

Bounding Stochastic Safety: Leveraging Freedman's Inequality With Discrete-Time Control Barrier Functions

Ryan K. Cosner^{1b}, Graduate Student Member, IEEE, Preston Culbertson^{1b},
and Aaron D. Ames^{1b}, Fellow, IEEE

Abstract—When deployed in the real world, safe control methods must be robust to unstructured uncertainties such as modeling error and external disturbances. Typical robust safety methods achieve their guarantees by always assuming that the worst-case disturbance will occur. In contrast, this letter utilizes Freedman's inequality in the context of discrete-time control barrier functions (DTCBFs) and c-martingales to provide stronger (less conservative) safety guarantees for stochastic systems. Our approach accounts for the underlying disturbance distribution instead of relying exclusively on its worst-case bound and does not require the barrier function to be upper-bounded, which makes the resulting safety probability bounds more useful for intuitive safety constraints such as signed distance. We compare our results with existing safety guarantees, such as input-to-state safety (ISSf) and martingale results that rely on Ville's inequality. When the assumptions for all methods hold, we provide a range of parameters for which our guarantee is stronger. Finally, we present simulation examples, including a bipedal walking robot, that demonstrate the utility and tightness of our safety guarantee.

Index Terms—Constrained control, Lyapunov methods, robotics, stochastic systems, uncertain systems.

I. INTRODUCTION

SAFETY—characterized as the forward-invariance of a safe set [1]—has become a popular area of study within control theory, with broad applications to autonomous vehicles, medical and assistive robotics, aerospace systems, and beyond. Ensuring safety for these systems requires one to account for unpredictable, real-world effects. Historically, control theory has treated the problem of safety under uncertainty using deterministic methods, often seeking safety guarantees in the presence of bounded disturbances. This problem has been studied using a variety of safe control approaches including control barrier functions (CBFs) [2], Hamilton-Jacobi (HJ) reachability [3], and state-constrained model-predictive control (MPC) [4]. However, this worst-case analysis often leads

Manuscript received 15 March 2024; revised 19 May 2024; accepted 29 May 2024. Date of publication 3 June 2024; date of current version 16 July 2024. This work was supported in part by BP and in part by NSF CPS under Award 1932091. Recommended by Senior Editor V. Ugrinovskii. (Corresponding author: Ryan K. Cosner.)

The authors are with the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 9125 USA (e-mail: rkcosner@caltech.edu; pbulbert@caltech.edu; ames@caltech.edu).

Digital Object Identifier 10.1109/LCSYS.2024.3409105

2475-1456 © 2024 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See <https://www.ieee.org/publications/rights/index.html> for more information.




Fig. 1. Safety results for a bipedal robot navigating around an obstacle. Details are provided in Section IV. (Top) Visualization of the Hybrid Linear Inverted Pendulum (HLIP) model. Yellow indicates the center-of-mass (COM), blue is the stance foot, and red is the swing foot. The states x_k are the global COM position, the relative COM position, and COM velocity, and the input is the relative position of the feet at impact. (Bottom) A table with variable maximum disturbance value ($d_{\max}$) and controller parameter (α) shows our (dashed lines) theoretical bound on safety failure from Thm. 3, (dotted lines) the shortest first-violation time based on the worst-case disturbance approximation, and (solid lines) approximated probabilities from 5000 trials (lower is safer). On the left, the trajectories of the COM are shown walking from bottom left towards the top right while avoiding the obstacle with each color corresponding to a different $d_{\max}$. The robot attempts to avoid the obstacle (black).

to conservative performance since it ensures robustness to adversarial disturbances which are uncommon in practice.

Stochastic methods provide an alternative to the worst-case bounding approach. Instead of a conservative uncertainty bound, these methods consider a distribution of possible disturbances. Although they do not provide absolute, risk-free safety guarantees, they allow for smooth degradation of safety via variable, risk-aware levels of conservatism. A wide variety of stochastic safety methods exist including: reachability-based optimal safety [6], [7], constrained coherent risk measures [8], sampling-based general risk measures [9], and martingale-based methods [10], [11] amongst many others. In this letter, we will focus on martingale-based methods due to their ability to generate trajectory-long guarantees and their relative simplicity as a method which relies primarily on only a distribution's first-moment.

