_\infty \leq k)$ for several budgets k.

For each benchmark, we report the synthesis time, the optimal SOS value γ, the certified lower bound on $\mathbb{P}_x(D_\infty \leq k)$, and the threshold $k_{0.95}$, defined as the smallest budget k for which the certified lower bound is at least 0.95. For

TABLE I: Experimental results for discrete-time stochastic systems.

(a) (k, λ)-Tolerant Safety Results (Discrete Time)

Benchmark	dim	q	time(s)	γ	lower bound on $\mathbb{P}_x(D_\infty \leq k)$ for different k											$k_{0.95}$
					1	2	5	10	15	20	30	40	50	80	100	
randomwalk1	1	1.00	0.043	9.452	0.000	0.000	0.000	0.141	0.409	0.550	0.695	0.769	0.815	0.883	0.906	189
		0.95	0.031	35.454	0.000	0.000	0.000	0.000	0.000	0.056	0.435	0.661	0.797	0.956	0.984	78
randomwalk2	1	1.00	0.031	4.009	0.000	0.000	0.332	0.636	0.749	0.809	0.871	0.902	0.921	0.951	0.960	80
		0.90	0.039	14.717	0.000	0.000	0.000	0.224	0.542	0.730	0.906	0.967	0.989	1.000	1.000	37
equil [19]	2	1.00	4.740	7.797	0.000	0.000	0.000	0.291	0.513	0.629	0.748	0.810	0.847	0.904	0.923	155
		0.80	2.663	25.566	0.000	0.000	0.000	0.475	0.828	0.944	0.994	0.999	1.000	1.000	1.000	21
osc [50]	2	1.00	0.121	5.462	0.000	0.000	0.090	0.503	0.659	0.740	0.824	0.867	0.893	0.933	0.946	109
		0.60	0.112	37.017	0.000	0.000	0.263	0.943	0.996	1.000	1.000	1.000	1.000	1.000	1.000	11
liederivative [51]	2	0.95	0.569	13.010	0.000	0.000	0.000	0.061	0.274	0.438	0.663	0.798	0.879	0.974	0.991	68
		0.60	0.881	586.999	0.000	0.000	0.000	0.145	0.933	0.995	1.000	1.000	1.000	1.000	1.000	16
circuit [52]	2	1.00	0.082	14.973	0.000	0.000	0.000	0.000	0.064	0.287	0.517	0.635	0.706	0.815	0.852	299
		0.95	0.113	25.056	0.000	0.000	0.000	0.000	0.008	0.233	0.541	0.725	0.835	0.965	0.987	74
arch [53]	2	0.95	8.326	0.379	0.080	0.126	0.251	0.420	0.552	0.653	0.792	0.876	0.926	0.984	0.994	58
		0.90	1.572	0.915	0.116	0.204	0.420	0.657	0.798	0.881	0.958	0.985	0.995	1.000	1.000	29
		0.60	2.485	215.782	0.000	0.000	0.000	0.683	0.975	0.998	1.000	1.000	1.000	1.000	1.000	14
lyapunov [54]	3	1.00	76.257	4.013	0.000	0.000	0.331	0.635	0.749	0.809	0.871	0.902	0.921	0.950	0.960	80
		0.95	80.644	5.062	0.000	0.000	0.079	0.287	0.448	0.573	0.744	0.847	0.908	0.980	0.993	62
		0.90	77.976	6.020	0.000	0.000	0.149	0.497	0.703	0.825	0.939	0.979	0.993	1.000	1.000	32
stable [55]	4	0.95	183.048	0.068	0.094	0.140	0.262	0.429	0.558	0.658	0.795	0.877	0.927	0.984	0.994	58
		0.90	166.486	0.088	0.183	0.265	0.464	0.683	0.813	0.890	0.962	0.987	0.995	1.000	1.000	28
		0.80	192.830	0.126	0.344	0.475	0.731	0.912	0.971	0.991	0.999	1.000	1.000	1.000	1.000	13
		0.60	190.722	0.326	0.593	0.756	0.947	0.996	1.000	1.000	1.000	1.000	1.000	1.000	1.000	6

(b) (k, λ)-Tolerant Reach-Avoid Results (Discrete Time)

