Robotics: Science and Systems 2023
Daegu, Republic of Korea, July 10-July 14, 2023

Robust Safety under Stochastic Uncertainty with Discrete-Time Control Barrier Functions

Ryan K. Cosner[1], Preston Culbertson[1], Andrew J. Taylor[2], and Aaron D. Ames[1]

[1]Department of Mechanical and Civil Engineering, [2]Computing and Mathematical Sciences Department
California Institute of Technology, Pasadena, California 91125
{rkcosner, pculbert, ajtaylor, ames}@caltech.edu

***Abstract*—Robots deployed in unstructured, real-world environments operate under considerable uncertainty due to imperfect state estimates, model error, and disturbances. Given this real-world context, the goal of this paper is to develop controllers that are provably safe under uncertainties. To this end, we leverage Control Barrier Functions (CBFs) which guarantee that a robot remains in a "safe set" during its operation—yet CBFs (and their associated guarantees) are traditionally studied in the context of continuous-time, deterministic systems with bounded uncertainties. In this work, we study the safety properties of discrete-time CBFs (DTCBFs) for systems with discrete-time dynamics and unbounded stochastic disturbances. Using tools from martingale theory, we develop probabilistic bounds for the safety (over a finite time horizon) of systems whose dynamics satisfy the discrete-time barrier function condition in expectation, and analyze the effect of Jensen's inequality on DTCBF-based controllers. Finally, we present several examples of our method synthesizing safe control inputs for systems subject to significant process noise, including an inverted pendulum, a double integrator, and a quadruped locomoting on a narrow path.**

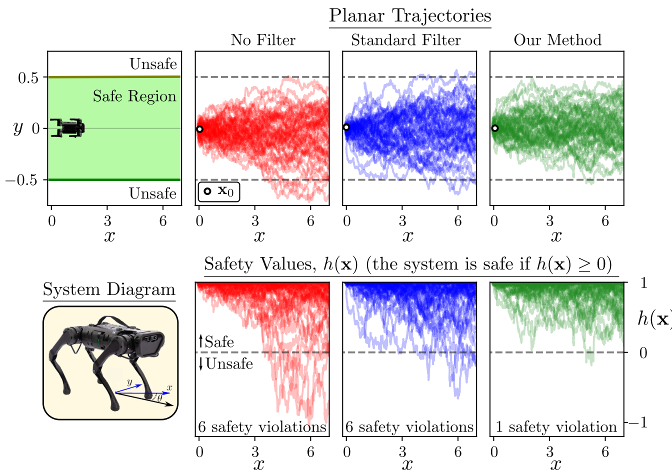

Fig. 1. Safety of a simulated quadrupedal robot locomoting on a narrow path for a variety of controllers. **(Top Left)** The safe region that the quadruped is allowed to traverse. **(Bottom Left)** A system diagram depicting the states of the quadruped $[x, y, \theta]^\top$. **(Top Right)** 50 trajectories for 3 controllers: one without any knowledge of safety ($\mathbf{k}_{\text{nom}}$), one with a standard safety filter (DTCBF-OP), and finally our method which accounts for stochasticity (JED). **(Bottom Right)** Plots of $h(\mathbf{x})$, a scalar value representing safety. The system is safe (i.e., in the green safe region) if $h(\mathbf{x}) \geq 0$.

I. Introduction

Safety is critical for a multitude of modern robotic systems, from autonomous vehicles, to medical and assistive robots, to aerospace systems. When deployed in the real world, these systems face sources of uncertainty such as imperfect perception, approximate models of the world and the system, and unexpected disturbances. In order to achieve the high degrees of safety necessary for these robots to be deployed at scale, it is essential that controllers can not only guarantee safe behavior, but also provide robustness to these uncertainties.

In the field of control theory, safety is often defined as the forward invariance of a "safe set" [6]. In this view, a closed-loop system is considered safe if all trajectories starting inside the safe set will remain in this set for all time. Several tools exist for generating controllers which can guarantee this forward-invariance property, including Control Barrier Functions (CBFs) [7], reachability-based controllers [9], and state-constrained Model-Predictive Controller (MPC) approaches [19]. Considerable advancements have been made in guaranteeing safety or stability in the presence of bounded uncertainties [37, 11, 8, 29, 20, 5]. Yet, less attention has been paid to the case of unbounded uncertainties, where the aforementioned methods generally do not apply.

Obtaining robust safety in the case of unbounded disturbances is particularly important when considering systems subject to stochastic disturbances, since these disturbances are often modeled as continuous random variables with unbounded support (e.g., zero-mean, additive Gaussian noise). For such systems, it is impossible to give an absolute bound on the disturbance magnitude. Existing methods for unbounded, random disturbances fall into two categories. The first is to impose step-wise chance constraints on a given safety criterion (e.g., a state constraint in MPC [19] or CBF-based controllers [4]), which in turn provide one-step safety guarantees. The other class of approaches [21, 26, 27, 17, 30] use Lyapunov or barrier function techniques to provide bounds on the safety probabilities for trajectories over a fixed time horizon; existing approaches, however, often assume the presence of a stabilizing controller, or model the system in continuous-time (i.e., assume the controller has, in effect, infinite bandwidth).

In order to best represent the uncertainty that might appear from sources such as discrete-time perception errors or sampled-data modeling errors, we focus our work on generating probabilistic bounds of safety for discrete-time (DT) stochastic systems. While MPC state constraints are generally enforced in discrete time, CBFs, normally applied in continuous time, have a discrete-time counterpart (DTCBFs) that were first introduced in [1] and have gained popularity due to their compatibility with planners based on MPC [36, 23, 35], reinforcement learning [15], and Markov decision processes

[3]. In a stochastic setting, martingale-based techniques have been leveraged to establish safety guarantees [27, 30], yet these works have limited utility when analyzing the safety of discrete-time CBF-based controllers.

In particular, the "c-martingale" condition used in [30] does not admit a multiplicative scaling of the barrier function, and therefore, at best, provides a weak worst-case safety bound for CBF-based controllers that grows linearly in time. The work of [27] (which builds upon [21], as does this paper) is largely focused on offline control synthesis to achieve a desired safety bound (as opposed to the online, optimization-based control studied in this work). Also, this method can only generate discrete-time controllers for affine barriers, which severely limits its applicability to general barrier functions. Both [30] and [27] depend on sum-of-squares (SoS) programming [25] for control synthesis/system verification, thereby requiring an offline step that scales poorly with the state dimension. The goal of this paper is to extend the results of [21] in a different direction, and thereby enable the synthesis of online controllers that can be realized on robotic systems.

The main contribution of this paper is to apply martingale-based probability bounds in the context of discrete-time CBFs to guarantee robust safety under stochastic uncertainty. To this end, we leverage the bounds originally presented in the seminal work by Kushner [21]. Our first key contribution is the translation of these results from a Lyapunov setting to a CBF one. To this end, we present a new proof of the results in [21] which we believe to be more complete and intuitive and which relates to the existing results of Input-to-State Safety (ISSf) for systems with bounded uncertainties [20]. Furthermore, we present a method (based on Jensen's inequality) to account for the effects of process noise on a DTCBF-based controller. Finally, we apply this method to a variety of systems in simulation to analyze the tightness of our bound and demonstrate its utility. These experiments range from simple examples that illustrate the core mathematics—a single- and double-integrator and a pendulum—to a high fidelity simulation of a quadrupedal robot locomoting along a narrow path with the uncertainty representing the gap between the simplified and full-order dynamics models.

II. Background

In this section we provide a review of safety for discrete-time nonlinear systems via control barrier functions (CBFs), and review tools from probability theory useful for studying systems with stochastic disturbances.

A. Safety of Discrete-time Systems

Consider a discrete-time (DT) nonlinear system with dynamics given by:

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k), \quad \forall k \in \mathbb{N}, \tag{1}$$

with state $\mathbf{x}_k \in \mathbb{R}^n$, input $\mathbf{u}_k \in \mathbb{R}^m$, and continuous dynamics $\mathbf{F} : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$. A continuous state-feedback controller $\mathbf{k} : \mathbb{R}^n \to \mathbb{R}^m$ yields the DT closed-loop system:

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k)), \quad \forall k \in \mathbb{N}. \tag{2}$$

We formalize the notion of safety for systems of this form using the concept of forward invariance:

Definition 1 (Forward Invariance & Safety [11])**.** *A set $\mathcal{C} \subset \mathbb{R}^n$ is forward invariant for the system* (2) *if $\mathbf{x}_0 \in \mathcal{C}$ implies that $\mathbf{x}_k \in \mathcal{C}$ for all $k \in \mathbb{N}$. In this case, we call the system* (2) *safe with respect to the set $\mathcal{C}$.*

Discrete-time barrier functions (DTBFs) are a tool for guaranteeing the safety of discrete-time systems. Consider a set $\mathcal{C} \triangleq \{\mathbf{x} \in \mathbb{R}^n \mid h(\mathbf{x}) \geq 0\}$ expressed as the 0-superlevel set of a continuous function $h : \mathbb{R}^n \to \mathbb{R}$. We refer to such a function h as a DTBF[1] if it satisfies the following properties:

Definition 2 (Discrete-Time Barrier Function (DTBF) [1])**.** *Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a continuous function $h : \mathbb{R}^n \to \mathbb{R}$. The function h is a discrete-time barrier function (DTBF) for* (2) *on $\mathcal{C}$ if there exists an $\alpha \in [0, 1]$ such that for all $\mathbf{x} \in \mathbb{R}^n$, we have that:*

$$h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x}))) \geq \alpha h(\mathbf{x}). \tag{3}$$

This inequality mimics that of discrete-time Lyapunov functions [12], and similarly regulates the evolution of h based on its previous value. DTBFs serve as a certificate of forward invariance as captured in the following theorem:

Theorem 1 ([1])**.** *Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a continuous function $h : \mathbb{R}^n \to \mathbb{R}$. If h is a DTBF for* (2) *on $\mathcal{C}$, then the system* (2) *is safe with respect to the set $\mathcal{C}$.*

Intuitively, the value of $h(\mathbf{x}_k)$ can only decay as fast as the geometric sequence $\alpha^k h(\mathbf{x}_0)$, which is lower-bounded by 0, thus ensuring the safety (i.e., forward invariance) of $\mathcal{C}$.