Continuous-time martingale-based stochastic safety methods have successfully achieved strong probabilistic safety guarantees [12], [13], [14], [15], but require controllers with functionally infinite bandwidth, a strong assumption for real-world systems with discrete-time sensing and actuation. Alternatively, discrete-time methods have shown success while also capturing the sampled-data complexities of real-world systems [10], [16], [17], [18]. In this letter we focus on extending the theory of discrete-time martingale-based stochastic safety involving discrete-time control barrier functions (DTCBFs) and c -martingales.

The stochastic discrete-time martingale-based safety literature has shown theoretical success [10], [11], [14], [16], [18] in generating risk-based safety guarantees and in deploying these guarantees to real-world systems [19]. In this letter we seek to extend these existing safety techniques by utilizing a different (and often stronger) concentration inequality that can provide sharper safety probability bounds. Where other works have traditionally relied on Ville's inequality [20], we instead turn to Freedman's inequality [21]. By assuming that the martingale differences and predictable quadratic variation are bounded, this inequality relaxes the nonnegativity assumption required by Ville's inequality while also providing generally tighter bounds that degrade smoothly with increasing uncertainty.

This letter combines discrete-time martingale-based safety techniques with Freedman's inequality to obtain tighter bounds on stochastic safety. We make three key contributions: (1) introducing Freedman-based safety probabilities for DTCBFs and c -martingales, (2) providing a range of parameter values where our bound is tighter than existing discrete-time martingale-based safety results, and (3) validating our method in simulation. We apply our results to bipedal obstacle avoidance (Fig. 1), using a reduced-order model of the step-to-step dynamics. This case study shows the utility of our probability bounds, which decay smoothly with increasing uncertainty and enable non-conservative, collision avoidance.

II. BACKGROUND

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let $\mathcal{F}_0 \subset \mathcal{F}_1 \subset \dots \subset \mathcal{F}$ be a filtration of $\mathcal{F}$. Consider discrete-time dynamical systems of the form:

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

where $\mathbf{x}_k \in \mathbb{R}^n$ is the state, $\mathbf{u}_k \in \mathbb{R}^m$ is the input, $\mathbf{d}_k$ is an $\mathcal{F}_{k+1}$ measurable random disturbance which takes values in $\mathbb{R}^\ell$, and $\mathbf{F} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{R}^\ell \rightarrow \mathbb{R}^n$ is the dynamics. Throughout this letter we assume that all random variables and functions of random variables are integrable.

To create a closed-loop system, we add a state-feedback controller $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$:

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

The goal of this letter is to provide probabilistic safety guarantees for this closed-loop system.

A. Safety and Discrete-Time Control Barrier Functions

To make guarantees regarding the safety of system (2), we first formalize our notion of safety as the forward invariance of a user-defined “safe set”, $\mathcal{C} \subset \mathbb{R}^n$, as is common in the robotics and control literature [1], [3], [4], [22].

Definition 1 (Forward Invariance and Safety)¹: A set $\mathcal{C} \subset \mathbb{R}^n$ is forward invariant for system (2) if $\mathbf{x}_0 \in \mathcal{C} \implies \mathbf{x}_k \in \mathcal{C}$

¹For this letter we will focus on the safety of (2) exclusively at samples times as in [4] and [23]. We refer to [24] for an analysis of intersample safety.

for all $k \in \mathbb{N}$. We define “safety” with respect to $\mathcal{C}$ as the forward invariance of $\mathcal{C}$.

One method for ensuring safety is through the use of Discrete-Time Control Barrier Functions (DTCBFs). For DTCBFs, we consider safe sets that are 0-superlevel sets [1] of a continuous function $h : \mathbb{R}^n \rightarrow \mathbb{R}$:

$$\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n \mid h(\mathbf{x}) \geq 0\}. \quad (3)$$

In particular the DTCBF is defined as:

Definition 2 (Discrete-Time Control Barrier Function (DTCBF)) [23]: Let $\mathcal{C} \subset \mathbb{R}^n$ be the 0-superlevel set of some function $h : \mathbb{R}^n \rightarrow \mathbb{R}$. The function h is a DTCBF for $\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{u}, \mathbf{0})$ if there exists an $\alpha \in [0, 1]$ such that:

$$\sup_{\mathbf{u} \in \mathbb{R}^m} h(\mathbf{F}(\mathbf{x}, \mathbf{u}, \mathbf{0})) > \alpha h(\mathbf{x}), \quad \forall \mathbf{x} \in \mathcal{C} \quad (4)$$