Benchmark	dim	q	time(s)	γ	lower bound on $\mathbb{P}_x(T < \infty,\ D_{T-1} \leq k)$ for different k											$k_{0.95}$
					1	2	5	10	15	20	30	40	50	80	100	
randomdescent1	1	1.00	0.027	4.910	0.000	0.000	0.182	0.554	0.693	0.766	0.842	0.880	0.904	0.939	0.951	98
		0.60	0.009	72.119	0.000	0.000	0.000	0.892	0.992	0.999	1.000	1.000	1.000	1.000	1.000	12
randomdescent2	1	1.00	0.011	2.711	0.000	0.096	0.548	0.754	0.831	0.871	0.913	0.934	0.947	0.967	0.973	54
		0.60	0.019	8.712	0.000	0.031	0.791	0.984	0.999	1.000	1.000	1.000	1.000	1.000	1.000	8
vanderpol [50]	2	0.95	5.321	9.416	0.000	0.000	0.000	0.163	0.353	0.499	0.700	0.820	0.892	0.977	0.992	65
		0.60	2.650	934.876	0.000	0.000	0.000	0.000	0.894	0.992	1.000	1.000	1.000	1.000	1.000	17
equil [19]	2	0.95	6.329	9.716	0.000	0.000	0.000	0.155	0.346	0.494	0.697	0.819	0.891	0.977	0.992	66
		0.80	0.976	31.304	0.000	0.000	0.000	0.376	0.796	0.933	0.993	0.999	1.000	1.000	1.000	22
liederivative [51]	2	1.00	0.622	3.285	0.000	0.000	0.452	0.701	0.795	0.844	0.894	0.920	0.936	0.959	0.967	65
		0.95	0.625	3.699	0.000	0.000	0.129	0.326	0.478	0.596	0.758	0.855	0.913	0.981	0.993	61
		0.60	0.611	12.612	0.000	0.000	0.718	0.978	0.998	1.000	1.000	1.000	1.000	1.000	1.000	9
circuit [52]	2	1.00	0.028	14.409	0.000	0.000	0.000	0.000	0.099	0.314	0.535	0.649	0.717	0.822	0.857	288
		0.95	0.027	23.512	0.000	0.000	0.000	0.000	0.042	0.259	0.556	0.734	0.841	0.966	0.988	73
		0.80	0.055	191.186	0.000	0.000	0.000	0.000	0.000	0.638	0.961	0.996	1.000	1.000	1.000	29
cwh [56]	4	1.00	0.202	8.386	0.000	0.000	0.000	0.238	0.476	0.601	0.729	0.795	0.836	0.896	0.917	167
		0.95	0.200	11.521	0.000	0.000	0.000	0.104	0.306	0.463	0.679	0.808	0.885	0.975	0.991	67
		0.90	0.227	18.931	0.000	0.000	0.000	0.092	0.464	0.683	0.890	0.962	0.987	0.999	1.000	38
		0.80	0.249	98.047	0.000	0.000	0.000	0.000	0.420	0.810	0.980	0.998	1.000	1.000	1.000	26

readability, the reported certified lower bounds are rounded to three decimal places. Smaller values of $k_{0.95}$ therefore indicate tighter guarantees. The same reporting format is used for the other three settings. For all reach-avoid benchmarks, the reachability probability α in Corollaries 1 and 2 is certified to be 1; this does not mean the benchmarks are restrictive, but rather helps isolate and better evaluate the tolerant part of our method. In the continuous-time case, unsafe exposure is measured by occupation time, and we report the analogous threshold $\ell_{0.95}$. We further report additional empirical findings structured as research questions (RQs).

RQ1: How effective is our approach across diverse stochastic benchmarks?

Answer: Our approach is effective across a broad and diverse collection of stochastic benchmarks. As shown in Tables I and II, the synthesis procedure successfully constructs a tolerant barrier certificate for every benchmark considered, covering both discrete-time and continuous-time systems as well as tolerant safety and tolerant reach-avoid properties. These results suggest that the proposed framework is not confined to a narrow class of examples, but instead applies consistently across systems with different dynamics, dimensions, and specification types. In particular, the method remains successful on higher-dimensional and nonlinear benchmarks such as `lyapunov`, `stable`, and `cwh`, highlighting both the practical effectiveness and the generality of our approach.

TABLE II: Experimental results for continuous-time systems.

(a) (t, λ)-Tolerant Safety Results (Continuous Time)