Discrete-time control barrier functions (DTCBFs) provide a tool for constructively synthesizing controllers that yield closed-loop systems that possess a DTBF:

Definition 3 (Discrete-Time Control Barrier Function (DTCBF) [1])**.** *Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a continuous function $h : \mathbb{R}^n \to \mathbb{R}$. The function h is a discrete-time control barrier function (DTCBF) for* (1) *on $\mathcal{C}$ if there exists an $\alpha \in [0, 1]$ such that for each $\mathbf{x} \in \mathbb{R}^n$, there exists a $\mathbf{u} \in \mathbb{R}^m$ such that:*

$$h(\mathbf{F}(\mathbf{x}, \mathbf{u})) \geq \alpha h(\mathbf{x}). \tag{4}$$

Given a CBF h for (1) and a corresponding $\alpha \in [0, 1]$, we define the point-wise set of control values:

$$\mathscr{K}_{\mathrm{CBF}}(\mathbf{x}) = \{\mathbf{u} \in \mathbb{R}^m \mid h(\mathbf{F}(\mathbf{x}, \mathbf{u})) \geq \alpha h(\mathbf{x})\}. \tag{5}$$

This yields the following result:

Theorem 2 ([2])**.** *Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of a continuous function $h : \mathbb{R}^n \to \mathbb{R}$. If h is a DTCBF for* (1) *on $\mathcal{C}$, then the set $\mathscr{K}_{\mathrm{CBF}}(\mathbf{x})$ is non-empty for all $\mathbf{x} \in \mathbb{R}^n$, and for any continuous state-feedback controller $\mathbf{k}$ with $\mathbf{k}(\mathbf{x}) \in \mathscr{K}_{\mathrm{CBF}}(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^n$, the function h is a DTBF for* (2) *on $\mathcal{C}$.*

Given a continuous nominal controller $\mathbf{k}_{\mathrm{nom}} : \mathbb{R}^n \times \mathbb{N} \to \mathbb{R}^m$ and a DTCBF h for (1) on $\mathcal{C}$, a controller $\mathbf{k}$ satisfying

[1]The state constraint $\mathbf{x}_k \in \mathcal{C}$, when expressed as $h(\mathbf{x}_k) \geq 0$, is the special case of a DTBF with $\alpha = 0$.

$\mathbf{k}(\mathbf{x},k) \in \mathscr{K}_{\text{CBF}}(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^n$ and $k \in \mathbb{N}$ can be specified via the following optimization problem:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) = \underset{\mathbf{u}\in\mathbb{R}^m}{\operatorname{argmin}} \quad & \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x},k)\|^2 \\ \text{s.t.} \quad & h(\mathbf{F}(\mathbf{x},\mathbf{u})) \geq \alpha h(\mathbf{x}). \end{aligned} \tag{DTCBF-OP}$$

We note that unlike the affine inequality constraint that arises with continuous-time CBFs [7], the DTCBF inequality constraint (4) is not necessarily convex with respect to the input, preventing it from being integrated into a convex optimization-based controller. To solve this issue, it is often assumed that the function $h \circ \mathbf{F} : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ is concave with respect to its second argument [1, 3, 36]. This assumption was shown to be well motivated for concave h [31].

B. Stochastic Preliminaries

We now review tools from probability theory that will allow us to utilize information about the distribution of a stochastic disturbance signal in constructing a notion of stochastic safety and corresponding safety-critical controllers. We choose to provide this background material at the level necessary to understand our later constructions of stochastic safety and safety-critical controllers, but refer readers to [18] for a precise measure-theoretic presentation of the following concepts.

The key tool underlying our construction of a notion of stochastic safety is a nonnegative supermartingale, a specific type of expectation-governed random process:

Definition 4. *Let $\mathbf{x}_k$ be a sequence of random variables that take values in $\mathbb{R}^n$, $W : \mathbb{R}^n \times \mathbb{N} \to \mathbb{R}$, and suppose that $\mathbb{E}\big[|W(\mathbf{x}_k,k)|\big] < \infty$ for $k \in \mathbb{N}$. The process $W_k \triangleq W(\mathbf{x}_k,k)$ is a supermartingale if:*

$$\mathbb{E}[W_{k+1} \mid \mathbf{x}_{0:k}] \leq W_k \text{ almost surely for all } k \in \mathbb{N}, \tag{6}$$

where $\mathbf{x}_{0:k}$ indicates the random variables $\{\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_k\}$. If, additionally, $W_k \geq 0$ for all $k \in \mathbb{N}$, W_k is a nonnegative supermartingale. If the process is non-decreasing in expectation, the process W_k is a submartingale. If the inequality (6) *holds with equality, the process W_k is a martingale.*

An important result from martingale theory that we will use to develop probabilistic safety guarantees is *Ville's inequality*, which allows us to bound the probability that a nonnegative supermartingale will rise above a certain value:

Theorem 3 (Ville's Inequality [33])**.** *Let W_k be a nonnegative supermartingale. Then for all $\lambda \in \mathbb{R}_{>0}$,*

$$\mathbb{P}\left\{\sup_{k\in\mathbb{N}} W_k > \lambda\right\} \leq \frac{\mathbb{E}[W_0]}{\lambda}. \tag{7}$$

Intuitively, Ville's inequality can be compared with Markov's inequality for nonnegative random variables; since the process W_k is nonincreasing in expectation, Ville's inequality allows us to control the probability the process instead moves upward above λ.

Lastly, as we will see when synthesizing safety-critical controllers in the presence of stochastic disturbances, we will need to enforce conditions on the expectation of a DCTBF. In doing so, we will need to relate the expectation of the DCTBF $h(\mathbf{x}_{k+1})$ to the expectation of the state $\mathbf{x}_{k+1}$. This will be achieved using Jensen's inequality:

Theorem 4 (Jensen's Inequality [22])**.** *Consider a continuous function $h : \mathbb{R}^n \to \mathbb{R}$ and a random variable $\mathbf{x}$ that takes values in $\mathbb{R}^n$ with $\mathbb{E}[\|\mathbf{x}\|] < \infty$. We have that:*

$$\begin{cases} \text{if } h \text{ is convex,} & \text{then } \mathbb{E}[h(\mathbf{x})] \geq h(\mathbb{E}[\mathbf{x}]), \\ \text{if } h \text{ is concave,} & \text{then } \mathbb{E}[h(\mathbf{x})] \leq h(\mathbb{E}[\mathbf{x}]). \end{cases} \tag{8}$$

III. Safety of Discrete-Time Stochastic Systems

In this section we provide one of our main results in the form of a bound on the probability that a system with stochastic disturbances will exit a given superlevel set of a DTBF over a finite time horizon.

Consider the following modification of the DT system (1):

$$\mathbf{x}_{k+1} = \mathbf{F}_d(\mathbf{x}_k, \mathbf{u}_k, \mathbf{d}_k), \quad \forall k \in \mathbb{N}, \tag{9}$$

where the dynamics $\mathbf{F}_d : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^\ell \to \mathbb{R}^n$ now also includes a disturbance $\mathbf{d}_k$ taking values in $\mathbb{R}^\ell$. For any state-feedback controller, this has an associated closed-loop system:

$$\mathbf{x}_{k+1} = \mathbf{F}_d(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k), \mathbf{d}_k), \quad \forall k \in \mathbb{N}. \tag{10}$$

We assume that $\mathbf{x}_0$ is known and the disturbances $\mathbf{d}_k$ are a sequence of independent and identically distributed (with distribution $\mathcal{D}$) random variables[2] with (potentially unbounded) support on $\mathbb{R}^\ell$, generating the random process $\mathbf{x}_{1:k}$. To study the safety of this system, we will use the following definition:

Definition 5 (K-Step Exit Probability)**.** *Let $h : \mathbb{R}^n \to \mathbb{R}$ be a continuous function. For any $K \in \mathbb{N}$, $\gamma \in \mathbb{R}_{\geq 0}$, and initial condition $\mathbf{x}_0 \in \mathbb{R}^n$, the K-step exit probability of the closed-loop system* (10) *is given by:*

$$P_u(K, \gamma, \mathbf{x}_0) = \mathbb{P}\left\{\min_{k\in\{0,\ldots,K\}} h(\mathbf{x}_k) < -\gamma\right\}. \tag{11}$$

which describes the probability that the system will leave the $-\gamma$ superlevel set of h within K steps. This probability is directly related to the robust safety concept of Input-to-State Safety (ISSf) [20] which reasons about the superlevel set of h which is rendered safe in the presence of bounded disturbances. For the remainder of this work, we will omit the dependence of P_u on K, γ, and $\mathbf{x}_0$ for notational simplicity.

Remark 1. The finite time aspect of K-step exit probabilities is critical since systems exposed to unbounded disturbances will exit a bounded set with probability $P_u = 1$ over an infinite horizon [30, 16]. Intuitively, this is because a sufficiently large sample will eventually be drawn from the tail of the distribution that forces the system out in a single step.

Given this definition, we now provide one of our main results relating DTBFs to K-step exit probabilities. We note that this result is a reframing of the stochastic invariance theorem in [21, 27]. Our reframing features three key components.

[2]This implies the dynamics define a Markov process, i.e. $\mathbb{E}[h(\mathbf{F}_d(\mathbf{x}_k,\mathbf{u}_k,\mathbf{d}_k)) \mid \mathbf{x}_{0:k}] = \mathbb{E}[h(\mathbf{F}_d(\mathbf{x}_k,\mathbf{u}_k,\mathbf{d}_k)) \mid \mathbf{x}_k]$, since the state $\mathbf{x}_{k+1}$ at time $k+1$ only depends on the state $\mathbf{x}_k$, input $\mathbf{u}_k$, and disturbance $\mathbf{d}_k$ at time k.