DTCBFs differ from their continuous-time counterparts in that they satisfy an inequality constraint on their *finite difference* instead of their derivative.² On the other hand, they are similar in their ability to create *safety filters* for nominal controllers $\mathbf{k}_{\text{nom}} : \mathbb{R}^n \times \mathbb{Z} \rightarrow \mathbb{R}^m$ of the form:

$$\begin{aligned} \mathbf{k}(\mathbf{x}) = \operatorname{argmin}_{\mathbf{u} \in \mathbb{R}^m} \quad & \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x}, k)\|^2 \\ \text{s.t.} \quad & h(\mathbf{F}(\mathbf{x}, \mathbf{u}, \mathbf{0})) \geq \alpha h(\mathbf{x}). \end{aligned} \quad (5)$$

Assuming feasibility, $\mathbf{k}(\mathbf{x})$ guarantees safety for the undisturbed system by selecting inputs that satisfy (4) [23, Pro. 1].

For deterministic systems, infinite-horizon safety guarantees are common. However, for discrete-time stochastic systems, when the disturbance is bounded, infinite horizon guarantees fail to capture the nuances of variable risk levels and, when the disturbance is unbounded, infinite-horizon guarantees can be impossible to achieve³ [25, Sec. IV]. In order to provide an achievable risk-based guarantee we choose to analyze finite-time safety probabilities as in [10], [11], [14], [15] instead of infinite-time safety guarantees.

Definition 3 (K-step Exit Probability): For any $K \in \mathbb{N}_1$ and initial condition $\mathbf{x}_0 \in \mathbb{R}^n$, the K -step exit probability of the set $\mathcal{C}$ for the closed-loop system (2) is:

$$P_u(K, \mathbf{x}_0) = \mathbb{P}\{\mathbf{x}_k \notin \mathcal{C} \text{ for some } k \leq K\} \quad (6)$$

This describes the probability that the system will leave the safe set $\mathcal{C}$ within K time steps given that it started at $\mathbf{x}_0$.

B. Existing Martingale-Based Safety Methods

In this letter, we will generate bounds on K -step exit probabilities using martingale-based concentration inequalities. Martingales are a class of stochastic processes which satisfy a relationship between their mean and previous value.

Definition 4 (Martingale) [10], [26]: Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space with a filtration $\{\mathcal{F}_0, \mathcal{F}_1, \dots, \mathcal{F}\}$. A stochastic process W_k that is adapted to the filtration and is integrable at each k is a martingale if

$$\mathbb{E}[W_{k+1} \mid \mathcal{F}_k] = W_k, \quad \forall k \in \mathbb{Z} \quad \text{a.s.} \quad (7)$$

Additionally, if W_k satisfies:

$$\mathbb{E}[W_{k+1} \mid \mathcal{F}_k] \leq W_k + c, \quad \forall k \in \mathbb{Z} \quad (\text{a.s.}), \quad (8)$$

²The standard continuous-time CBF condition $\dot{h}(\mathbf{x}) \leq -\bar{\gamma}h(\mathbf{x})$ for $\bar{\gamma} > 0$ becomes $h(\mathbf{x}_{k+1}) - h(\mathbf{x}_k) \geq -\gamma h(\mathbf{x}_k)$ for $\gamma \in [0, 1]$ in discrete-time; defining $\alpha = 1 - \gamma$ recovers the condition $h(\mathbf{x}_{k+1}) \geq \alpha h(\mathbf{x}_k)$.

³Consider the system: $\mathbf{x}_{k+1} = \mathbf{u}_k + \mathbf{d}_k$, where $\mathbf{x} \in \mathbb{R}$, $\mathbf{u} \in \mathbb{R}$, $\mathbf{d} \sim \mathcal{N}(0, 1)$, and $\mathcal{C} = \{\mathbf{x} \in \mathbb{R} \mid |\mathbf{x}| < 1\}$. At every time step, $\mathbb{P}\{\mathbf{x}_{k+1} \in \mathcal{C}\}$ is maximized with $\mathbf{u}_k = 0$, but then even over a single discrete step, there is at least 30% chance of failure. As time continues, this constant risk of failure at every step makes infinite horizon guarantees impossible to achieve.

with $c = 0$ then it is a supermartingale and if it satisfies (8) with $c \geq 0$ then it is a c -martingale.

Many concentration inequalities can be used to bound the spread of a martingale over time. One particularly useful bound is Ville's [20] which bounds the probability that a supermartingale W_k rises above a threshold $\lambda > 0$.