Benchmark	dim	η	time(s)	γ	lower bound on $\mathbb{P}_x(A_\infty \leq \ell)$ for different l											$\ell_{0.95}$
					1	2	3	4	5	6	7	10	20	50	100	
randomwalk1	1	0.00	0.035	9.410	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.059	0.529	0.812	0.906	188.206
randomwalk2	1	0.00	0.071	3.853	0.000	0.000	0.000	0.037	0.229	0.358	0.450	0.615	0.807	0.923	0.961	77.069
		0.10	0.051	11.728	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.201	0.706	0.985	1.000	37.717
equil [19]	2	0.00	0.236	8.016	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.198	0.599	0.840	0.920	160.330
		0.50	0.266	87.917	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.697	0.998	1.000	1.000	13.603
osc [50]	2	0.00	0.132	5.340	0.000	0.000	0.000	0.000	0.000	0.110	0.237	0.466	0.733	0.893	0.947	106.810
		0.80	0.121	76.673	0.000	0.000	0.000	0.000	0.000	0.487	0.769	0.979	1.000	1.000	1.000	8.910
liederivative [51]	2	0.10	0.101	16.810	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.014	0.637	0.982	1.000	39.819
		0.80	0.069	28089.530	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.997	1.000	1.000	16.270
circuit [52]	2	0.10	0.241	11.225	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.219	0.713	0.986	1.000	37.483
		0.20	0.111	20.695	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.305	0.906	1.000	1.000	23.163
arch [53]	2	0.10	0.443	0.517	0.048	0.139	0.221	0.295	0.362	0.423	0.478	0.613	0.858	0.993	1.000	30.461
		0.20	1.197	0.923	0.030	0.206	0.350	0.468	0.564	0.643	0.708	0.840	0.978	1.000	1.000	15.825
		0.50	0.456	15.330	0.000	0.000	0.000	0.000	0.289	0.569	0.738	0.942	1.000	1.000	1.000	10.310
lyapunov [54]	3	0.00	0.881	1.989	0.000	0.005	0.337	0.503	0.602	0.668	0.716	0.801	0.901	0.960	0.980	39.785
		0.10	0.911	2.442	0.000	0.000	0.078	0.166	0.245	0.317	0.382	0.542	0.832	0.992	1.000	32.143
		0.50	1.703	16.110	0.000	0.000	0.000	0.000	0.257	0.549	0.727	0.939	1.000	1.000	1.000	10.398
stable [55]	4	0.10	0.430	0.089	0.087	0.174	0.253	0.324	0.388	0.446	0.499	0.629	0.863	0.993	1.000	30.046
		0.20	0.425	0.123	0.161	0.313	0.438	0.540	0.623	0.691	0.747	0.861	0.981	1.000	1.000	15.100
		0.50	0.467	0.326	0.295	0.572	0.740	0.843	0.905	0.942	0.965	0.992	1.000	1.000	1.000	6.294
		0.80	0.658	4.061	0.000	0.142	0.615	0.827	0.922	0.965	0.984	0.999	1.000	1.000	1.000	5.553

(b) (t, λ)-Tolerant Reach-Avoid Results (Continuous Time)

Benchmark	dim	η	time(s)	γ	lower bound on $\mathbb{P}_x(\tau < \infty,\ A_\tau \leq \ell)$ for different l											$\ell_{0.95}$
					1	2	3	4	5	6	7	10	20	50	100	
randomdescent1	1	0.00	0.015	5.139	0.000	0.000	0.000	0.000	0.000	0.144	0.266	0.486	0.743	0.897	0.949	102.770
		0.80	0.025	180.423	0.000	0.000	0.000	0.000	0.000	0.000	0.463	0.951	1.000	1.000	1.000	9.968
randomdescent2	1	0.00	0.013	2.542	0.000	0.000	0.153	0.364	0.492	0.576	0.637	0.746	0.873	0.949	0.975	50.849
		0.80	0.010	16.767	0.000	0.000	0.000	0.412	0.736	0.881	0.947	0.995	1.000	1.000	1.000	7.080
vanderpol [50]	2	0.10	0.333	9.157	0.000	0.000	0.000	0.000	0.000	0.000	0.049	0.295	0.741	0.987	1.000	36.458
		0.80	0.179	1332.818	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.642	1.000	1.000	1.000	12.461
equil [19]	2	0.10	0.263	11.986	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.191	0.702	0.985	1.000	37.836
		0.50	0.160	99.422	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.658	0.998	1.000	1.000	13.844
liederivative [51]	2	0.00	0.117	2.600	0.000	0.000	0.133	0.350	0.480	0.567	0.629	0.740	0.870	0.948	0.974	52.005
		0.10	0.115	3.087	0.000	0.000	0.030	0.123	0.206	0.282	0.350	0.519	0.823	0.991	1.000	32.648
		0.80	0.115	14.390	0.000	0.000	0.000	0.490	0.771	0.897	0.954	0.996	1.000	1.000	1.000	6.903
circuit [52]	2	0.00	0.076	7.555	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.245	0.622	0.849	0.924	151.091
		0.10	0.082	11.213	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.220	0.713	0.986	1.000	37.478
		0.50	0.028	234.536	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.203	0.995	1.000	1.000	15.537
cwh [56]	4	0.00	0.039	6954.075	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	139081

Moreover, the approach's effectiveness goes beyond mere feasibility. For several benchmarks, the synthesized certificates provide strong guarantees, with certified lower bounds approaching 1.000 for sufficiently large exposure budgets, after rounding to three decimal places. This behavior appears across multiple settings, showing that the method can, in some cases, certify essentially perfect tolerant guarantees. Representative examples include `arch`, `stable`, and `lyapunov` in discrete-time safety; `liederivative` and `cwh` in discrete-time reach-avoid; and `randomwalk1`, `arch`, `stable`, and `cwh` in the continuous-time experiments. Overall, these results demonstrate both the broad applicability of our framework and its strong empirical effectiveness. ◁

RQ2: How efficient and scalable is our synthesis process?