First, we develop our results using the standard formulation of DTBFs covered in the background. Second, we produce a probability bound not only for $\mathcal{C}$ (defined as the 0-superlevel set of h, such that $\gamma = 0$), but for all non-positive superlevel sets of h ($\gamma \geq 0$), a stochastic variant of ISSf [20]. Third, we present a complete proof of our result, with the goal of illuminating how to leverage tools from martingale theory to reason about the safety of discrete-time stochastic systems.

Theorem 5. *Let $h : \mathbb{R}^n \to \mathbb{R}$ be a continuous, upper-bounded function with upper bound $M \in \mathbb{R}_{>0}$. Suppose there exists an $\alpha \in (0, 1)$ and*[3] *$\delta \leq M(1-\alpha)$ such that the closed-loop system* (10) *satisfies:*

$$\mathbb{E}[\ h(\mathbf{F}_d(\mathbf{x}, \mathbf{k}(\mathbf{x}), \mathbf{d})) \mid \mathbf{x}\ ] \geq \alpha h(\mathbf{x}) + \delta, \tag{12}$$

for all $\mathbf{x} \in \mathbb{R}^n$, with $\mathbf{d} \sim \mathcal{D}$. For any $K \in \mathbb{N}$ and $\gamma \in \mathbb{R}_{\geq 0}$, if $\delta < -\gamma(1-\alpha)$, we have that:

$$P_u \leq \left(\frac{M - h(\mathbf{x}_0)}{M+\gamma}\right)\alpha^K + \frac{M(1-\alpha)-\delta}{M+\gamma}\sum_{i=1}^{K}\alpha^{i-1}. \tag{13}$$

Alternatively if $\delta \geq -\gamma(1-\alpha)$, then:

$$P_u \leq 1 - \frac{h(\mathbf{x}_0)+\gamma}{M+\gamma}\left(\frac{M\alpha+\gamma+\delta}{M+\gamma}\right)^K. \tag{14}$$

Remark 2. The upper bound $\delta \leq M(1-\alpha)$ is relatively nonrestrictive, as not only is δ typically negative, but it must hold such that, in expectation, $h(\mathbf{x}_{k+1})$ cannot rise above the upper bound M on h. The switching condition between (13) and (14) of $\delta = \gamma(1-\alpha)$ corresponds to whether, in expectation, the one-step evolution of the system remains in the set $\mathcal{C}_\gamma = \{\mathbf{x} \in \mathbb{R}^n \mid h(\mathbf{x}) \geq -\gamma\}$ when it begins on the boundary of $\mathcal{C}_\gamma$.

To make our argument clear at a high level, we begin with a short proof sketch before proceeding in detail.

Proof sketch: The key tool in proving Theorem 5 is Ville's inequality (7). Since $h(\mathbf{x}_k)$, in general, is not a super- or submartingale, we will first construct a nonnegative supermartingale, $W_k \triangleq W(\mathbf{x}_k, k)$, by scaling and shifting $h(\mathbf{x}_k)$. We can then apply Ville's inequality (7) to bound the probability of W_k going above any $\lambda > 0$. Next we find a particular value of λ, denoted λ^*, such that:

$$\max_{k\in\{0,\ldots,K\}} W_k \leq \lambda^* \implies \min_{k\in\{0,\ldots,K\}} h(\mathbf{x}_k) \geq -\gamma. \tag{15}$$

Intuitively, this means that any sequence W_k that remains below λ^* ensures that the corresponding sequence $h(\mathbf{x}_k)$ remains (safe) above $-\gamma$. This allows us to bound the K-step exit probability P_u of our original process $h(\mathbf{x}_k)$ with the probability that W_k will rise above λ^*:

$$P_u \leq \mathbb{P}\left\{\max_{k\in\{0,\ldots,K\}} W_k > \lambda^*\right\} \leq \frac{\mathbb{E}[W_0]}{\lambda^*} = \frac{W_0}{\lambda^*}, \tag{16}$$

where the last equality will follow as it is assumed $\mathbf{x}_0$ is known *a priori*. Particular choices of W and λ^* will yield the bounds stated in the theorem, completing the proof.

[3] The original presentation of Theorem 5 in [21] considers variable δ_k for $k \in \{0, \ldots, K\}$, which are known *a priori*. In most practical applications, one assumes a lower bound that holds for all δ_k, motivating our use of a constant δ. Moreover, the use of a constant δ significantly clarifies the proof.

A. Proof: Constructing a Nonnegative Supermartingale

We will begin by constructing a nonnegative supermartingale, allowing us to use Ville's inequality. To construct this supermartingale, we first note that by rearranging terms in the inequality in (12), we can see the process $M - h(\mathbf{x}_k)$ resembles a supermartingale:

$$\begin{aligned}\mathbb{E}[M - h(\mathbf{x}_{k+1}) \mid \mathbf{x}_k] &\leq \alpha(M - h(\mathbf{x}_k)) + M(1-\alpha) - \delta,\\ &\triangleq \alpha(M - h(\mathbf{x}_k)) + \varphi,\end{aligned} \tag{17}$$

but with a scaling α and additive term $\varphi \triangleq M(1-\alpha) - \delta$ that makes $\mathbb{E}\left[M - h(\mathbf{x}_{k+1}) \mid \mathbf{x}_k\right] \not\leq M - h(\mathbf{x}_k)$ in general. To remove the effects of α and φ, consider the function $W : \mathbb{R}^n \times \mathbb{N} \to \mathbb{R}$ defined as:

$$W(\mathbf{x}_k, k) \triangleq \underbrace{(M - h(\mathbf{x}_k))\theta^k}_{\text{negate and scale}} \underbrace{- \varphi\sum_{i=1}^{k}\theta^i}_{\text{cancel } \varphi} + \underbrace{\varphi\sum_{i=1}^{K}\theta^i}_{\text{ensure } W \geq 0}, \tag{18}$$

where $\theta \in [1, \infty)$ will be used to cancel the effect of α, but is left as a free variable that we will later use to tighten our bound on P_u. Denoting $W_k \triangleq W(\mathbf{x}_k, k)$, we now verify W_k is a nonnegative supermartingale. We first show that $W_k \geq 0$ for all $k \in \{0, \ldots, K\}$. Combining the two sums in (18) yields:

$$W_k = (M - h(\mathbf{x}_k))\theta^k + \varphi\sum_{i=k+1}^{K}\theta^i, \tag{19}$$

which is nonnegative as $h(\mathbf{x}) \leq M$ for all $\mathbf{x} \in \mathbb{R}^n$, $\theta \geq 1$, and $\varphi \geq 0$ since $\delta \leq M(1-\alpha)$ by assumption. We now show that W_k satisfies the supermartingale inequality (6):

$$\mathbb{E}[W_{k+1} \mid \mathbf{x}_{0:k}] = \mathbb{E}[W_{k+1} \mid \mathbf{x}_k], \tag{20}$$

$$= (M - \mathbb{E}[h(\mathbf{x}_{k+1}) \mid \mathbf{x}_k])\theta^{k+1} + \varphi\sum_{i=k+2}^{K}\theta^i, \tag{21}$$

$$\leq (M - \alpha h(\mathbf{x}_k) - \delta)\theta^{k+1} + \varphi\sum_{i=k+2}^{K}\theta^i, \tag{22}$$

$$= \alpha\theta(M - h(\mathbf{x}_k))\theta^k + \theta^{k+1}\underbrace{((1-\alpha)M - \delta)}_{=\varphi} + \varphi\sum_{i=k+2}^{K}\theta^i,$$

$$= \underbrace{\alpha\theta}_{\text{req.} \leq 1}(M - h(\mathbf{x}_k))\theta^k + \varphi\sum_{i=k+1}^{K}\theta^i \leq W_k, \tag{23}$$

where (20) is due to the Markovian nature of system (10), (21) comes from using (19) to write W_{k+1}, (22) follows from (12), and (23) follows from the preceding line using the definition of φ and assuming the further requirement that $\theta \leq \frac{1}{\alpha}$. Thus, we have shown that W_k is a nonnegative supermartingale.

B. Proof: Bounding the Exit Probability via Ville's Inequality

Since W_k is a nonnegative supermartingale, we can apply Ville's inequality to establish:

$$\mathbb{P}\left\{\max_{k\in\{0,\ldots,K\}} W_k > \lambda\right\} \leq \frac{\mathbb{E}[W_0]}{\lambda} = \frac{W_0}{\lambda}. \tag{24}$$

for all $\lambda \in \mathbb{R}_{>0}$. To relate this bound to the K-step exit probability P_u, we seek a value of λ, denoted λ^*, such that:

$$\max_{k\in\{0,\dots,K\}} W_k \le \lambda^*. \implies \min_{k\in\{0,\dots,K\}} h(\mathbf{x}_k) \ge -\gamma. \quad (25)$$

In short, we will choose a value of λ^* that is small enough to ensure that all trajectories of W_k that remain below λ^* must also have $h_k \ge -\gamma$. To this end, we use the geometric series identity[4] $\sum_{i=1}^{k} \theta^{i-1} = \frac{1-\theta^k}{1-\theta}$ to rewrite W_k as:

$$W_k = (M - h(\mathbf{x}_k))\theta^k + \varphi\theta\frac{\theta^K - \theta^k}{\theta - 1}. \quad (26)$$

Let us define:

$$\lambda_k = \left(\gamma + M - \frac{\varphi\theta}{\theta-1}\right)\theta^k + \frac{\varphi\theta}{\theta-1}\theta^K > 0, \quad (27)$$

which, intuitively, applies the same time-varying scaling and shift to a constant, $-\gamma$, that was applied to $h(\mathbf{x}_k)$ to yield W_k (26). Let us choose:

$$\lambda^* \triangleq \min_{k\in\{0,\dots,K\}} \lambda_k. \quad (28)$$

Since we assume $\max_{k\in\{0,\dots,K\}} W_k \le \lambda^*$, we can write, for all $k \in \{0,\dots,K\}$:

$$0 \ge W_k - \lambda^* \ge W_k - \lambda_k = (-\gamma - h_k)\theta^k. \quad (29)$$

Since $\theta > 1$, this implies that $-\gamma - h_k \le 0$ for all $k \in \{0,\dots,K\}$, and thus $\min_{k\in\{0,\dots,K\}} h(\mathbf{x}_k) \ge -\gamma$, as needed.

