2022 American Control Conference (ACC)
Atlanta, USA, June 8-10, 2022

Myopically Verifiable Probabilistic Certificate for Long-term Safety

Zhuoyuan Wang[1†], Haoming Jing[1†], Christian Kurniawan[1], Albert Chern[2], Yorie Nakahira[1*]

Abstract— **In this paper, we consider using barrier function-based approaches for the safe control problem in stochastic systems. In the presence of stochastic uncertainties, a myopic controller that ensures safe probability in infinitesimal time intervals may suffer from the accumulation of unsafe probability over time and result in a small long-term safe probability. Meanwhile, increasing the outlook time horizon may lead to significant computation burdens and delayed reactions, which also compromises safety. To tackle this challenge, we define a new notion of forward invariance on 'probability space' as opposed to the safe regions on state space. This new notion allows the long-term safe probability to be framed into a forward invariance condition, which can be efficiently evaluated. We use this safety condition to propose a controller that evaluates infinitesimal outlook horizon and guarantees long-term safe probability or fast recovery probability. The proposed controller ensures the safe probability does not decrease over time or informs the exposed levels of risks (unsafe probability) when it becomes infeasible. The performance of the proposed controller is evaluated in numerical simulations. Finally, we show that this framework can also be adapted to characterize the speed and probability of forward convergent behaviors, which can be of use to finite-time Lyapunov analysis in stochastic systems.**

I. Introduction

Autonomous systems (*e.g.,* robots and self-driving vehicles) must make safe control decisions in real-time and in the presence of various uncertainties. The control of such safety- and delay-critical systems rely extensively on barrier function-based approaches. For deterministic or worst-case systems that possess small and bounded noise, barrier function-based approaches can provide provable safety with low computational cost [1], [2], [3]. Its computation efficiency mainly arises from two features: computation efficiency arising from a myopic controller (feature 1) and from the use of analytical/affine safety conditions (feature 2). However, these two features did not necessarily translate to stochastic systems whose uncertainty is captured by random variables with unbounded support, as we will discuss below. This paper overcomes this difficulty by characterizing a sufficient condition for 'invariance' in the probability space. This condition guarantees that the unsafe probability stays below a tolerable level while preserving the above two features, as detailed below.

Feature 1: Computation efficiency arising from a myopic controller. In a deterministic system, safety can be guaranteed if the state never moves outside the safe set within an infinitesimal outlook time interval. This property allows a myopic controller, which only evaluates the infinitesimal outlook time interval (immediate future time), to keep the system safe at all times. A myopic evaluation requires much less computation than methods that evaluate a long time horizon because the computational load to evaluate possible future trajectories significantly increases with the outlook time horizon.

In a stochastic system whose uncertainty has unbounded support, however, the probability of staying within the safe set in the infinitesimal outlook time interval is strictly less than one. In other words, there will always be a non-zero tail probability to move outside of the safe set. This tail probability can accumulate over time and result in a small long-term safe probability. This problem persists even when stochastic systems are modeled in a worst-case framework, and the tail probability beyond the assumed size of uncertainties is sufficiently small. The lack of long-term safety guarantees suggests the need for a more refined *temporal* characterization of long-term safe/unsafe probabilities.

Feature 2: Computation efficiency arising from the use of analytical/affine safety conditions. In a deterministic system, the condition for the state to stay within the safe set in an infinitesimal time can be translated as requiring the vector field of the state to stay within the tangent cone of the safe set [4]. A sufficient condition of this requirement for affine control systems can be expressed using analytic inequalities that are affine in the control action and thereby can be integrated into quadratic programs (see [5] and references therein).

In a stochastic system, however, constraining the mean trajectory to satisfy this condition without bounding the higher moments has no control over the tail probability of the state moving outside of the safe region. This suggests the need for a more refined *spatial* characterization of unsafe behaviors and state distribution.

Therefore, ensuring safety in a stochastic system needs more refined temporal and spatial characterization of safe/unsafe behaviors during a long outlook time interval. However, refined *temporal* characterization requires tracing the long-term evolution of complex dynamics, environmental changes, control actions, as well as their couplings, and refined *spatial* characterization requires characterization of the state distribution, tails, and conditional value at risk.

*This work is supported by JST, PRESTO Grant Number JPMJPR2136, Japan, and by the CONIX Research Center, one of six centers in JUMP, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

†These authors contributed equally.

[1]Zhuoyuan Wang, Haoming Jing, Christian Kurniawan and Yorie Nakahira are with the Department of Electrical and Computer Engineering, Carnegie Mellon Universty, {zhuoyuaw,haomingj,ckurniaw,ynakahir}@andrew.cmu.edu.

[2]Albert Chern is with the Department of Computer Science and Engineering, University of California San Diego, alchern@ucsd.edu.

∗To whom correspondence should be addressed.

Both compromise the above two features and can impose a significant computational burden. Such heavy computation can even compromise safety due to slower response, despite the use of more optimized actions.