Answer: The proposed synthesis procedure is computationally practical for the benchmark suite considered in this work. All reported instances, including the four-dimensional benchmarks `stable` and `cwh`, are solved within roughly 200 seconds per tested value of q or η. These suggest that SOS-based tolerant barrier synthesis is practical for the stochastic systems of moderate dimension and nontrivial nonlinear dynamics. To assess scalability beyond the low-dimensional benchmarks, we additionally consider a coupled stochastic sink family with dimensions $n \in \{6, 8, 10, 12, 14, 16\}$. Its discrete-time dynamics are $X_{k+1} = (A_n + \xi_k D_n)X_k$, where $\{\xi_k\}_{k\geq 0}$ are i.i.d. random variables with $\xi_k \sim \mathcal{U}[-1, 1]$, $(A_n)_{ij} = 0.58\,\mathbf{1}_{\{j=i\}} + 0.04\,\mathbf{1}_{\{j=i-1\}} - 0.035\,\mathbf{1}_{\{j=i+1\}} + 0.015\,\mathbf{1}_{\{j=i-2\}}$, and $D_n = \mathrm{diag}_{i=1,\ldots,n}\left((-1)^{i+1}0.025(0.92)^{i-1}\right)$. The continuous-time counterpart is SDE $dX_t = (A_n - I)X_t\,dt + \frac{1}{\sqrt{3}}D_n X_t\,dW_t$, using the same initial, unsafe, and target regions. Table III reports the corresponding synthesis times, showing that the

TABLE III: Scalability results for the coupled stochastic sink family.

Benchmark	dim	tolerant safety (discrete)			tolerant reach-avoid (discrete)			tolerant safety (continuous)			tolerant reach-avoid (continuous)		
		q	time(s)	γ	q	time(s)	γ	η	time(s)	γ	η	time(s)	γ
coupledsink	6	0.80	0.200	3.634	0.80	0.178	3.073	0.80	0.221	4.363	0.80	0.191	4.262
	8	0.80	0.685	3.723	0.80	0.608	3.094	0.80	0.578	4.479	0.80	0.575	4.370
	10	0.80	3.080	3.775	0.80	2.402	3.103	0.80	2.062	4.535	0.80	2.301	4.425
	12	0.80	11.472	3.801	0.80	8.512	3.107	0.80	7.980	4.545	0.80	7.489	4.441
	14	0.80	49.982	3.822	0.80	37.811	3.114	0.80	37.637	4.573	0.80	35.571	4.470
	16	0.80	136.887	3.835	0.80	109.516	3.112	0.80	118.537	4.575	0.80	111.208	4.471

TABLE IV: Comparison with counter augmentation for discrete-time tolerant safety.

Benchmark	Method	time(s)	$\mathbb{P}_x(D_\infty \le k)$ lower bound				
			1	5	10	20	30
randomwalk2	Tolerant barrier	0.096	0.000	0.332	0.636	0.809	0.906
	Counter augmentation	0.821	0.561	0.774	0.894	0.984	0.998
arch [53]	Tolerant barrier	19.044	0.859	0.953	0.974	0.998	1.000
	Counter augmentation	124.689	0.971	0.982	0.995	0.999	1.000
stable [55]	Tolerant barrier	14.084	0.528	0.843	0.981	1.000	1.000
	Counter augmentation	169.281	0.916	0.975	0.998	1.000	1.000

proposed SOS-based method remains practical up to $n = 16$ across all four verification routines.

A further observation is that continuous-time verification is consistently much faster than discrete-time verification. The main reason is structural: in the discrete-time setting, the SOS constraints must encode the one-step expectation $E[B(f(x,\xi))]$, which requires composing the barrier with the dynamics and integrating over the stochastic disturbance, thereby increasing the effective polynomial degree and the size of the resulting semidefinite program. In contrast, the continuous-time setting relies on the infinitesimal generator $\mathcal{L}B$ obtained by differentiation, therefore producing substantially smaller SOS programs, which explains the pronounced runtime advantage of the continuous-time experiments. ◁

RQ3: How does tolerant-barrier synthesis compare with counter augmentation?