C. Proof: Choosing θ to Minimize the Ville's Bound

Since our supermartingale W_k includes a free parameter $\theta \in (1, \frac{1}{\alpha}]$, we will choose a value of θ in this interval which provide the tightest bound on P_u.

Case 1: Consider the first case where $\delta < -\gamma(1-\alpha)$, implying $\varphi > (M+\gamma)(1-\alpha)$. In this case $\frac{1}{\alpha} < \frac{M+\gamma}{M+\gamma-\varphi}$ and thus all of the allowable choices of $\theta \in (1, \frac{1}{\alpha})$ are such that $\theta < \frac{M+\gamma}{M+\gamma-\varphi}$. Denoting k^* such that $\lambda^* = \lambda_{k^*}$, we have that:

$$\lambda^* = \underbrace{\left(\gamma + M - \frac{\varphi\theta}{\theta-1}\right)}_{\le 0}\theta^{k^*} + \frac{\varphi\theta}{\theta-1}\theta^K. \quad (30)$$

Thus, we know $\min_{k\in\{0,\dots,K\}} \lambda_k$ occurs at $k^* = K$ and so:

$$P_u \le \frac{W_0}{\lambda^*} = \frac{M - h(\mathbf{x}_0) + \frac{\varphi\theta}{\theta-1}\left(\theta^K - 1\right)}{(M+\gamma)\theta^K}. \quad (31)$$

Since this bound is a decreasing function of θ (as shown in Lemma 2 in Appendix A), we choose the largest allowable value $\theta^* = \frac{1}{\alpha}$ to achieve the bound:

$$P_u \le \frac{W_0}{\lambda^*} = \frac{M - h(\mathbf{x}_0) + \frac{\varphi}{1-\alpha}\left(\alpha^{-K} - 1\right)}{(M+\gamma)\alpha^{-K}}, \quad (32)$$

$$= \left(\frac{M - h(\mathbf{x}_0)}{M+\gamma}\right)\alpha^K + \frac{M(1-\alpha) - \delta}{M+\gamma}\sum_{i=1}^{K}\alpha^{i-1}, \quad (33)$$

where we again use the geometric series identity.

Case 2: Now consider the second case where $\delta \ge -\gamma(1-\alpha)$, so $\varphi \le (M+\gamma)(1-\alpha)$, which implies that the set $[\frac{M+\gamma}{M+\gamma-\varphi}, \frac{1}{\alpha}]$ is nonempty. Choosing a value of θ in this set ensures that:

$$\lambda^* = \underbrace{\left(\gamma + M - \frac{\varphi\theta}{\theta-1}\right)}_{\ge 0}\theta^{k^*} + \frac{\varphi\theta}{\theta-1}\theta^K. \quad (34)$$

Thus $\min_{k\in\{0,\dots,K\}} \lambda_k$ occurs at $k^* = 0$ and:

$$P_u \le \frac{W_0}{\lambda} = \frac{(M - h(\mathbf{x}_0)) + \frac{\varphi\theta}{\theta-1}\left(\theta^K - 1\right)}{(M+\gamma) + \frac{\varphi\theta}{\theta-1}\left(\theta^K - 1\right)}, \quad (35)$$

$$= 1 - \frac{h(\mathbf{x}_0) + \gamma}{M + \gamma + \frac{\varphi\theta}{\theta-1}\left(\theta^K - 1\right)}. \quad (36)$$

Since this bound is increasing in θ (as shown in Lemma 3 in Appendix A), we choose $\theta^* = \frac{M+\gamma}{M+\gamma-\varphi}$ to achieve the bound:

$$P_u \le 1 - \left(\frac{h(\mathbf{x}_0) + \gamma}{M+\gamma}\right)\left(\frac{M\alpha + \gamma + \delta}{M+\gamma}\right)^K. \quad (37)$$

If, alternatively, we choose $\theta \in \left(1, \frac{M+\gamma}{M+\gamma-\varphi}\right]$, then the inequality in (30) holds, $k^* = K$, and the bound is decreasing in θ as in Case 1. Evaluating this bound for the minimizing value $\theta^* = \frac{M+\gamma}{M+\gamma-\varphi}$ again yields:

$$P_u \le \frac{M - h(\mathbf{x}_0) + (M+\gamma)(\theta^K - 1)}{(M+\gamma)\theta^K}, \quad (38)$$

$$= 1 - \left(\frac{h(\mathbf{x}_0) + \gamma}{M+\gamma}\right)\left(\frac{M\alpha + \gamma + \delta}{M+\gamma}\right)^K. \quad (39)$$

■

IV. Practical Considerations for Enforcing Stochastic DTCBFs

Theorem 5 allows us to reason about the finite-time safety of systems governed by DTBFs. To utilize the results of this theorem in a control setting, we aim to use DTCBFs to develop control methods which enforce the expectation condition:

$$\mathbb{E}[h(\mathbf{F}_d(\mathbf{x}_k, \mathbf{u}_k, \mathbf{d}_k)) \mid \mathbf{x}_k] \ge \alpha h(\mathbf{x}_k). \quad (40)$$

If such a condition can be enforced, then the result of Theorem 5 can be directly applied to provide probabilistic bounds on the system's safety.[5] However, since the composition of system dynamics with the disturbance may make computing this expectation difficult, we instead focus on systems with additive disturbances of the form:

$$\mathbf{x}_k = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{d}_k, \quad (41)$$

where $\mathbf{d}_k$ takes values in $\mathbb{R}^n$ and the expectation condition for Theorem 5 becomes,

$$\mathbb{E}[h(\mathbf{F}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{d}_k) \mid \mathbf{x}_k] \ge \alpha h(\mathbf{x}_k). \quad (42)$$

[4]At $\theta = 1$, the fraction $\frac{1-\theta^k}{1-\theta}$ is not well defined. However, the proof can be carried out using the summation notation. In this case $\lambda^* = M + \gamma$, and (24) yields $P_u \le 1 - \frac{h(\mathbf{x}_0)+\gamma-\varphi K}{M+\gamma}$.

[5]We note that the bound provided by Theorem 5 assumes that there exists a control $\mathbf{u}_k$ satisfying the constraint (40) for any state $\mathbf{x}_k$. This may not be possible for systems with limited actuation or underactuated dynamics – algorithms for verifying this existence for a particular choice of h, α is an interesting direction for future work.

Like the DTCBF-OP controller, we seek to enforce this constraint using an optimization-based controller that enforces safety while achieving pointwise minimal deviation from a nominal controller $\mathbf{k}_{\text{nom}}$ in the form of an Expectation-based DTCBF (ED) Controller:

$$\begin{aligned} \mathbf{k}_{\text{ED}}(\mathbf{x}_k) = \underset{\mathbf{u}\in\mathbb{R}^m}{\operatorname{argmin}} \quad & \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x}_k, k)\|^2 \\ \text{s.t.} \quad & \mathbb{E}[h(\mathbf{F}(\mathbf{x}_k, \mathbf{u}) + \mathbf{d}_k) \mid \mathbf{x}_k] \geq \alpha h(\mathbf{x}_k). \end{aligned} \tag{ED}$$

The expectation in (ED) adds complexity that is not generally considered in the application of deterministic DTCBFs. More commonly, CBF-based controllers solve "certainty-equivalent" optimization programs, like this Certainty-Equivalent DTCBF (CED) controller, that replaces the expected barrier value $\mathbb{E}[h(\mathbf{x}_{k+1}) \mid \mathbf{x}_k]$ with the barrier evaluated at the expected next state, $h(\mathbb{E}[\mathbf{x}_{k+1} \mid \mathbf{x}_k])$:

$$\begin{aligned} \mathbf{k}_{\text{CED}}(\mathbf{x}_k) = \underset{\mathbf{u}\in\mathbb{R}^m}{\operatorname{argmin}} \quad & \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x}_k, k)\|^2 \\ \text{s.t.} \quad & h(\mathbf{F}(\mathbf{x}_k, \mathbf{u}) + \mathbb{E}[\mathbf{d}_k]) \geq \alpha h(\mathbf{x}_k). \end{aligned} \tag{CED}$$

where $\mathbb{E}[\mathbf{F}(\mathbf{x}_k, \mathbf{u}_k)|\mathbf{x}_k] = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k)$ and $\mathbb{E}[\mathbf{d}_k|\mathbf{x}_k] = \mathbb{E}[\mathbf{d}_k]$. This constraint is often easier to evaluate than (40) since it allows control actions to be selected with respect to the expected disturbance $\mathbb{E}[\mathbf{d}_k]$ without needing to model the disturbance distribution $\mathcal{D}$. If the disturbance is zero-mean, then this form of the constraint is implicitly enforced by DTCBF controllers such as those presented in [1, 36]. However, when replacing ED with CED it is important to consider the effect of Jensen's inequality in Theorem 4.

If the "certainty-equivalent" constraint in CED is strictly concave[6], then we can apply the results of Theorem 5 directly since Jensen's inequality tightens the constraint and ensures satisfaction of the expectation condition (12). Unfortunately, using such a controller is a non-convex optimization program which can be impractical to solve. If, instead, the constraint is convex, then CED is a convex program, but does not necessarily enforce the expectation condition (12) in Theorem (5) due to the gap introduced by Jensen's inequality.