Lemma 1 (Ville's Inequality [20]): If W_k is a nonnegative supermartingale, then for all $\lambda > 0$,

$$\mathbb{P}\left\{\sup_{k \in \mathbb{Z}} W_k > \lambda\right\} \leq \frac{\mathbb{E}[W_0]}{\lambda} \quad (9)$$

Critically, Ville's inequality assumes *nonnegativity* which manifests as a requirement that h be upper-bounded, e.g., (10).

For safety applications of Ville's inequality, we consider the case where $h(\mathbf{x}_k)$ is upper bounded by $B > 0$ and satisfies one of the following expectation conditions

$$\mathbb{E}[h(\mathbf{F}(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k), \mathbf{d}_k)) \mid \mathcal{F}_k] \geq \alpha h(\mathbf{x}_k), \quad (\text{DTCBF})$$

$$\mathbb{E}[h(\mathbf{F}(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k), \mathbf{d}_k)) \mid \mathcal{F}_k] \geq h(\mathbf{x}_k) - c, \quad (c\text{-mart.})$$

for some $\alpha \in (0, 1)$ or $c \geq 0$. In this case, we can achieve the following bound on the K -step exit probability, $P_u(K, \mathbf{x}_0)$:

Theorem 1 (Safety Using Ville's Inequality, [10], [11], [14], [16])⁴: If, for some $B > 0$ and $K \in \mathbb{N}_1$, the function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ satisfies:

$$h(\mathbf{x}) \leq B, \quad \text{for all } \mathbf{x} \in \mathbb{R}^n, \quad (10)$$

$$\text{then:} \quad P_u(K, \mathbf{x}_0) \leq 1 - \frac{\lambda}{B}, \quad (11)$$

$$\text{where } \lambda = \begin{cases} \alpha^K h(\mathbf{x}_0), & \text{if (2) satisfies (DTCBF) } \forall k \leq K \\ h(\mathbf{x}_0) - cK, & \text{if (2) satisfies (c-mart.) } \forall k \leq K. \end{cases}$$

This guarantees that the risk of the becoming unsafe is upper bounded by a function which decays to 1 with time and which depends on the system's initial safety "fraction", $h(\mathbf{x}_0)/B$.

III. SAFETY GUARANTEES USING FREEDMAN'S INEQUALITY

This section presents our main result: K -step exit probability bounds for DTCBFS and c -martingales generated using Freedman's inequality, a particularly strong martingale concentration inequality. Here, we use the simpler, historical version as presented by Freedman [21]; see [28] for historical context and a new, tighter alternative which could also be used. After presenting this result, this section explores comparisons with existing Ville's-based methods and input-to-state safety.

Before presenting Freedman's inequality, we must define the predictable quadratic variation (PQV) of a process which is a generalization of variance for stochastic processes.

Definition 5 (Predictable Quadratic Variation (PQV) [26]): The PQV of a martingale W_k at $K \in \mathbb{N}_1$ is:

$$\langle W \rangle_K \triangleq \sum_{i=1}^K \mathbb{E}[(W_i - W_{i-1})^2 \mid \mathcal{F}_{i-1}] \quad (12)$$

Unlike Ville's inequality, Freedman's inequality does not require nonnegativity of the martingale W_k , thus removing the upper-bound requirement (10) on h . In place of nonnegativity, we require two alternative assumptions:

Assumption 1 (Upper-Bounded Differences): We assume that the martingale differences are upper-bounded by 1 (i.e., $W_{k+1} - W_k \leq 1$, similar to Azuma-Hoeffding methods [26]).

Assumption 2 (Bounded PQV): We assume that the PQV is upper-bounded by $\xi^2 > 0$.

Given the PQV of the process, Freedman's inequality provides the following bound:

Theorem 2 (Freedman's Inequality [21, Th. 4.1]): If, for some $K \in \mathbb{N}_1$ and $\xi > 0$, W_k is a supermartingale with $W_0 = 0$ such that:

$$(W_k - W_{k-1}) \leq 1 \text{ for all } k \leq K, \quad (\text{Assumption 1})$$

$$\langle W \rangle_K \leq \xi^2, \quad (\text{Assumption 2})$$

then, for any $\lambda \geq 0$,

$$\mathbb{P}\left\{\max_{k \leq K} W_k \geq \lambda\right\} \leq H(\lambda, \xi) \triangleq \left(\frac{\xi^2}{\lambda + \xi^2}\right)^{\lambda + \xi^2} e^\lambda. \quad (13)$$