A. Related Work

Prior work has yielded diverse approaches for finer time/space characterization in stochastic systems. These approaches can be approximately classified into three main types based on their choice of tradeoffs: long-term safety with heavy computation (approach A), myopic safety with low computation (approach B), and long-term conservative safety with low computation (approach C). It shall be noted that all wrestle with the above-mentioned tradeoffs between longer-term safety vs. faster response (computational burden).

Approach A: long-term safety with heavy computation. There exists extensive literature that considers a long time horizon and/or the state distribution (or higher moments of the state distribution) at the expense of high computation costs. For example, various model predictive control (MPC) and chance-constrained optimization include safety constraints in a long time horizon (see [6], [7] and references therein). Reachability-based techniques use the characterization of reachable states over a finite/infinite time horizon to constrain the control action so that the state reaches or avoids certain regions [8]. Within barrier function-based approaches, the safety condition can be formulated as constraints on the control action that involve the conditional value-at-risk (CVaR) of the barrier function values [9]. While these techniques can find more optimal control actions that are safe in the long term, they often come with significant computation costs. The cause is twofold: first, possible trajectories often scale exponentially with the length of the outlook time horizon; and second, tails or CVaR involve the probability and mean of rare events, which are more challenging to estimate than nominal events. Such stringent tradeoffs between estimating longer-term safe probability vs. computation burden limit the utility of these techniques in delay-critical systems for more expansive (longer time scale or precise characterization of the state distribution) control action evaluation.

Approach B: myopic safety with low computation.: Motivated by the latency requirement in real-time safety-critical control, a few approaches use myopic controllers that constrain the probability of unsafe events in an infinitesimal time interval. For example, the stochastic control barrier function use a sufficient condition for ensuring that the state, on average, moves within the tangent cone of the safe set [10]. The probabilistic barrier certificate ensures certain conditions of the barrier functions to be satisfied with high probability [11], [12]. The myopic nature of these methods achieves a significant reduction in computational cost but can result in unsafe behaviors in a longer time horizon due to the accumulation of tail probabilities of unsafe events.

Approach C: long-term conservative safety with low computation.: To have a faster response but still achieve longer-term safety, other approaches use probability and/or martingale inequalities to derive sufficient conditions for constraining the evolution of barrier function values in a given time interval [3], [13], [14]. These sufficient conditions are given analytically and are elegantly integrated into the convex optimization problems to synthesize controllers offline or verify control actions online. The controllers based on these techniques often require less online computation to find the action that guarantees longer-term safety. However, due to the approximate nature of the probabilistic inequalities, the control actions can be conservative and unnecessarily compromise nominal performance.

Contribution of this paper

This paper proposes an efficient algorithm that ensures safety during a fixed or receding time horizon. The algorithm is based on a new safety condition that is sufficient to control the unsafe probability in a given time interval to stay above the tolerable risk levels.[1] This safety condition is constructed by translating probabilistic safety specifications into a forward invariance condition on the level sets of the safe probability. The use of forward invariance allows safety at all time points to be guaranteed by a myopic controller that only evaluates the state evolution in an infinitesimal future time interval. For affine control systems, the condition is affine to the control action and thus can be evaluated using convex/quadratic programs. Our framework may be useful in characterizing the speed and probability of forward convergence in finite-time Lyapunov analysis of stochastic systems.

Below, we summarize the advantages of the proposed algorithms.

Advantage 1: Computation efficiency. The proposed method only myopically evaluates the immediate future using closed-form safety constraints. Thus, it can have reduced computational burdens than approach A.

Advantage 2: Provable guarantee in long-term safe probability. The closed-form safety constraints are derived from the safe probability during a receding or fixed time horizon. Thus, the proposed method can have more direct control over the probability of accumulating tail events than approach B.

Advantage 3: Intuitive parameter tuning using exact safety vs. performance tradeoffs. The proposed method uses exact characterizations of safe probability. Thus, it allows the aggressiveness towards safety to be directly tuned based on the exact probability, as opposed to probabilistic bounds or martingale approximations used in approach C. Moreover, when the condition becomes infeasible, one can deduce the exposed risk levels (unsafe probability). This property contrasts with the cases in a deterministic framework in that risk levels after a violation of the safety condition may be unknown and can largely depend on the choice of barrier functions.[2]

[1]Here, we consider two types of unsafe probability: the probability of exiting the safe set in a time interval when originated inside and the probability of recovering to the safe set when originated outside.

[2]This is because the barrier functions may not necessarily have physical or operational meanings.

II. Preliminary

Let $\mathbb{R}$, $\mathbb{R}_+$, $\mathbb{R}^n$, and $\mathbb{R}^{m\times n}$ be the set of real numbers, the set of non-negative real numbers, the set of n-dimensional real vectors, and the set of $m \times n$ real matrices, respectively. Let $x[k]$ be the k-th element of vector x. Let $f : \mathcal{X} \to \mathcal{Y}$ represent that f is a mapping from space $\mathcal{X}$ to space $\mathcal{Y}$. Let $\mathbb{1}\{\mathcal{E}\}$ be an indicator function, which takes 1 when condition $\mathcal{E}$ holds and 0 otherwise. Let $\mathbf{0}_{m\times n}$ be an $m \times n$ matrix with all entries 0. Given events $\mathcal{E}$ and $\mathcal{E}_c$, let $\mathbb{P}(\mathcal{E})$ be the probability of $\mathcal{E}$ and $\mathbb{P}(\mathcal{E}|\mathcal{E}_c)$ be the conditional probability of $\mathcal{E}$ given the occurrence of $\mathcal{E}_c$. Given random variables X and Y, let $\mathbb{E}[X]$ be the expectation of X and $\mathbb{E}[X|Y = y]$ be the conditional expectation of X given $Y = y$. We use uppercase letters (*e.g.,* Y) to denote random variables and lowercase letters (*e.g.,* y) to denote their specific realizations. The Hessian of a scalar valued function $\phi(x)$ is denoted as $\operatorname{Hess}\phi(x)$.