In order to apply the results of Theorem 5 to controllers of the form (CED) with convex constraints, we must first provide a bound on the gap introduced by Jensen's inequality. In particular, for any concave function $h : \mathbb{R}^n \to \mathbb{R}$ and random variable $\mathbf{d} \sim \mathcal{D}$, we seek to determine a value $\psi \in \mathbb{R}_{\geq 0}$ such that, for all $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{u} \in \mathbb{R}^m$:

$$\mathbb{E}[h(\mathbf{F}(\mathbf{x}, \mathbf{u}) + \mathbf{d}) \mid \mathbf{x}] \geq h(\mathbf{F}(\mathbf{x}, \mathbf{u}) + \mathbb{E}[\mathbf{d}]) - \psi, \tag{43}$$

thus quantifying the gap introduced by Jensen's inequality.

A large body of work has studied methods for finding the smallest possible ψ that satisfies (43). Here we adapt a result in [10] to achieve a relatively loose, but straightforward bound:

Lemma 1. *Consider a twice-continuously differentiable, concave function $h : \mathbb{R}^n \to \mathbb{R}$ with $\sup_{\mathbf{x}\in\mathbb{R}^n} \|\nabla^2 h(\mathbf{x})\|_2 \leq \lambda_{\max}$ for some $\lambda_{\max} \in \mathbb{R}_{\geq 0}$, and a random variable $\mathbf{x}$ that takes values in $\mathbb{R}^n$ with $\mathbb{E}[\|\mathbf{x}\|] < \infty$ and $\|\text{cov}(\mathbf{x})\| < \infty$. Then we have that:*

$$\mathbb{E}[h(\mathbf{x})] \geq h(\mathbb{E}[\mathbf{x}]) - \frac{\lambda_{\max}}{2}\text{tr}(\text{cov}(\mathbf{x})). \tag{44}$$

[6] The constraint $h(\mathbf{x}_k + \mathbf{u}) \geq \alpha h(\mathbf{x}_k)$ is concave in $\mathbf{u}$ when h is convex and it is convex in $\mathbf{u}$ when h is concave.

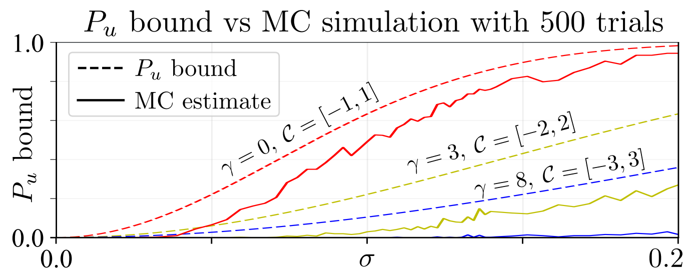

Fig. 2. The dashed lines represent the theoretical probability bounds for the system as in Theorem 5. The solid lines represent the Monte Carlo (MC) estimated P_u across 500 experiments.

The proof is included in Appendix B. We note that although this value of $\psi = \frac{\lambda_{\max}}{2}\text{tr}(\text{cov}(\mathbf{x}))$ is easy to interpret, tighter bounds exist which have less restrictive assumptions than a globally bounded Hessian [22]. We also note that one could also use sampling-based methods to approximately satisfy the constraint (43) by estimating ψ empirically.

Next we present a controller which combines the mean-based control of the "certainty equivalent" (CED) while also accounting for Jensen's inequality. This Jensen-Enhanced DTCBF Controller (JED) includes an additional control parameter $c_{\text{J}} \geq 0$ to account for Jensen's inequality:

$$\begin{aligned} \mathbf{k}_{\text{JED}}(\mathbf{x}_k) = \underset{\mathbf{u}\in\mathbb{R}^m}{\operatorname{argmin}} \quad & \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x}_k, k)\|^2 \\ \text{s.t.} \quad & h(\mathbf{F}(\mathbf{x}_k, \mathbf{u}_k) + \mathbb{E}[\mathbf{d}_k]) - c_{\text{J}} \geq \alpha h(\mathbf{x}_k). \end{aligned} \tag{JED}$$

Given this controller and a method for bounding ψ, we can now apply Theorem 5 while accounting for (or analyzing) the effects of Jensen's inequality on the (JED) controller:

Theorem 6. *Consider the system (41) and let $h : \mathbb{R}^n \to \mathbb{R}$ be a twice-continuously differentiable, concave function such that $\sup_{\mathbf{x}\in\mathbb{R}^n} h(\mathbf{x}) \leq M$ for $M \in \mathbb{R}_{>0}$ and $\sup_{\mathbf{x}\in\mathbb{R}^n} \|\nabla^2 h(\mathbf{x})\|_2 \leq \lambda_{\max}$ for $\lambda_{\max} \in \mathbb{R}_{\geq 0}$. Suppose there exists an $\alpha \in (0, 1)$ and a $c_{\text{J}} \in [0, \frac{\lambda_{\max}}{2}\text{tr}(\text{cov}(\mathbf{d})) + M(1-\alpha)]$ such that:*

$$h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \mathbb{E}[\mathbf{d}]) - c_{\text{J}} \geq \alpha h(\mathbf{x}), \tag{45}$$

for all $\mathbf{x} \in \mathbb{R}^n$ with $\mathbf{d} \sim \mathcal{D}$. Then we have that:

$$\mathbb{E}[\ h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \mathbf{d}) \mid \mathbf{x}\ ] \geq \alpha h(\mathbf{x}) + \delta, \tag{46}$$

for all $\mathbf{x} \in \mathbb{R}^n$ with $\mathbf{d} \sim \mathcal{D}$ and $\delta = c_{\text{J}} - \frac{\lambda_{\max}}{2}\text{tr}(\text{cov}(\mathbf{d}_k))$.

Proof: Given $\mathbf{x} \in \mathbb{R}^n$, Lemma 1 ensures that:

$$\begin{aligned} 0 &\leq h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \mathbb{E}[\mathbf{d}]) - c_{\text{J}} - \alpha h(\mathbf{x}) & (47) \\ &\leq \mathbb{E}[h(\mathbf{F}(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \mathbf{d}) \mid \mathbf{x}] + \psi - c_{\text{J}} - \alpha h(\mathbf{x}) & (48) \end{aligned}$$

where $\psi = \frac{\lambda_{\max}}{2}\text{tr}(\text{cov}(\mathbf{d}))$. Letting $\delta = c_{\text{J}} - \frac{\lambda_{\max}}{2}\text{tr}(\text{cov}(\mathbf{d}))$ yields the the desired result. ■

V. Practical Examples

In this section we consider a variety of simulation examples that highlight the key features of our approach.

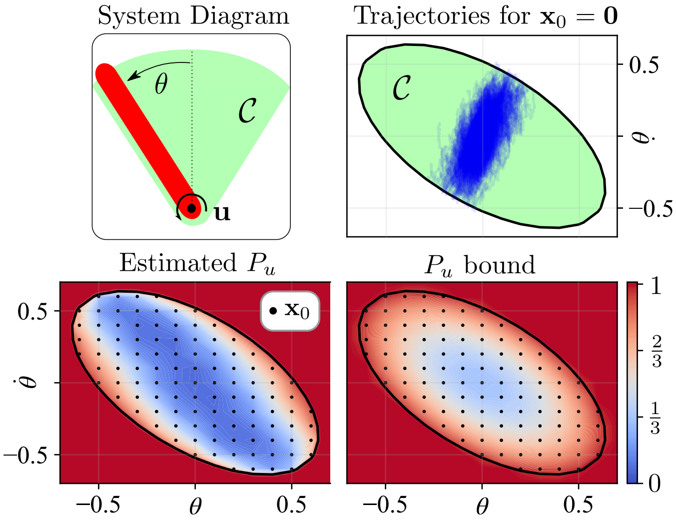

Fig. 3. **(Top Left)** System diagram of the inverted pendulum. **(Top Right)** 500 one second long example trajectories starting at $\mathbf{x}_0 = 0$. **(Bottom Left)** Monte Carlo estimates of P_u for $\gamma = 0$ using 500 one second long trials for each initial conditions represented by a black dot. **(Bottom Right)** Our (conservative) theoretical bounds on P_u from Theorem 5

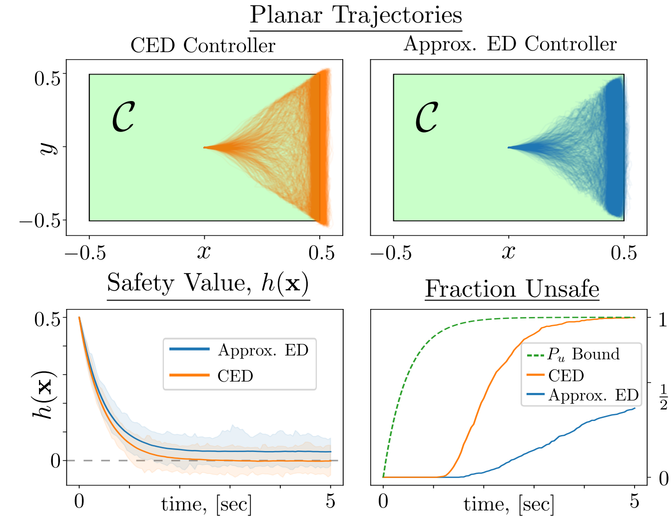

Fig. 4. Simulation results for double integrator over 500 trials. **(Top left):** Planar (x, y) trajectories for the approximated ED controller, with the safe set (a unit square) plotted in green. **(Top right):** Planar (x, y) trajectories for a CED controller. **(Bottom left):** The $h(\mathbf{x}_k)$ for both controllers, with the max and min values shaded. **(Bottom right):** Percent of trajectories that have remained safe over time. We also plot our (conservative) bound (14) on the unsafe probability P_u.