To prove this, we first bound the moment generating function $\mathbb{E}[e^{\gamma X}] \leq e^{(\gamma^2 - 1 - \gamma)\text{Var}(X)}$ for $\gamma \geq 0$ and 0-mean random variables X bounded by 1. By choosing X as the martingale differences, this is used to prove that $\int_{\{\tau < \infty\}} e^{\gamma W_\tau - (\gamma^2 - 1 - \gamma)\langle W \rangle_\tau} d\mathbb{P} \leq 1$ for W_k satisfying assp. (1) where τ any stopping time [21, Prop. 3.3]. Bounding with Assp. (2) and optimizing γ then yields the desired bound. A restatement of this proof can be found in [27, Appx. D].

A. Main Result: Freedman's Inequality for Safety

Next we present the key contribution of this letter: the application of Freedman's inequality to stochastic safety.

Theorem 3: If, for some $K \in \mathbb{N}_1$, $\sigma > 0$, and $\delta > 0$, the following bounds⁵ on the difference⁶ between the true and predictable update (14) and the conditional variance (15) hold for all $k \leq K$:

$$\mathbb{E}[h(\mathbf{x}_k) \mid \mathcal{F}_{k-1}] - h(\mathbf{x}_k) \leq \delta, \quad (14)$$

$$\text{Var}(h(\mathbf{x}_{k+1}) \mid \mathcal{F}_k) \leq \sigma^2, \quad (15)$$

then the K -step exit probability is bounded as:

$$P_u(K, \mathbf{x}_0) \leq H\left(\frac{\lambda}{\delta}, \frac{\sigma\sqrt{K}}{\delta}\right), \quad (16)$$

$$\text{where } \lambda = \begin{cases} \alpha^K h(\mathbf{x}_0), & \text{if (2) satisfies (DTCBF) } \forall k \leq K, \\ h(\mathbf{x}_0) - cK, & \text{if (2) satisfies (c-mart.) } \forall k \leq K. \end{cases}$$

To apply Thm. 2 to achieve Thm. 3 we follow this proof structure: **(Step 1)** normalize h and use it to construct a candidate supermartingale W_k , **(Step 2)** verify that W_k is indeed a supermartingale with $W_0 = 0$, **(Step 3)** use Doob's decomposition [26, Th. 12.1.10] to produce a martingale M_k from W_k in order to remove the negative effect of safe, predictable jumps from the PQV, **(Step 4)** verify that M_k satisfies Assp.s 1 and 2, **(Step 5)** choose $\lambda \geq 0$ such that a safety failure implies $\{\max_{k \leq K} W_k \geq \lambda\}$ as in (13), and **(Step 6)** specialize to specific values of α and c for each case.

Proof (Step 1): Consider the case, for $\tilde{\alpha} \in (0, 1]$ and $\tilde{c} \geq 0$,

$$\text{where } \mathbb{E}[h(\mathbf{x}_{k+1}) \mid \mathcal{F}_k] \geq \tilde{\alpha} h(\mathbf{x}_k) - \tilde{c}, \quad \text{for all } k \leq K. \quad (17)$$

⁴For a discussion notational differences between this presentation of Thm. 1 and that in [17] and [10]. Also, see [16, Th. 5] for probability bounds associated with the general condition $\mathbb{E}[h(\mathbf{F}(\mathbf{x}_k, \mathbf{k}(\mathbf{x}_k), \mathbf{d}_k)) \mid \mathcal{F}_k] \geq \alpha h(\mathbf{x}_k) - c$ using Ville's inequality.

⁵Only upper-bounds on δ and σ^2 are required for (16) to hold and this guarantee is robust to changes in distribution that still satisfy (14) and (15). For real-world systems, distribution-learning can be employed, similar to [19].

⁶See [27, Appx. G] for a constructive method for determining δ and σ .

First, define the normalized safety function $\eta(\mathbf{x}) \triangleq \frac{h(\mathbf{x})}{\delta}$ to ensure that the martingale differences will be bounded by 1. Next, use η to define the candidate supermartingale⁷

$$W_k \triangleq -\tilde{\alpha}^{K-k}\eta(\mathbf{x}_k) + \tilde{\alpha}^K\eta(\mathbf{x}_0) - \sum_{i=1}^k \tilde{\alpha}^{K-i}\frac{\tilde{c}}{\delta} \quad (18)$$