Definition 1 (Infinitesimal Generator)**.** The infinitesimal generator A of a stochastic process $\{Y_t \in \mathbb{R}^n\}_{t\in\mathbb{R}_+}$ is

$$AF(y) = \lim_{h\to 0} \frac{\mathbb{E}[F(Y_h)|Y_0 = y] - F(y)}{h} \tag{1}$$

whose domain is the set of all functions $F : \mathbb{R}^n \to \mathbb{R}$ such that the limit of (1) exists for all $y \in \mathbb{R}^n$.

III. Problem Statement

Here, we introduce the control system in subsection III-A, characterize two types of safety in subsection III-B, state the controller design goals in subsection III-C, and outline the control policy in subsection III-D.

A. Control System Description

We consider the following time-invariant stochastic affine control system:

$$dX_t = (f(X_t) + g(X_t)U_t)\, dt + \sigma(X_t)dW_t, \tag{2}$$

where $X_t \in \mathbb{R}^n$ is the system state, $U_t \in \mathbb{R}^m$ is the control input, and $W_t \in \mathbb{R}^\omega$ captures the system uncertainties. Here, X_t can include both the controllable states of the system and the uncontrollable environmental variables such as moving obstacles. We assume that W_t is the standard Wiener process with 0 initial value, *i.e.,* $W_0 = 0, W_t \sim \mathcal{N}(0, t)$. The value of $\sigma(X_t)$ is determined based on the size of uncertainty in the state, unmodeled dynamics, and environmental variables. The control action U_t is determined at each time by the control policy introduced in subsection III-D. We assume that accurate information of the system state can be used for control.

The safe region of the state is specified by the zero super level set of a continuous barrier function $\phi(x) : \mathbb{R}^n \to \mathbb{R}$, *i.e.,*

$$\mathcal{C}(0) = \{x \in \mathbb{R}^n : \phi(x) \geq 0\}. \tag{3}$$

We use

$$\mathcal{C}(L) := \{x \in \mathbb{R}^n : \phi(x) \geq L\} \tag{4}$$

to denote the set with safety margin L. Accordingly, we use $\operatorname{int}\mathcal{C}(0) = \{x \in \mathbb{R}^n : \phi(x) > 0\}$ to denote the interior of the safe set, $\mathcal{C}(0)^c = \{x \in \mathbb{R}^n : \phi(x) < 0\}$ to denote the unsafe set, $\partial\mathcal{C}(L) = \{x \in \mathbb{R}^n : \phi(x) = L\}$ to denote the boundary of the L super level set.

B. Probabilistic Characterization of Safe Behaviours

The system must satisfy two types of probabilistic safety specifications: forward invariance (type I) and forward convergence (type II).

1) Forward Invariance (type I): The forward invariance property refers to the system's ability to keep its state within a set when the state originated from the set. The probabilistic forward invariance to a set $\mathcal{C}(L)$ can be quantified using

$$\mathbb{P}\left(\, X_\tau \in \mathcal{C}(L), \forall\tau \in [t, t+T] \mid X_t = x \,\right) \tag{5}$$

for some time interval $[t, t + T]$ conditioned on an initial condition $X_t = x \in \mathcal{C}(L)$. Probability (5) can be computed from the distribution of the following two random variables:[3]

$$\Phi_x(T) := \inf\{\phi(X_t) \in \mathbb{R} : t \in [0, T], X_0 = x\}, \tag{6}$$
$$\Gamma_x(L) := \inf\{t \in \mathbb{R}_+ : \phi(X_t) < L, X_0 = x\}. \tag{7}$$

Here, $\Phi_x(T)$ is the worst-case safety margin from the boundary of the safe set $\partial\mathcal{C}(0)$ during $[0, T]$, and $\Gamma_x(L)$ is the time when the system exits from $\mathcal{C}(L)$ for the first time. We can rewrite (5) using the two random variables (6) and (7) as

$$\begin{aligned}
&\mathbb{P}\left(\, X_\tau \in \mathcal{C}(L), \forall\tau \in [t, t+T] \mid X_t = x \,\right) && (8)\\
&= \mathbb{P}\left(X_\tau \in \mathcal{C}(L), \forall\tau \in [0, T] \mid X_0 = x\right) && (9)\\
&= \mathbb{P}(\Phi_x(T) \geq L) && (10)\\
&= \mathbb{P}(\Gamma_x(L) > T) = 1 - \mathbb{P}(\Gamma_x(L) \leq T). && (11)
\end{aligned}$$

Here, equality (9) holds due to the time-invariant nature of the system and control policies.