A. Linear 1D System

Here we analyze our bounds by considering the case of unbounded i.i.d. disturbances $d_k \sim \mathcal{N}(0, 1)$ for the one dimensional system $(x, u, \in \mathbb{R})$ and safe set:

$$x_{k+1} = x_k + 2 + u_k + \sigma d_k, \;\; \mathcal{C} = \{x \mid 1 - x^2 \geq 0\}. \quad (49)$$

The Jensen gap for this system and DTCBF is bounded by $\psi = \sigma^2$. For simulation, we employ the JED controller with $c_J = \sigma^2$, $\alpha = 1 - \sigma^2$, and nominal controller $\mathbf{k}_{\text{nom}}(\mathbf{x}_k, k) = 0$. Figure 2 shows the results of 500 one second long trials run with a variety of $\sigma \in [0, 0.2]$ and also displays how the bound on P_u decreases as γ increases.

B. Simple Pendulum

Next we consider an inverted pendulum about its upright equilibrium point with the DT dynamics:

$$\begin{bmatrix} \theta_{k+1} \\ \dot{\theta}_{k+1} \end{bmatrix} = \begin{bmatrix} \theta_k + \Delta t \dot{\theta}_k \\ \dot{\theta}_k + \Delta t \sin(\theta_k) \end{bmatrix} + \begin{bmatrix} 0 \\ \Delta t \mathbf{u} \end{bmatrix} + \mathbf{d}_k, \quad (50)$$

with time step $\Delta_t = 0.01$ sec, i.i.d disturbances $\mathbf{d}_k \sim \mathcal{N}(\mathbf{0}_2, \text{Diag}([0.005^2, 0.025^2]))$, and safe set[7]:

$$\mathcal{C} = \left\{ \mathbf{x} \in \mathbb{R}^n \;\middle|\; \underbrace{1 - \frac{6^2}{\pi^2} \mathbf{x}^\top \begin{bmatrix} 1 & 3^{-\frac{1}{2}} \\ 3^{-\frac{1}{2}} & 1 \end{bmatrix} \mathbf{x}}_{h_{\text{pend}}(\mathbf{x})} \geq 0 \right\} \quad (51)$$

which is constructed using the continuous-time Lyapunov equation as in [31] and for which $|\theta| \leq \pi/6$ for all $\mathbf{x} \in \mathcal{C}$. Figure 3 shows the results of 500 one second long trials for each $\mathbf{x}_0 \in \mathcal{C}$ using the JED controller with parameters $\alpha = 1 - \psi, c_J = \psi$, where $\psi = \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{d}_k))$. This figure highlights the influence of $\mathbf{x}_0$ and shows how the bound on P_u increases as $h(\mathbf{x}_0)$ decreases.

C. Double Integrator

We also consider the problem of controlling a planar system with unit-mass double-integrator dynamics to remain inside a convex polytope (in particular, a unit square centered at the origin). Using Heun's method, the dynamics are given by:

$$\mathbf{x}_{k+1} = \begin{bmatrix} \mathbf{I}_2 & \Delta t\, \mathbf{I}_2 \\ \mathbf{0}_2 & \mathbf{I}_2 \end{bmatrix} \mathbf{x}_k + \begin{bmatrix} \frac{\Delta t^2}{2} \mathbf{I}_2 \\ \Delta t \mathbf{I}_2 \end{bmatrix} \mathbf{u}_k + \mathbf{d}_k, \quad (52)$$

$$\triangleq \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k + \mathbf{d}_k, \quad (53)$$

where Δt is the integration time step and $\mathbf{d}_k \sim \mathcal{N}(\mathbf{0}_4, \mathbf{Q})$ is a zero-mean Gaussian process noise added to the dynamics. Here we use $\Delta t = 0.05$ sec, and $\mathbf{Q} = \mathbf{B}\mathbf{B}^T$, which corresponds to applying a disturbance force $\mathbf{f}_k \sim \mathcal{N}(0, \mathbf{I}_2)$ to the system at each timestep.

To keep the system inside a convex polytope, we seek to enforce the affine inequalities $\mathbf{C}\mathbf{x} \leq \mathbf{w}$ for $\mathbf{C} \in \mathbb{R}^{n_c \times n}, \mathbf{w} \in \mathbb{R}^{n_c}$. Thus, we define our barrier $h(\mathbf{x}) = -\max(\mathbf{C}\mathbf{x} - \mathbf{w})$, where $\max(\cdot)$ defines the element-wise maximum, and $h(\mathbf{x}) \geq 0$ if and only if the constraint $\mathbf{C}\mathbf{x} \leq \mathbf{w}$ holds. Implementing the ED controller for this system is non-trivial, since the expectation of $h(\mathbf{x})$ for a Gaussian-distributed $\mathbf{x}$ does not have a closed form. Similarly, implementing the JED controller to account for Jensen's inequality is non-trivial since h is not twice continuously differentiable. We instead choose to enforce a conservative approximation of the barrier condition (40) using the *log-sum-exp* function. As we show in Appendix C, this approximation yields an analytic lower bound (derived using the moment-generating function of Gaussian r.v.s) on $\mathbb{E}[h(\mathbf{x}_{k+1})]$ which can be imposed via a convex constraint.

Figure 4 plots the results of 500 simulated trajectories for the double integrator system using the proposed ED controller,

[7] Diag: $\mathbb{R}^n \to \mathbb{R}^{n \times n}$ generates a square diagonal matrix with its argument along the main diagonal.

and the certainty equivalent CED controller that neglects the presence of process noise. Both controllers have a nominal controller $\mathbf{k}_{\text{nom}}(\mathbf{x}) = [50, 0]$ which seeks to drive the system into the right wall. All trajectories start from the origin. We note the proposed controller is indeed more conservative than the CED controller, yielding both fewer and smaller violations of the safe set. In the bottom right, we also plot our bound as a function of the time horizon, which we note is quite conservative compared to our Monte Carlo estimate of the safety probability, motivating future work.

D. Quadruped

Finally, we consider the problem of controlling a simulated quadrupedal robot locomoting along a narrow path. The simulation is based on a Unitree A1 robot as shown in Figure 1 which has 18 degrees of freedom and 12 actuators. with configuration space coordinates $\mathbf{q} \in \mathbb{R}^{18}$, the full state is given by $\mathbf{x} = (\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^{36}$. For simulated walking, a no-slip condition, $\mathbf{c}(\mathbf{q}) = \mathbf{0} \in \mathbb{R}^{c}$, is enforced on the feet where c depends on the number of feet in contact with the ground. As discussed in [32], when $\mathbf{c}(\mathbf{q})$ is differentiated twice, D'Alembert's principle applied to the constrained Euler-Lagrange equations yields the robotic system dynamics:

$$\mathbf{D}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{H}(\mathbf{q}, \dot{\mathbf{q}}) = \mathbf{B}\mathbf{u} + \mathbf{J}(\mathbf{q})^\top \boldsymbol{\lambda}, \tag{54}$$

$$\mathbf{J}(\mathbf{q})\ddot{\mathbf{q}} + \dot{\mathbf{J}}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} = 0, \tag{55}$$

where $\mathbf{D}(\mathbf{q}) \in \mathbb{R}^{18\times 18}$ is the mass-inertia matrix, $\mathbf{H}(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^{18}$ contains the Coriolis and gravity terms, $\mathbf{B} \in \mathbb{R}^{18\times 12}$ is the actuation matrix, $\mathbf{J}(\mathbf{q}) = \partial \mathbf{c}(\mathbf{q})/\partial \mathbf{q} \in \mathbb{R}^{c\times 18}$ is the Jacobian of the holonomic constraints, and $\boldsymbol{\lambda} \in \mathbb{R}^{c}$ is the constraint wrench. These full-system dynamics including ground contacts were used in our simulations.

In order to represent the error caused by uncertain terrain, zero mean Gaussian disturbances are added to the quadruped's (x, y) body position and velocity at 1kHz with variances of 2.25×10^{-6} and 0.01 respectively. This noise was chosen to qualitatively match the rough-terrain walking that we have observed in experiments; a video comparing our simulated walking to rough-terrain walking can be found here.

For joint-level torque control, an ID-QP controller designed using concepts in [14] and implemented at 1kHz is used with dynamics (54, 55) to track center-of-mass velocities and angle rates with swing legs following a Reibert-style trajectory in a diagonal walking gait using the motion primitive framework in [32]. We simulate the entire quadruped's dynamics (54, 55), but follow a similar reduced-order-modeling methodology to [24] and consider the following simplified discrete-time single-integrator system for DTCBF-based control:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \Delta t \begin{bmatrix} \cos\theta_k & -\sin\theta_k & 0 \\ \sin\theta_k & \cos\theta_k & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} v_k^x \\ v_k^y \\ \theta_k \end{bmatrix} + \mathbf{d}_k. \tag{56}$$

where $\mathbf{x}_k = \begin{bmatrix} x_k, & y_k, & \theta_k \end{bmatrix}^\top$ and $\Delta t = 0.05$ seconds. Here the i.i.d. random process $\mathbf{d}_k$ models the random disturbances introduced to the planar (x, y) position and velocity as well as the dynamics-mismatch between the full-order quadrupedal dynamics (54, 55) and the simplified model (56).

Using the motion primitive framework presented in [32], the quadruped is commanded to stand and then traverse a 7 meter path that is 1 meter wide, with the safe set $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n \mid 0.5^2 - y^2 \geq 0\}$. For this simulation, three controllers are compared: a simple nominal controller $\mathbf{k}_{\text{nom}}(\mathbf{x}) = \begin{bmatrix} 0.2, & 0, & -\theta \end{bmatrix}^\top$ with no understanding of safety, the DTCBF-OP controller with $\alpha = 0.99$, and our proposed JED controller with $\alpha = 0.99$ and $c_{\text{J}} = \psi$ using the mean and covariance estimates, $\mathbb{E}[\mathbf{d}_k] \approx \begin{bmatrix} -0.0132, & -0.0034, & -0.0002 \end{bmatrix}^\top$ and $\text{tr}(\text{cov}(\mathbf{d}_k)) \approx \psi = 0.000548$, which were estimated using 15 minutes of 20 Hz walking data controlled by $\mathbf{k}_{\text{nom}}$ and which characterize the effect of both the planar disturbances and the model-mismatch between (54, 55) and (56).