Answer: In general, counter augmentation fixes a tolerance budget K and synthesizes a budget-specific family $V_0, \dots, V_{K+1}$ on the augmented state space. Hence it can be tighter for a chosen budget, but it is not reusable when K changes. Table IV compares the bounds obtained from a single tolerant barrier, evaluated at each displayed budget, with counter-augmentation certificates synthesized separately for those budgets.

Counter augmentation is effective for fixed-budget discrete safety and may produce tighter certificates for a chosen K. However, its complexity grows with the budget, since both the number of component barriers and the constraints increase with K. Tolerant barriers instead synthesize one certificate on the original state space and provide budget-parametric tail bounds. Thus the two approaches are complementary: counter augmentation is useful for fixed-budget comparison, whereas tolerant barriers provide budget-parametric certificates on the original state space for the safety and reach-avoid settings considered in this paper. ◁

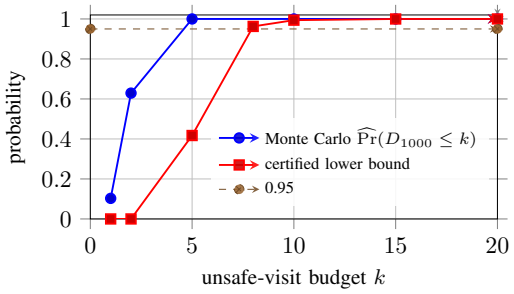

Fig. 3: Empirical CDF and certified lower bound for osc. The empirical 95% threshold is $k = 5$, while the certificate (with $q = 0.4$) gives $k_{0.95} = 8$.

TABLE V: Normalized unsafe exposure for reach-avoid benchmarks. The denominator is the empirical mean target hitting time from 10,000 simulations truncated at 2000 steps.

Benchmark	dim	certified $k_{0.95}$	$\widehat{\mathbb{E}}[T]$	$k_{0.95}/\widehat{\mathbb{E}}[T]$
vanderpol [50]	2	17	102.0	0.167
liederivative [51]	2	9	30.9	0.291
equil [19]	2	22	71.8	0.307

RQ4: How tight are the certified exposure bounds, and are the certified budgets small relative to execution time?

Answer: To assess the conservatism of the certified bounds in Tables I and II, we compare the certified 95% exposure thresholds with Monte Carlo estimates. For tolerant safety, the property is infinite-horizon, so there is no meaningful finite execution length by which to normalize the certified threshold. We therefore use finite-window simulations as an empirical diagnostic. In particular, for the representative discrete-time tolerant safety benchmark in Fig. 3, we plot the empirical CDF of D_{1000}, the number of unsafe visits during the first 1000 simulated steps. For reach-avoid, exposure is accumulated only until the first target hit, so the target hitting time provides the natural scale; accordingly, we report $k_{0.95}/\widehat{\mathbb{E}}[T]$ in Table V.

These comparisons show that the SOS-based certificates can be conservative, as expected from sufficient conditions, but they are not merely vacuous: on representative benchmarks, the certified thresholds remain on the same scale as the empirical unsafe-exposure distribution. The normalized reach-avoid results further indicate that the certified 95% unsafe-exposure budgets amount to only a modest fraction of the empirical target-hitting time. ◁

RQ5: How does the choice of the parameters q and η affect the tightness of the certified lower bounds and the

thresholds $k_{0.95}$ and $\ell_{0.95}$?

Answer: The experiments show that the parameter choice induces a clear trade-off in bound tightness. In the corner case $q = 1$ (or $\eta = 0$ in continuous time), the certificate is Markov-type, which typically yields a smaller certificate value γ and therefore tighter certified lower bounds for small unsafe-exposure budgets. By contrast, when $0 < q < 1$ (or $\eta > 0$), the certificate becomes exponential-type: the stronger decrease requirement inside the unsafe set often leads to a larger γ, so the certified lower bound can initially be looser at small budgets. This crossover is visible in `randomwalk1`, where the Markov barrier gives a better bound at $k = 10$, while the exponential barrier is worse in that low-budget regime.

For moderate and large budgets, however, the exponential-type case usually becomes superior because it yields geometric/exponential tail decay, whereas the Markov-type case provides only reciprocal decay. As a result, smaller q in discrete time, and analogously larger η in continuous time, generally lead to much tighter certified lower bounds and substantially smaller thresholds $k_{0.95}$ and $\ell_{0.95}$. For example, in `randomwalk1`, the threshold drops from $k_{0.95} = 189$ to 78 in discrete time, and from $\ell_{0.95} = 188.545$ to 4.414 in continuous time. ◁

VII. Conclusion

We presented tolerant safety and tolerant reach-avoid properties for both discrete-time stochastic systems and continuous-time stochastic differential dynamics. These properties generalize classical safety and reach-avoid by allowing bounded unsafe exposure, measured as the number of unsafe visits in discrete time and as unsafe occupation time in continuous time. We introduced tolerant barrier certificates that convert unsafe-exposure budgets into probabilistic tail bounds, and showed how to synthesize such certificates using SOS programming. Future work includes tighter bounds via higher-order moments and concentration inequalities, as well as extensions to stochastic hybrid systems.