2) Forward Convergence (type II): The forward convergence property indicates the system's capability for its state to enter a set when the state originated from outside the set. This probabilistic forward convergence can be quantified using

$$\mathbb{P}\left(\, \exists\tau \in [t, t+T] \text{ s.t. } X_\tau \in \mathcal{C}(L) \mid X_t = x \,\right) \tag{12}$$

for some time interval $[t, t + T]$ conditioned on an initial condition $X_t = x \in \mathcal{C}(L)^c$. Similar to the case of forward invariance, probability (12) can also be computed from the distribution of the following two random variables:[3]

$$\Theta_x(T) := \sup\{\phi(X_t) \in \mathbb{R} : t \in [0, T], X_0 = x\}, \tag{13}$$
$$\Psi_x(L) := \inf\{t \in \mathbb{R}_+ : \phi(X_t) \geq L, X_0 = x\}. \tag{14}$$

Here, $\Theta_x(T)$ indicates the distance to the boundary of the safe set $\partial\mathcal{C}(0)$, and $\Psi_x(L)$ is the duration for the state to

[3] These random variables are previously introduced and analyzed in [15].

enter the set $\mathcal{C}(L)$ for the first time. We can also rewrite (12) using the two random variables (13) and (14) as

$$\mathbb{P}\left(\exists \tau \in [t, t+T] \text{ s.t. } X_\tau \in \mathcal{C}(L) \mid X_t = x\right) \tag{15}$$
$$= \mathbb{P}\left(\exists \tau \in [0, T] \text{ s.t. } X_\tau \in \mathcal{C}(L) \mid X_0 = x\right) \tag{16}$$
$$= \mathbb{P}(\Theta_x(T) \geq L) \tag{17}$$
$$= \mathbb{P}(\Psi_x(L) \leq T). \tag{18}$$

C. Design goals

The design goal is to ensure long-term safety guarantees given as probabilistic forward invariance or convergence conditions. When the goal is to guarantee probabilistic forward invariance (type I), we aim to ensure the following condition: for each time $t \in \mathbb{R}_+$,

$$\mathbb{P}\big(X_\tau \in \mathcal{C}(L_t), \forall \tau \in [t, t+T_t]\big) \geq 1 - \epsilon, \tag{19}$$

for some $\epsilon \in (0, 1)$. From now on, all probabilities are conditioned on the initial condition $X_0 = x$ unless otherwise noted. Here, L_t is the desired safety margin, and T_t is the outlook time horizon. For each time t, condition (19) constrains the probability of staying within the safe set with margin L_t during the time interval $[t, t+T_t]$ to be above $1-\epsilon$.

When the goal is to guarantee probabilistic forward convergence (type II), we aim to ensure the following condition: for each time $t \in \mathbb{R}_+$,

$$\mathbb{P}\big(\exists \tau \in [t, t+T_t] \text{ s.t. } X_\tau \in \mathcal{C}(L_t)\big) \geq 1 - \epsilon, \tag{20}$$

for some $\epsilon \in (0, 1)$. Here, T_t is the convergence time and L_t is the convergence accuracy.

In both cases, the value of $\epsilon \in (0, 1)$ is chosen based on risk tolerance. In (19) and (20), the probabilities are taken over the distribution of X_t and its future trajectories $\{X_\tau\}_{\tau \in (t, t+T_t]}$ conditioned on $X_0 = x$. The distribution of X_t is generated based on the closed-loop system of (2) and (24), which are defined in subsection III-D. The distribution of $\{X_t\}_{t \in (t, t+T_t]}$ are allowed to be defined by a probability measure generated by other closed (or open) loop systems. For example, one can use the measure defined by the closed-loop system involving (2) and (23), the measure defined by optimal control actions with respect to some objective functions, or the measure defined by reachability methods [16], [17]. The scope of this paper is to present a general framework. We leave the computational techniques of any specific choices of measures as future work.

We consider either fixed time horizon or receding time horizon. In the fixed time horizon, safety is evaluated at each time t for a time interval $[t, t+H]$ of fixed length H. In the receding time horizon, we evaluate, at each time t, safety only for the remaining time $[t, H]$. The outlook time horizon for each case is given by

$$T_t = \begin{cases} H, & \text{for fixed time horizon,} \\ H - t, & \text{for receding time horizon.} \end{cases} \tag{21}$$

The safety margin is assumed to be either fixed or time varying. Fixed margin refers to when the margin remains constant at all time, *i.e.,* $L_t = \ell$. For time-varying margin, we consider the margin L_t that evolves according to

$$dL_t = f_\ell(L_t),\ L_0 = \ell, \tag{22}$$

for some continuously differentiable function f_ℓ.[4] The values of T_t and $\{L_t\}_{t \in [0, \infty)}$ are determined based on the design choice.