The results of 50 trials for each controller can be seen in Figure 1. As expected, $\mathbf{k}_{\text{nom}}$ generated the largest safety violations and JED the smallest and fewest safety violations.

VI. Conclusion

In this work, we developed a bound for the finite-time safety of stochastic discrete-time systems using discrete-time control barrier functions. Additionally, we presented a method for practically implementing convex optimization-based controllers which satisfy this bound by accounting for or analyzing the effect of Jensen's inequality. We presented several examples which demonstrate the efficacy of our bound and our proposed ED and JED controllers,

This paper offers a large variety of directions for future work. In particular, in our practical examples, we find the safety bound presented here is often quite conservative in practice. One way forward would be to find other supermartingale transformations of the process $h(\mathbf{x}_k)$ (perhaps programatically, as in [30]) that can yield tighter bounds than those in Theorem 5. Another potential avenue may consider alternative martingale inequalities to the Ville's inequality used in this work. Another important open question is how to incorporate state uncertainty into our framework. This would allow us to reason about the safety of CBF-based controllers that operate in tandem with state estimators such as Kalman Filters or SLAM pipelines. Similarly, our methods may have interesting applications in handling the dynamics errors introduced in sampled-data control which can perhaps be modeled as a random variable or learned using a distribution-generating framework such as a state-dependent Gaussian processes or Bayesian neural networks. Finally, we assume that the disturbance distribution $\mathcal{D}$ is known exactly, *a priori*; it would be interesting to consider a "distributionally robust" variant of the stochastic barrier condition (40) that can provide safety guarantees for a class of disturbance distributions.

Acknowledgments

The authors would like to thank Alexander De Capone and Victor Dorobantu for their invaluable discussion, Joel Tropp for his course on probability theory and stochastic processes, and Wyatt Ubellacker for generously providing his fantastic quadruped simulation environment.

APPENDIX

A. Lemmas for Theorem 5

The following lemmas are used to prove optimality of the bound in Theorem 5 Cases 1 and 2. These lemmas were originally stated without proof in [21].

Lemma 2. *For $M \in \mathbb{R}_{>0}$, $\gamma, \varphi \in \mathbb{R}_{\geq 0}$, $h(\mathbf{x}_0) \in [-\gamma, M]$, and $K \in \mathbb{N}_{\geq 1}$, the function $\Psi_1 : (1, \infty) \to \mathbb{R}$ defined as:*

$$\Psi_1(\theta) = \frac{M - h(\mathbf{x}_0) + \frac{\varphi\theta}{\theta-1}\left(\theta^K - 1\right)}{(M+\gamma)\theta^K}, \tag{57}$$

is monotonically decreasing.

Proof: The geometric series identity yields:

$$\Psi_1(\theta) = \frac{M - h(\mathbf{x}_0)}{M+\gamma}\theta^{-K} + \frac{\varphi}{(M+\gamma)}\sum_{i=1}^{K}\theta^{i-K}, \tag{58}$$

$$\frac{d\Psi_1}{d\theta} = -\frac{M - h(\mathbf{x}_0)}{M+\gamma}K\theta^{-K-1} - \varphi\sum_{i=1}^{K}\frac{(K-i)\theta^{i-K-1}}{M+\gamma},$$

$$\leq 0, \tag{59}$$

for all $\theta \in (1, \infty)$. ■

Lemma 3. *For $M \in \mathbb{R}_{>0}, \gamma, \varphi \in \mathbb{R}_{\geq 0}$, $h(\mathbf{x}_0) \in [-\gamma, M]$, and $K \in \mathbb{N}_{\geq 1}$, the function $\Psi_2 : (1, \infty) \to \mathbb{R}$ defined as:*

$$\Psi_2(\theta) = 1 - \frac{h(\mathbf{x}_0) + \gamma}{M + \gamma + \frac{\varphi\theta}{\theta-1}\left(\theta^K - 1\right)}, \tag{60}$$

is monotonically increasing.

Proof: The geometric series identity yields:

$$\Psi_2(\theta) = 1 - \frac{h(\mathbf{x}_0) + \gamma}{M + \gamma + \varphi\sum_{i=1}^{K}\theta^i}, \tag{61}$$

$$\frac{d\Psi_2}{d\theta} = \frac{(h(\mathbf{x}_0)+\gamma)\left(\varphi\sum_{i=1}^{K} i\theta^{i-1}\right)}{\left(M+\gamma+\varphi\sum_{i=1}^{K}\theta^i\right)^2}, \tag{62}$$

$$\geq 0, \tag{63}$$

for all $\theta \in (1, \infty)$. ■

B. Lemma 1

Here we present a proof of Lemma 1.

Proof: Consider the convex, twice-continuously differentiable function $\eta : \mathbb{R}^n \to \mathbb{R}$ defined as $\eta = -h$. The intermediate value theorem implies that for all $\mathbf{y}, \mathbf{z} \in \mathbb{R}^n$, there exists an $\omega \in [0, 1]$ such that:

$$\eta(\mathbf{z}) = \eta(\mathbf{y}) + \nabla\eta(\mathbf{y})^\top\mathbf{e} + \frac{1}{2}\mathbf{e}^\top\nabla^2\eta(\mathbf{c})\mathbf{e}, \tag{64}$$

where $\mathbf{e} \triangleq \mathbf{z} - \mathbf{y}$, $\mathbf{c} \triangleq \omega\mathbf{z} + (1-\omega)\mathbf{z}$, and $\nabla^2\eta(\mathbf{c})$ is the Hessian of η evaluated at $\mathbf{c}$. We then have that:

$$\eta(\mathbf{z}) = \eta(\mathbf{y}) + \nabla\eta(\mathbf{y})^\top\mathbf{e} + \frac{1}{2}\mathrm{tr}\left(\nabla^2\eta(\mathbf{c})\mathbf{e}\mathbf{e}^\top\right), \tag{65}$$

$$\leq \eta(\mathbf{y}) + \nabla\eta(\mathbf{y})^\top\mathbf{e} + \frac{1}{2}\|\nabla^2\eta(\mathbf{c})\|_2\mathrm{tr}\left(\mathbf{e}\mathbf{e}^\top\right), \tag{66}$$

$$\leq \eta(\mathbf{y}) + \nabla\eta(\mathbf{y})^\top\mathbf{e} + \frac{\lambda_{\max}}{2}\mathrm{tr}\left(\mathbf{e}\mathbf{e}^\top\right), \tag{67}$$

where the first inequality is a property of the trace operator for positive semi-definite matrices [28] (and $\nabla^2\eta(\mathbf{c})$ is positive semi-definite as η is convex), and the second inequality follows by our definition of $\lambda_{\max}$. Let $\mathbf{x}$ be a random variable taking values in $\mathbb{R}^n$ with probability density function $p : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$, and let $\boldsymbol{\mu} \triangleq \mathbb{E}[\mathbf{x}]$. We then have that:

$$\mathbb{E}[\eta(\mathbf{x})] - \eta(\mathbb{E}[\mathbf{x}]) = \int_{\mathbb{R}^n}(\eta(\mathbf{x}) - \eta(\boldsymbol{\mu}))p(\mathbf{x})d\mathbf{x}, \tag{68}$$

$$\leq \int_{\mathbb{R}^n}\nabla\eta(\boldsymbol{\mu})^\top\mathbf{e} + \frac{\lambda_{\max}}{2}\mathrm{tr}\left(\mathbf{e}\mathbf{e}^\top\right)p(\mathbf{x})d\mathbf{x}, \tag{69}$$

$$= \frac{\lambda_{\max}}{2}\mathrm{tr}(\mathrm{cov}(\mathbf{x})), \tag{70}$$

where $\mathbf{e} = \mathbf{x} - \boldsymbol{\mu}$. Replacing η with $-h$ yields:

$$\mathbb{E}[h(\mathbf{x})] \geq h(\mathbb{E}[\mathbf{x}]) - \frac{\lambda_{\max}}{2}\mathrm{tr}(\mathrm{cov}(\mathbf{x})). \tag{71}$$

■

C. Derivation of Convex Approximation for Polytopic Barrier

Here we derive a conservative approximation of the constraint $\mathbb{E}[h(\mathbf{x}_{k+1})] \geq \alpha h(\mathbf{x}_k)$ for barriers of the form $h(\mathbf{x}) = -\max(\mathbf{C}\mathbf{x} - \mathbf{w})$ and systems with linear-Gaussian dynamics (53). The key idea is to use the *log-sum-exp* function as a smooth, convex upper bound of the pointwise maximum in the barrier function, which yields a closed-form expression for Gaussian random variables.

In particular, if L is the *log-sum-exp* function, for any $t > 0$, $\max(\mathbf{x}) \leq \frac{1}{t}L(t\mathbf{x}) \triangleq \frac{1}{t}\log(\sum_{i=1}^{n}\exp(tx_i))$ [13, Chapter 3]. We can use this to upper bound the expectation of $-h$,

$$-\mathbb{E}[h(\mathbf{x}_{k+1})] = \mathbb{E}\left[\max(\mathbf{C}\mathbf{x}_{k+1} - \mathbf{w})\right] \tag{72}$$

$$\leq \frac{1}{t}\mathbb{E}\left[L\Big(t(\mathbf{C}\mathbf{x}_{k+1} - \mathbf{w})\Big)\right] \tag{73}$$

$$\leq \frac{1}{t}\log\left(\sum_{i=1}^{n_c}\mathbb{E}\left[\exp(t\mathbf{r}_i)\right]\right), \tag{74}$$

for $\mathbf{r}_i \triangleq \mathbf{c}_i^T\mathbf{x} - w_i$, where $\mathbf{c}_i$ is the i^{th} row of $\mathbf{C}$, w_i is the i^{th} entry of $\mathbf{w}$, and the last inequality follows from Jensen's inequality and the concavity of the natural logarithm. Further, since we have linear-Gaussian dynamics, it is easy to show that $\mathbf{r}_i \sim \mathcal{N}(\mathbf{c}_i^T(\mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k) - w_i, \mathbf{c}_i^T\mathbf{Q}\mathbf{c}_i)$. The expression $\mathbb{E}[\exp(t\mathbf{X})]$ is the "moment-generating function" of a random variable $\mathbf{X}$, and for a Gaussian r.v. $\mathbf{X} \sim \mathcal{N}(\mu, \sigma^2)$, it has a closed form, $\mathbb{E}[\exp(t\mathbf{X})] = \exp(t\mu + \frac{t^2}{2}\sigma^2)$ [34, Chapter 6].