References

[1] J. R. Norris, *Markov chains*. Cambridge university press, 1998, no. 2.

[2] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, ser. Wiley Series in Probability and Statistics. Wiley, 1994.

[3] D. Koller and N. Friedman, *Probabilistic Graphical Models - Principles and Techniques*. MIT Press, 2009.

[4] B. Xue, M. Fränzle, N. Zhan, S. Bogomolov, and B. Xia, "Safety verification for random ordinary differential equations," *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.*, vol. 39, no. 11, pp. 4090–4101, 2020.

[5] B. Øksendal, *Stochastic Differential Equations: An introduction with applications*. Springer Science & Business Media, 2013.

[6] G. Pola, M. Bujorianu, J. Lygeros, and M. D. D. Benedetto, "Stochastic hybrid models: An overview," in *ADHS*, vol. 36, no. 6. Elsevier, 2003, pp. 45–50.

[7] D. Bertsekas and S. E. Shreve, *Stochastic optimal control: The discrete-time case*. Athena Scientific, 1996, vol. 5.

[8] W. Paul and J. Baschnagel, *Stochastic Processes: From Physics to Finance*. Springer, 2013.

[9] J. M. Steele, *Stochastic calculus and financial applications*. Springer, 2001, vol. 1.

[10] L. J. Allen, *An introduction to stochastic processes with applications to biology*. CRC press, 2010.

[11] H. C. Tuckwell, *Stochastic processes in the neurosciences*. SIAM, 1989.

[12] S. Prajna and A. Jadbabaie, "Safety verification of hybrid systems using barrier certificates," in *HSCC*, ser. LNCS, vol. 2993. Springer, 2004, pp. 477–492.

[13] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control.*, vol. 62, no. 8, pp. 3861–3876, 2016.

[14] P. Wieland and F. Allgöwer, "Constructive safety using control barrier functions," *IFAC Proceedings Volumes*, vol. 40, no. 12, pp. 462–467, 2007.

[15] H. Kong, F. He, X. Song, W. N. N. Hung, and M. Gu, "Exponential-condition-based barrier certificate generation for safety verification of hybrid systems," in *CAV*, ser. LNCS, vol. 8044. Springer, 2013, pp. 242–257.

[16] X. Zeng, W. Lin, Z. Yang, X. Chen, and L. Wang, "Darboux-type barrier certificates for safety verification of nonlinear hybrid systems," in *EMSOFT*. ACM, 2016, pp. 1–10.

[17] A. Sogokon, K. Ghorbal, Y. K. Tan, and A. Platzer, "Vector barrier certificates and comparison systems," in *FM*, ser. LNCS, vol. 10951. Springer, 2018, pp. 418–437.

[18] Q. Wang, M. Chen, B. Xue, N. Zhan, and J.-P. Katoen, "Encoding inductive invariants as barrier certificates: Synthesis via difference-of-convex programming," *Information and Computation*, vol. 289, p. 104965, 2022.

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

[20] J. Steinhardt and R. Tedrake, "Finite-time regional verification of stochastic non-linear systems," *Int. J. Robotics Res.*, vol. 31, no. 7, pp. 901–923, 2012.

[21] P. Jagtap, S. Soudjani, and M. Zamani, "Formal synthesis of stochastic systems via control barrier certificates," *IEEE Trans. Autom. Control.*, vol. 66, no. 7, pp. 3097–3110, 2020.

[22] C. Santoyo, M. Dutreix, and S. Coogan, "A barrier function approach to finite-time stochastic system verification and control," *Automatica*, vol. 125, p. 109439, 2021.

[23] S. Feng, M. Chen, B. Xue, S. Sankaranarayanan, and N. Zhan, "Unbounded-time safety verification of stochastic differential dynamics," in *CAV (II)*, ser. LNCS, vol. 12225. Springer, 2020, pp. 327–348.

[24] A. Chakarov and S. Sankaranarayanan, "Probabilistic program analysis with martingales," in *CAV*, ser. LNCS, vol. 8044. Springer, 2013, pp. 511–526.

[25] K. Chatterjee, H. Fu, P. Novotný, and R. Hasheminezhad, "Algorithmic analysis of qualitative and quantitative termination problems for affine probabilistic programs," in *POPL*. ACM, 2016, pp. 327–342.