D. Control Policy

The control policy is composed of a nominal controller and additional modification scheme to ensure the safety specifications illustrated in subsections III-B and III-C. The nominal controller is represented by

$$U_t = N(X_t), \tag{23}$$

which can be arbitrary depending on the original control task and does not necessarily account for the safety specifications defined below. To adhere to the safety specifications, the output of the nominal controller is then modified by another scheme. The overall control policy involving the nominal controller and the modification scheme is represented by

$$U_t = K_N(X_t, L_t, T_t), \tag{24}$$

where $K_N : \mathbb{R}^n \times \mathbb{R} \times \mathbb{R} \to \mathbb{R}^m$ is a deterministic function of the current state X_t, safety margin L_t, and time horizon T_t to the current control action U_t. In K_N, the subscript N represents its dependence on the nominal controller N. The policy of the form (24) assumes that the decision rule is time-invariant,[5] and that the control action can be uniquely determined for each (X_t, L_t, T_t). This policy is also assumed to be memory-less in the sense that it does not use the past history of the state $\{X_\tau\}_{\tau < t}$ to produce the control action U_t. The assumption for memory-less controller is reasonable because the state evolution dX_t of system (2) only depends on the current system state X_t.[6] We restrict ourselves to the settings when f, g, σ, N, and K_N have sufficient regularity conditions such that both the closed loop system involving (2) and (23) and the system involving (2) and (24) have unique strong solutions.[7]

IV. Proposed Method

Here, we present a sufficient condition to achieve the safety requirements in subsection IV-A. Based on this condition, we propose a safe controller in subsection IV-B.

Before presenting these results, we first define a few notations. To capture the time-varying nature of T_t and L_t, we augment the state space as

$$Z_t := [T_t, L_t, X_t^\intercal]^\intercal \in \mathbb{R}^{n+2}. \tag{25}$$

The dynamics of Z_t satisfies the following SDE:

$$dZ_t = (\tilde{f}(Z_t) + \tilde{g}(Z_t)U_t)dt + \tilde{\sigma}(Z_t)dW_t. \tag{26}$$

[4] This representation also captures fixed margin by setting $f_\ell(L_t) \equiv 0$.

[5] The mappings N, K_N do not change over time.

[6] Note that $f(X_t)$, $g(X_t)$, and $\sigma(X_t)$ are time-invariant functions of the system state.

[7] Conditions required to have a unique strong solution can be found in [18, Chapter 5], [19, Chapter 1], [20, Chapter II.7] and references therein.

Here, $\tilde{f}$, $\tilde{g}$, and $\tilde{\sigma}$ are defined to be

$$\tilde{f}(Z_t) := \begin{bmatrix} f_T \\ f_\ell(L_t) \\ f(X_t) \end{bmatrix} \in \mathbb{R}^{(n+2)}, \tag{27}$$

$$\tilde{g}(Z_t) := \begin{bmatrix} \mathbf{0}_{2\times n} \\ g(X_t) \end{bmatrix} \in \mathbb{R}^{(n+2)\times m}, \tag{28}$$

$$\tilde{\sigma}(Z_t) := \begin{bmatrix} \mathbf{0}_{2\times n} \\ \sigma(X_t) \end{bmatrix} \in \mathbb{R}^{(n+2)\times \omega}. \tag{29}$$

In (27), the scalar f_T is given by

$$f_T := \begin{cases} 0, & \text{in fixed time horizon,} \\ -1, & \text{in receding time horizon,} \end{cases} \tag{30}$$

and the function f_ℓ is given by (22).

Remark 1. The Lie derivative of a function $\phi(x)$ along the vector field $f(x)$ is denoted as $\mathcal{L}_f\phi(x) = f(x)\cdot\nabla\phi(x)$. The Lie derivative $(\mathcal{L}_g\phi(x))$ along a matrix field $g(x)$ is interpreted as a row vector such that $(\mathcal{L}_g\phi(x))\,u = (g(x)u)\cdot\nabla\phi(x)$.

A. Conditions to Assure Safety

When safety is specified using forward invariance condition (type I) and is given in (19), we define

$$\mathbf{F}(Z_t) := \mathbb{P}\left(\Phi_{X_t}(T_t) \geq L_t\right) = \mathbb{P}\left(\Gamma_{X_t}(L_t) > T_t\right). \tag{31}$$

Otherwise, when safety is specified using forward convergence condition (type II) and is given in (20), we define

$$\mathbf{F}(Z_t) := \mathbb{P}\left(\Theta_{X_t}(T_t) \geq L_t\right) = \mathbb{P}\left(\Psi_{X_t}(L_t) \leq T_t\right). \tag{32}$$

The probability is taken over the same distributions of $\{X_\tau\}_{\tau\in[t,T_t]}$ that are used in the safety requirement (19) and (20). The values of T_t and L_t (known and deterministic) are defined in (21) and (22) depending on the design choice of receding/fixed time-horizon and fixed/varying margin.