Thus, for $\boldsymbol{\mu} \triangleq \mathbf{C}(\mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k) - \mathbf{w}$, $\boldsymbol{\sigma} \triangleq \mathrm{diag}(\mathbf{A}\mathbf{Q}\mathbf{A}^T)$, where $\mathrm{diag}(\cdot)$ defines the diagonal of a square matrix,

$$-\mathbb{E}[h(\mathbf{x}_{k+1})] \leq \frac{1}{t}L\left(t\boldsymbol{\mu} + \frac{t^2}{2}\boldsymbol{\sigma}\right), \tag{75}$$

which implies that imposing the constraint $\frac{1}{t}L(t\boldsymbol{\mu} + \frac{t^2}{2}\boldsymbol{\sigma}) \leq -\alpha h(\mathbf{x}_k)$ ensures that the stochastic barrier condition (40) is satisfied. Finally, recognizing that our constraint is a perspective transform of $L(\boldsymbol{\mu} + \frac{t}{2}\boldsymbol{\sigma})$ by the scalar $\frac{1}{t}$, which preserves convexity [13, Chapter 3], our constraint is indeed convex. Thus an optimization-based controller such as ED can be used online to select control actions, and can jointly optimize over $\mathbf{u}_k, t$ to obtain the tightest bound on the expectation possible.

References

[1] Ayush Agrawal and Koushil Sreenath. Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation. In *Robotics: Science and Systems*, volume 13. Cambridge, MA, USA, 2017.

[2] Devansh R. Agrawal, Hardik Parwana, Ryan K. Cosner, Ugo Rosolia, Aaron D. Ames, and Dimitra Panagou. A Constructive Method for Designing Safe Multirate Controllers for Differentially-Flat Systems. *IEEE Control Systems Letters*, 6:2138–2143, 2022. ISSN 2475-1456. doi: 10.1109/LCSYS.2021.3136465. URL https://ieeexplore.ieee.org/document/9655322/.

[3] Mohamadreza Ahmadi, Andrew Singletary, Joel W Burdick, and Aaron D Ames. Safe policy synthesis in multi-agent pomdps via discrete-time barrier functions. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 4797–4803. IEEE, 2019.

[4] Mohamadreza Ahmadi, Xiaobin Xiong, and Aaron D. Ames. Risk-Averse Control via CVaR Barrier Functions: Application to Bipedal Robot Locomotion. *IEEE Control Systems Letters*, 6:878–883, 2022. ISSN 2475-1456. doi: 10.1109/LCSYS.2021.3086854. Conference Name: IEEE Control Systems Letters.

[5] Anil Alan, Andrew J Taylor, Chaozhe R He, Gábor Orosz, and Aaron D Ames. Safe controller synthesis with tunable input-to-state safe control barrier functions. *IEEE Control Systems Letters*, 6:908–913, 2021.

[6] Aaron D Ames, Xiangru Xu, Jessy W Grizzle, and Paulo Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8):3861–3876, 2016.

[7] Aaron D. Ames, Samuel Coogan, Magnus Egerstedt, Gennaro Notomista, Koushil Sreenath, and Paulo Tabuada. Control Barrier Functions: Theory and Applications. In *2019 18th European Control Conference (ECC)*, pages 3420–3431, June 2019. doi: 10.23919/ECC.2019.8796030.

[8] Jean-Pierre Aubin, Alexandre M Bayen, and Patrick Saint-Pierre. *Viability theory: new directions*. Springer Science & Business Media, 2011.

[9] Somil Bansal, Mo Chen, Sylvia Herbert, and Claire J Tomlin. Hamilton-jacobi reachability: A brief overview and recent advances. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 2242–2253. IEEE, 2017.

[10] Robert A Becker. The variance drain and jensen's inequality. 2012.

[11] Franco Blanchini and Stefano Miani. *Set-theoretic methods in control*, volume 78. Springer, 2008.

[12] Nicoletta Bof, Ruggero Carli, and Luca Schenato. Lyapunov theory for discrete time systems. *arXiv preprint arXiv:1809.05289*, 2018.

[13] Stephen P. Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2014. ISBN 978-0-521-83378-3. doi: 10.1017/CBO9780511804441. URL https://web.stanford.edu/%7Eboyd/cvxbook/.

[14] Jonas Buchli, Mrinal Kalakrishnan, Michael Mistry, Peter Pastor, and Stefan Schaal. Compliant quadruped locomotion over rough terrain. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 814–820, 2009. doi: 10.1109/IROS.2009.5354681.

[15] Richard Cheng, Gábor Orosz, Richard M Murray, and Joel W Burdick. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3387–3395, 2019.

[16] Albert Chern, Xiang Wang, Abhiram Iyer, and Yorie Nakahira. Safe Control in the Presence of Stochastic Uncertainties. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 6640–6645, December 2021. doi: 10.1109/CDC45484.2021.9683542. ISSN: 2576-2370.

[17] Andrew Clark. Control Barrier Functions for Complete and Incomplete Information Stochastic Systems. In *2019 American Control Conference (ACC)*, pages 2928–2935, July 2019. doi: 10.23919/ACC.2019.8814901. ISSN: 2378-5861.

[18] Geoffrey Grimmett and David Stirzaker. *Probability and random processes*. Oxford university press, 2020.

[19] Lukas Hewing, Kim P Wabersich, Marcel Menner, and Melanie N Zeilinger. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:269–296, 2020.

[20] Shishir Kolathaya and Aaron D Ames. Input-to-state safety with control barrier functions. *IEEE control systems letters*, 3(1):108–113, 2018.

[21] Harold J Kushner. Stochastic stability and control. Technical report, Brown Univ Providence RI, 1967.

[22] JG Liao and Arthur Berg. Sharpening jensen's inequality. *The American Statistician*, 2018.

[23] Shuo Liu, Jun Zeng, Koushil Sreenath, and Calin A Belta. Iterative convex optimization for model predictive control with discrete-time high-order control barrier functions. *arXiv preprint arXiv:2210.04361*, 2022.

[24] Tamas G. Molnar, Ryan K. Cosner, Andrew W. Singletary, Wyatt Ubellacker, and Aaron D. Ames. Model-Free Safety-Critical Control for Robotic Systems. *IEEE Robotics and Automation Letters*, 7(2):944–951, April 2022. ISSN 2377-3766, 2377-3774. doi: 10.1109/LRA.2021.3135569. URL https://ieeexplore.ieee.org/document/9652122/.

[25] Antonis Papachristodoulou and Stephen Prajna. A tutorial on sum of squares techniques for systems analysis. In *Proceedings of the 2005, American Control Conference, 2005.*, pages 2686–2700. IEEE, 2005.

[26] Stephen Prajna, Ali Jadbabaie, and George J Pappas. Stochastic safety verification using barrier certificates. In *2004 43rd IEEE conference on decision and control (CDC)(IEEE Cat. No. 04CH37601)*, volume 1, pages 929–934. IEEE, 2004.

[27] Cesar Santoyo, Maxence Dutreix, and Samuel Coogan. Verification and Control for Finite-Time Safety of Stochastic Systems via Barrier Functions, May 2019. URL http://arxiv.org/abs/1905.12077. arXiv:1905.12077 [cs].

[28] Khalid Shebrawi and Hussien Albadawi. Trace inequalities for matrices. *Bulletin of the Australian Mathematical Society*, 87(1):139–148, 2013.

[29] Eduardo D Sontag. Input to state stability: Basic concepts and results. In *Nonlinear and optimal control theory*, pages 163–220. Springer, 2008.

[30] Jacob Steinhardt and Russ Tedrake. Finite-time regional verification of stochastic non-linear systems. *The International Journal of Robotics Research*, 31(7):901–923, 2012.

[31] Andrew J Taylor, Victor D Dorobantu, Ryan K Cosner, Yisong Yue, and Aaron D Ames. Safety of sampled-data systems with control barrier functions via approximate discrete time models, 2022.

[32] Wyatt Ubellacker, Noel Csomay-Shanklin, Tamas G. Molnar, and Aaron D. Ames. Verifying safe transitions between dynamic motion primitives on legged robots. *arXiv preprint*, (arXiv:2106.10310), 2021.

[33] Jean Ville. Etude critique de la notion de collectif. *Bull. Amer. Math. Soc*, 45(11):824, 1939.

[34] D. Wackerly, W. Mendenhall, and R.L. Scheaffer. *Mathematical Statistics with Applications*. Cengage Learning, 2014. ISBN 978-1-111-79878-9.

[35] Adrian G Wills and William P Heath. Barrier function based model predictive control. *Automatica*, 40(8):1415–1422, 2004.

[36] Jun Zeng, Bike Zhang, and Koushil Sreenath. Safety-critical model predictive control with discrete-time control barrier function. In *2021 American Control Conference (ACC)*, pages 3882–3889. IEEE, 2021.

[37] Kemin Zhou and John Comstock Doyle. *Essentials of robust control*, volume 104. Prentice hall Upper Saddle River, NJ, 1998.