[26] K. Chatterjee, P. Novotný, and D. Zikelic, "Stochastic invariants for probabilistic termination," in *POPL*. ACM, 2017, pp. 145–160.

[27] T. Takisaka, Y. Oyabu, N. Urabe, and I. Hasuo, "Ranking and repulsing supermartingales for reachability in randomized programs," *ACM Trans. Program. Lang. Syst.*, vol. 43, no. 2, pp. 5:1–5:46, 2021.

[28] X. Chen, S. Chen, and V. M. Preciado, "Safety verification of nonlinear polynomial system via occupation measures," in *CDC*. IEEE, 2019, pp. 1159–1164.

[29] D. Henriques, J. G. Martins, P. Zuliani, A. Platzer, and E. M. Clarke, "Statistical model checking for Markov decision processes," in *QEST*. IEEE Computer Society, 2012, pp. 84–93.

[30] A. Legay, S. Sedwards, and L. Traonouez, "Scalable verification of Markov decision processes," in *SEFM Workshops*, ser. LNCS, vol. 8938. Springer, 2014, pp. 350–362.

[31] E. Bartocci, L. Kovács, and M. Stankovič, "Automatic generation of moment-based invariants for prob-solvable loops," in *International Symposium on Automated Technology for Verification and Analysis*. Springer, 2019, pp. 255–276.

[32] M. Moosbrugger, M. Stankovič, E. Bartocci, and L. Kovács, "This is the moment for probabilistic loops," *Proceedings of the ACM on Programming Languages*, vol. 6, no. OOPSLA2, pp. 1497–1525, 2022.

[33] 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.

[34] S. Summers and J. Lygeros, "Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem," *Autom.*, vol. 46, no. 12, pp. 1951–1961, 2010.

[35] M. Lahijanian, S. B. Andersson, and C. Belta, "Formal verification and synthesis for discrete-time stochastic systems," *IEEE Trans. Autom. Control.*, vol. 60, no. 8, pp. 2031–2045, 2015.

[36] C. Baier and J.-P. Katoen, *Principles of Model Checking*. MIT press, 2008.
[37] M. Z. Kwiatkowska, "Model checking for probability and time: From theory to practice," in *LICS*. IEEE Computer Society, 2003, p. 351.
[38] C. Baier, J. Klein, L. Leuschner, D. Parker, and S. Wunderlich, "Ensuring the reliability of your model checker: Interval iteration for Markov decision processes," in *CAV (II)*, ser. LNCS, vol. 10426. Springer, 2017, pp. 160–180.
[39] A. Hartmanns and B. L. Kaminski, "Optimistic value iteration," in *CAV (II)*, ser. LNCS, vol. 12225. Springer, 2020, pp. 488–511.
[40] T. Quatmann and J. Katoen, "Sound value iteration," in *CAV (I)*, ser. LNCS, vol. 10981. Springer, 2018, pp. 643–661.
[41] D. Williams, *Probability with Martingales*. Cambridge University Press, 1991.
[42] M. Anand, V. Murali, A. Trivedi, and M. Zamani, "k-inductive barrier certificates for stochastic systems," in *HSCC*. ACM, 2022, pp. 12:1–12:11.
[43] S. Prajna and A. Rantzer, "Convex programs for temporal verification of nonlinear dynamical systems," *SIAM J. Control. Optim.*, vol. 46, no. 3, pp. 999–1021, 2007.
[44] B. Xue, R. Li, N. Zhan, and M. Fränzle, "Reach-avoid analysis for stochastic discrete-time systems," in *ACC*. IEEE, 2021, pp. 4879–4885.
[45] R. Khasminskii, *Stochastic stability of differential equations*. Springer Science & Business Media, 2011, vol. 66.
[46] S. P. Meyn and R. L. Tweedie, *Markov chains and stochastic stability*. Springer Science & Business Media, 2012.
[47] B. Xue, N. Zhan, and M. Fränzle, "Reach-avoid analysis for polynomial stochastic differential equations," *IEEE Transactions on Automatic Control*, vol. 69, no. 3, pp. 1882–1889, 2023.
[48] P. A. Parrilo, "Semidefinite programming relaxations for semialgebraic problems," *Mathematical programming*, vol. 96, no. 2, pp. 293–320, 2003.
[49] L. Vandenberghe and S. Boyd, "Semidefinite programming," *SIAM review*, vol. 38, no. 1, pp. 49–95, 1996.
[50] B. Xue, N. Zhan, and M. Fränzle, "Reach-avoid analysis for stochastic differential equations," *arXiv preprint arXiv:2208.10752*, 2022.
[51] J. Liu, N. Zhan, and H. Zhao, "Computing semi-algebraic invariants for polynomial dynamical systems," in *EMSOFT'11*. ACM, 2011, pp. 97–106.
[52] M. Anand, V. Murali, A. Trivedi, and M. Zamani, "Safety verification of dynamical systems via k-inductive barrier certificates," in *CDC*. IEEE, 2021, pp. 1314–1320.
[53] A. Sogokon, K. Ghorbal, and T. T. Johnson, "Non-linear continuous systems for safety verification," in *ARCH@CPSWeek*, ser. EPiC Series in Computing, G. Frehse and M. Althoff, Eds., vol. 43, 2016, pp. 42–51.
[54] S. Ratschan and Z. She, "Providing a basin of attraction to a target region of polynomial systems by computation of lyapunov-like functions," *SIAM J. Control. Optim.*, vol. 48, no. 7, pp. 4377–4394, 2010.
[55] S. Sankaranarayanan, X. Chen, and E. Ábrahám, "Lyapunov function synthesis using handelman representations," in *NOLCOS*. IFAC, 2013, pp. 576–581.
[56] B. HomChaudhuri, M. Oishi, M. Shubert, M. Baldwin, and R. S. Erwin, "Computing reach-avoid sets for space vehicle docking under continuous thrust," in *CDC*. IEEE, 2016, pp. 3312–3318.
[57] M. Putinar, "Positive polynomials on compact semi-algebraic sets," *Indiana University Mathematics Journal*, vol. 42, no. 3, pp. 969–984, 1993.
[58] E. D. Andersen, C. Roos, and T. Terlaky, "On implementing a primal-dual interior-point method for conic quadratic optimization," *Math. Program.*, vol. 95, no. 2, pp. 249–277, 2003.