Additionally, we assume that f, g, σ, and ϕ are chosen such that $A\mathbf{F}(Z_t)$ exists for all $Z \in \mathbb{R}^{n+2}$ and $U \in \mathbb{R}^m$. When the control action U_t is given by the $U_t = u$, from Itô's Lemma,[8] its value can be computed as

$$\begin{aligned} A\mathbf{F}(Z_t) := {} & \mathcal{L}_{\tilde{f}}\mathbf{F}(Z_t) + \left(\mathcal{L}_{\tilde{g}}\mathbf{F}(Z_t)\right)u \\ & + \frac{1}{2}\mathrm{tr}\left(\left[\tilde{\sigma}(Z_t)\right]^\intercal \mathrm{Hess}\,\mathbf{F}(Z_t)\left[\tilde{\sigma}(Z_t)\right]\right). \end{aligned} \tag{33}$$

We propose to constrain the control action U_t to satisfy the following condition at each time t:

$$A\mathbf{F}(Z_t) \geq -\alpha\left(\mathbf{F}(Z_t) - (1-\epsilon)\right). \tag{34}$$

Here, $\alpha : \mathbb{R} \to \mathbb{R}$ is assumed to be a monotonically-increasing, strictly concave or linear function that satisfies $\alpha(0) \leq 0$. From (33), condition (34) is affine in U_t. This property allows us to integrate condition (34) into a convex/quadratic program.

[8]Itô's Lemma is stated as below: Given a n-dimensional real valued diffusion process $dX = \mu dt + \sigma dW$ and any twice differentiable scalar function $f : \mathbb{R}^n \to \mathbb{R}$, one has $df = \left(\mathcal{L}_\mu f + \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^\intercal\,\mathrm{Hess}\,f\right)\right)dt + \mathcal{L}_\mu\sigma dW$.

The existing safety conditions in deterministic systems are often designed to find control actions so that the vector field of the state does not point outside of the safe set around its boundary. In other words, the value of the barrier function will be non-decreasing in the infinitesimal future outlook time horizon whenever the state is close to the boundary of the safe set. However, such myopic decision-making may not account for the fact that different directions of the tangent cone of the safe set may lead to vastly different long-term safe probability. In contrast, the proposed condition (34) nests the long-term safe probability in $\mathbf{F}$, and are guaranteed to steer the state toward the direction with non-decreasing long-term safe probability when the tolerable long-term unsafe probability is about to be violated.

Assumption 1. *The mappings f, g, σ, N, and K_N have sufficient regularity conditions, such that $\mathbf{F}(z)$ in* (31) *or* (32) *is a continuously differentiable function of $z \in \mathbb{R}^{n+2}$ and $\mathbb{E}[\mathbf{F}(Z_t)]$ is differentiable in t.*

Lemma 1. *Consider the closed-loop system of* (2) *and* (24). *Assume that Assumption 1 holds. If system* (2) *originates at $X_0 = x$ with $\mathbf{F}(z) > 1-\epsilon$, and the control action satisfies* (34) *at all time, then the following condition holds:*[9]

$$\mathbb{E}\left[\mathbf{F}(Z_t)\right] \geq 1-\epsilon \tag{35}$$

for all time $t \in \mathbb{R}_+$.

Theorem 1. *Consider the closed-loop system of* (2) *and* (24). *Assume that Assumption 1 holds. Let $\mathbf{F}$ be defined as type I in* (31). *If the system state originates at $X_0 = x$ with $\mathbf{F}(z) > 1-\epsilon$, and the control action satisfies* (34) *at all time $t \in \mathbb{R}_+$, then condition* (19) *holds.*

Theorem 2. *Consider the closed-loop system of* (2) *and* (24). *Assume that Assumption 1 holds. Let $\mathbf{F}$ be defined as type II in* (32). *If the system state originates at $X_0 = x$ with $\mathbf{F}(z) > 1-\epsilon$, and the control action satisfies* (34) *at all time $t \in \mathbb{R}_+$, then condition* (20) *holds.*

The proofs of Lemma 1, Theorem 1, and Theorem 2 are given in the extended version of this paper [21].

B. Safe Controller

Here, we propose a safe controller based on the safety conditions introduced in subsection IV-A. In this controller, the value of $\mathbf{F}$ is defined as type I in (31) when the safety specification is given as forward invariance condition, and as type II in (32) when the safety specification is given as forward convergence condition.

We frame the safe controller as a constrained optimization of the form

$$\begin{aligned} K_N(X_t, L_t, T_t) = \quad & \underset{u}{\arg\min} \quad J(N(X_t), u) \\ & \text{s.t.} \quad (34). \end{aligned} \tag{36}$$

[9]Here, the expectation is taken over X_t conditioned on $X_0 = x$, and $\mathbf{F}$ in (31) or (32) gives the probability of forward invariance/convergence of the future trajectories $\{X_\tau\}_{(t,t+T_t]}$ starting at X_t.

Here, $J : \mathbb{R}^m \times \mathbb{R}^m \to \mathbb{R}$ is an objective function that penalizes the deviation from the desired performance, the nominal control action, and/or the costs. It is also designed to satisfy the assumptions of lemma 1 to comply with the safety specification (19) or (20). The constraint of (36) imposes that (34) holds at all time t, and can additionally capture other design restrictions in its objective function and constraints.[10] When $(\mathcal{L}_{\tilde{g}}\mathbf{F}(z)) \neq 0$ for any z, there always exists u that satisfy the constraint (34).

When $\mathbf{F}$ is defined based on the closed-loop system involving (2) and (23), its value can be computed offline. Given the offline evaluation, the controller only needs to myopically evaluate the closed-form inequality conditions (36) in real time execution. In such cases, its computation load is comparable to common myopic barrier function-based methods in a deterministic system.

V. Example Use Case

In this section, we show the efficacy of our proposed method in an example use case. The source code for simulation is available upon request.

A. Setting

We consider the control affine system (2) with $f(X_t) \equiv AX_t = 2X_t$, $g(X_t) \equiv 1$, $\sigma(X_t) \equiv 2$. The safe set is defined in (3) and the barrier function is chosen to be $\phi(x) := x - 1$. The safety specification is given as the forward invariance condition. The nominal controller is a proportional controller $N(X_t) = -KX_t$ with $K = 2.5$. The closed-loop system with this controller has an equilibrium at $x = 0$ and tends to move into the unsafe set in the state space. We run simulations with $dt = 0.1$ for all controllers. The initial state is set to be $x_0 = 3$.

We compare our proposed controller with three existing safe controllers designed for stochastic systems. Below, we present their simplified versions.

- Proposed controller: The safety condition is given by (34). We choose type I in (31) with fixed time horizon and time-invariant zero margin, *i.e.,* $\mathbb{P}\left(\Phi_{X_t}(H) \geq 0\right)$, with $H = 10$. The value of $A\mathbf{F}$ from (33) is acquired by sampling 10000 trajectories and applying the finite difference method.
- Stochastic control barrier functions (StoCBF) [10]: The safety condition is given by

$$A\phi(X_t) \geq -\alpha\phi(X_t), \tag{37}$$

  where $\alpha > 0$ is a constant. This condition constrains the average system state to move within the tangent cone of the safe set.
- Probabilistic safety barrier certificates (PrSBC) [11]: The safety condition is given by

$$\mathbb{P}\left(d\phi(X_t, U_t) + \alpha\phi(X_t) \geq 0\right) \geq 1 - \epsilon, \tag{38}$$

  where $\alpha > 0$ is a constant. This condition constrains the state to stay within the safe set in the infinitesimal future interval with high probability.

[10]For example, K_N is Lipschitz continuous when $J(N(x), u) = u^\intercal H(x)u$ with $H(x)$ being a positive definite matrix (pointwise in x).

- Conditional-value-at-risk barrier functions (CVaR) [9]: The safety condition is given by

$$\text{CVaR}_\beta\left(\phi(X_{t_{k+1}})\right) \geq \gamma\phi(X_{t_k}) \tag{39}$$

  where $\gamma \in (0, 1)$ is a constant, $\{t_0 = 0, t_1, t_2, \cdots\}$ is a discrete sampled time of equal sampling intervals. This is a sufficient condition to ensure the value of $\text{CVaR}^k_\beta(\phi(X_{t_k}))$ conditioned on $X_0 = x$ to be non-negative at all sampled time $t_{k\in\mathbb{Z}_+}$. The value of $\text{CVaR}^k_\beta(\phi(X_{t_k}))$ quantifies the evaluation made at time $t_0 = 0$ about safety at time t_k.

The parameter α in the proposed controller, StoCBF and PrSBC has a similar effect, and the parameter ϵ is the tolerable probability of unsafe events both in the proposed controller and PrSBC. Thus, we choosed the same values of α and ϵ for these algorithms for a fair comparison. Specifically, we use $\alpha = 1$ for all controllers except for CVaR, $\epsilon = 0.1$ for the proposed controller and PrSBC. We use $\gamma = 0.65$ and $\beta = 0.1$ for CVaR.

We consider the following two scenarios:

- **Worst-case safe control:** We use the controller that satisfies the safety condition with equality at all time to test the safety enforcement power of these safety constraint. Such control actions are the riskiest actions that are allowed by the safety condition. The use of such control actions allows us to evaluate the safety conditions separately from the impact of the nominal controllers.
- **Switching control:** We impose safe controller only when the nominal controller does not satisfy the safety constraint. We implement this by replacing the constraints in (36) with safety conditions for different controllers presented in section V-A.

B. Results

Fig. 1 shows the results in the worst-case setting. The proposed controller can keep the expected safe probability $\mathbb{E}[\mathbf{F}(X_t)]$ close to 0.9 all the time, while others fail to keep it at a high level. A major cause of such failures is due to the accumulation of rare event probability. This comparison shows the advantage of having a provable guarantee for non-decreasing long-term safe probability. For comparable parameters, the safety improves from StoCBF to PrSBC to CVaR. This is also expected as constraining the expectation has little control of higher moments, and constraining the tail is not as strong as constraining the tail and the mean values of the tail.