Shenghua Feng received his B.Sc degree in Mathematics from Nanjing University, Nanjing, China, in 2018, and a Ph.D degree in computer science from the Institute of Software, Chinese Academy of Sciences, Beijing, China, in 2023. From 2023 to 2025, he was an assistant researcher at Zhongguancun Laboratory, Beijing, China. Since then, he joined the Institute of Software, Chinese Academy of Sciences. His current research interests include formal methods, embedded and hybrid systems, probabilistic systems, and program verification.

Han Su received his B.Sc. degree in Control Science and Engineering from Shandong University, Jinan, China, in 2019, and his Ph.D degree in Computer Science from the Institute of Software, Chinese Academy of Sciences, in 2025.

Since then, he has been a Postdoctoral Researcher at the National Institute of Informatics (NII), Tokyo, Japan. His current research interests include cyber-physical systems, control synthesis and optimization, and runtime verification and enforcement.

Hao Wu received his B.Sc and Ph.D degrees both in computer science from the University of Chinese Academy of Sciences, Beijing, China, in 2020 and 2026, respectively. Since then, he joined the Institute of Science and Technology Austria (ISTA) as a post-doctoral researcher. His current research interests span a wide range of topics in theoretical computer science, including formal methods, complexity and computability.

Jie An received his Ph.D. degree in software engineering from Tongji University, Shanghai, China, in 2020. From 2020 to 2022, he worked as a postdoctoral researcher at the Max Planck Institute for Software Systems (MPI-SWS), Kaiserslautern, Germany. From 2022 to 2024, he was a project researcher and later a project assistant professor at the National Institute of Informatics (NII), Tokyo, Japan. Since June 2024, he joined the Institute of Software, Chinese Academy of Sciences, as an associate research professor. His research interests include real-time, embedded and hybrid systems, formal verification and synthesis, and artificial intelligence.

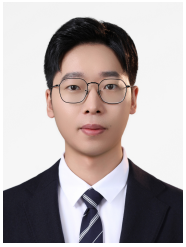

Mingshuai Chen received his Ph.D. degree in computer science from the Institute of Software, Chinese Academy of Sciences, Beijing, China, in 2019. He then worked as a postdoctoral researcher at RWTH Aachen University in Aachen, Germany. Since 2023, he joined the College of Computer Science and Technology at Zhejiang University as an assistant professor (ZJU100 Young Professor). His research interest lies in formal verification and synthesis, broadly construed in mathematical logic and theoretical computer science.

Naijun Zhan is a Boya distinguished professor in the School of Computer Science of Peking University. He got his PhD from the Institute of Software, Chinese Academy of Sciences (ISCAS). His research interests cover formal design of real-time, embedded and hybrid systems, and program verification. He is on the editorial boards of Journal of Automated Reasoning, Formal Aspects of Computing, Journal of Logical and Algebraic Methods in Programming, etc., a member of the steering committees of SETTA and MEMOCODE, the pc co-chairs of TACAS 2027, ICFEM 2025, FM 2021, and so on.