Fig. 2 shows the results in the switching control setting. We obtained the empirical safe probability by calculating the number of safe trajectories over the total trials. In this setting, the proposed controller can keep the state within the safe region with the highest probability compared to other methods, even when there is a nominal control that acts against safety criteria. This is because the proposed controller directly manipulates dynamically evolving state distributions to guarantee non-decreasing safe probability when the tolerable unsafe probability is about to be violated, as opposed to when the state is close to an unsafe region. Our novel use of forward invariance condition on the safe

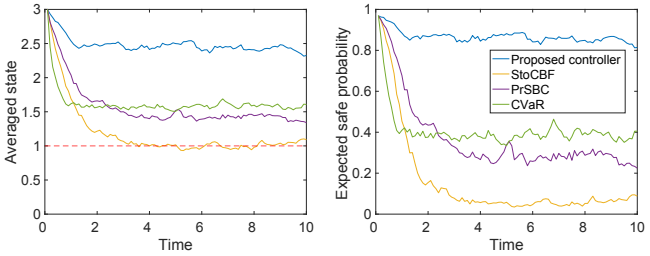

Fig. 1: Results in the worst-case setting. Left plot shows the average system state over 100 trajectories. Right plot shows the expected safe probability.

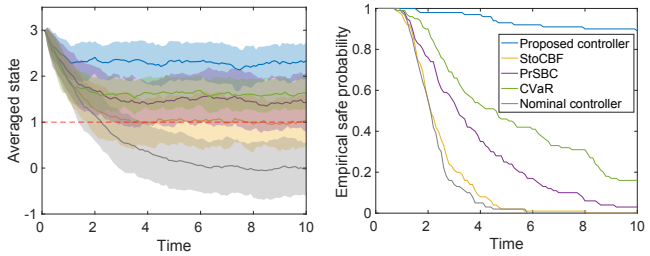

Fig. 2: Results in the switching control setting. Left plot shows the averaged system state of 100 trajectories with its standard deviation. Right plot shows the empirical safe probability.

probability allows a myopic controller to achieve long-term safe probability, which cannot be guaranteed by any myopic controller that directly imposes forward invariance on the safe set.

VI. Conclusion

In this paper, we considered the problem of ensuring long-term safety with high probability in stochastic systems. We proposed a sufficient condition to control the long-term safe probability of forward invariance (staying within the safe region) and forward convergence (recovering to the safe region). We then integrated the proposed sufficient condition into a computationally efficient myopic controller. Finally, we evaluated the performance of our proposed controller in a numerical example. Although beyond the scope of this paper, the proposed framework can also be used to characterize the speed and probability of system convergence and may be useful in finite-time Lyapunov analysis in stochastic systems.

References

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

[2] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," in *Autonomous robot vehicles*. Springer, 1986, pp. 396–404.

[3] 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.

[4] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.

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

[6] M. Farina, L. Giulioni, and R. Scattolini, "Stochastic linear model predictive control with chance constraints–a review," *Journal of Process Control*, vol. 44, pp. 53–67, 2016.

[7] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, "Learning-based model predictive control: Toward safe learning in control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 269–296, 2020.

[8] M. Chen and C. J. Tomlin, "Hamilton–jacobi reachability: Some recent theoretical advances and applications in unmanned airspace management," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 333–358, 2018.

[9] M. Ahmadi, X. Xiong, and A. D. Ames, "Risk-sensitive path planning via cvar barrier functions: Application to bipedal locomotion," *arXiv preprint arXiv:2011.01578*, 2020.

[10] A. Clark, "Control barrier functions for complete and incomplete information stochastic systems," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 2928–2935.

[11] W. Luo, W. Sun, and A. Kapoor, "Multi-robot collision avoidance under uncertainty with probabilistic safety barrier certificates," *arXiv preprint arXiv:1912.09957*, 2019.

[12] Y. Lyu, W. Luo, and J. M. Dolan, "Probabilistic safety-assured adaptive merging control for autonomous vehicles," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 10 764–10 770.

[13] S. Yaghoubi, K. Majd, G. Fainekos, T. Yamaguchi, D. Prokhorov, and B. Hoxha, "Risk-bounded control using stochastic barrier functions," *IEEE Control Systems Letters*, 2020.

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

[15] A. Chern, X. Wang, A. Iyer, and Y. Nakahira, "Safe control in the presence of stochastic uncertainties," *Accepted to 2021 60th Conference on Decision and Control*, 2021.

[16] J. J. Choi, D. Lee, K. Sreenath, C. J. Tomlin, and S. L. Herbert, "Robust control barrier-value functions for safety-critical control," *arXiv preprint arXiv:2104.02808*, 2021.

[17] J. F. Fisac, M. Chen, C. J. Tomlin, and S. S. Sastry, "Reach-avoid problems with time-varying dynamics, targets and constraints," in *Proceedings of the 18th international conference on hybrid systems: computation and control*, 2015, pp. 11–20.

[18] G. P. Moustris, S. C. Hiridis, K. M. Deliparaschos, and K. M. Konstantinidis, "Evolution of autonomous and semi-autonomous robotic surgical systems: a review of the literature," *The international journal of medical robotics and computer assisted surgery*, vol. 7, no. 4, pp. 375–392, 2011.

[19] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications*, 6th ed., ser. Universitext. Berlin Heidelberg: Springer-Verlag, 2003.

[20] A. N. Borodin, *Stochastic processes*. Springer, 2017.

[21] Z. Wang, H. Jing, C. Kurniawan, A. Chern, and Y. Nakahira, "Myopically verifiable probabilistic certificate for long-term safety," *arXiv preprint arXiv:2110.13380*, 2021.