(Step 2) This satisfies⁸ $W_0 = 0$ and is a supermartingale:

$$\begin{aligned} \mathbb{E}[W_{k+1}|\mathcal{F}_k] &= -\tilde{\alpha}^{K-(k+1)}\mathbb{E}[\eta(\mathbf{x}_{k+1})|\mathcal{F}_k] + \tilde{\alpha}^K\eta(\mathbf{x}_0) - \sum_{i=1}^{k+1} \tilde{\alpha}^{K-i}\frac{\tilde{c}}{\delta}, \\ &\leq -\tilde{\alpha}^{K-k}\eta(\mathbf{x}_k) + \tilde{\alpha}^K\eta(\mathbf{x}_0) - \sum_{i=1}^k \tilde{\alpha}^{K-i}\frac{\tilde{c}}{\delta} = W_k. \end{aligned} \quad (19)$$

which can be seen by applying the bound from (17).

(Step 3) The martingale from Doob's decomposition is:

$$\begin{aligned} M_k &\triangleq W_k + \sum_{i=1}^k (W_{i-1} - \mathbb{E}[W_i|\mathcal{F}_{i-1}]), \\ &= W_k + \underbrace{\sum_{i=1}^k \frac{\tilde{\alpha}^{K-i}}{\delta} (\mathbb{E}[h(\mathbf{x}_i)|\mathcal{F}_{i-1}] - \tilde{\alpha}h(\mathbf{x}_{i-1}) + \tilde{c})}_{\geq 0} \geq W_k \end{aligned} \quad (20)$$

where the bound comes from (17) and positivity of $\tilde{\alpha}$ and δ .

(Step 4) Furthermore, M_k satisfies Assp. 1:

$$\begin{aligned} M_k - M_{k-1} &= W_k - \mathbb{E}[W_k|\mathcal{F}_{k-1}], \\ &= \tilde{\alpha}^{K-k}(\mathbb{E}[\eta(\mathbf{x}_k)|\mathcal{F}_{k-1}] - \eta(\mathbf{x}_k)) \leq \tilde{\alpha}^{K-k}\frac{\delta}{\delta} \leq 1, \end{aligned} \quad (21)$$

since we assume in (14) that $\mathbb{E}[h(\mathbf{x}_k) | \mathcal{F}_{k-1}] - h(\mathbf{x}_k) \leq \delta$.

Next, $\tilde{\alpha} \in (0, 1]$ and (15) ensure that M_k satisfies Assp. 2:

$$\begin{aligned} \langle M \rangle_K &= \sum_{i=1}^K \mathbb{E}[\tilde{\alpha}^{2(K-i)}(\eta(\mathbf{x}_i) - \mathbb{E}[\eta(\mathbf{x}_i)|\mathcal{F}_{i-1}])^2|\mathcal{F}_{i-1}] \\ &= \sum_{i=1}^K \frac{\tilde{\alpha}^{2(K-i)}}{\delta^2} \text{Var}(h(\mathbf{x}_i)|\mathcal{F}_{i-1}) \leq \sum_{i=1}^K \tilde{\alpha}^{2(K-i)} \frac{\sigma^2}{\delta^2} \\ &\leq \frac{\sigma^2 K}{\delta^2}. \end{aligned} \quad (22)$$

(Step 5) Now, to relate the unsafe event $\{\min_{k \leq K} h(\mathbf{x}_k) < 0\}$ to our martingale M_k we consider the implications:

$$\min_{k \leq K} h(\mathbf{x}_k) < 0 \implies \min_{k \leq K} h(\mathbf{x}_k) \leq 0 \quad (25)$$

$$\iff \max_{k \leq K} -\tilde{\alpha}^{K-k}\eta(\mathbf{x}_k) \geq 0, \quad \text{since } \tilde{\alpha} > 0, \delta > 0 \quad (26)$$

$$\iff \max_{k \leq K} W_k \geq \tilde{\alpha}^K\eta(\mathbf{x}_0) - \sum_{i=1}^k \tilde{\alpha}^{K-i}\frac{\tilde{c}}{\delta} \quad (27)$$

$$\implies \max_{k \leq K} M_k \geq \tilde{\alpha}^K\eta(\mathbf{x}_0) - \sum_{i=1}^k \tilde{\alpha}^{K-i}\frac{\tilde{c}}{\delta} \quad (28)$$

$$\implies \max_{k \leq K} M_k \geq \tilde{\alpha}^K\eta(\mathbf{x}_0) - \sum_{i=1}^K \tilde{\alpha}^{K-i}\frac{\tilde{c}}{\delta}, \quad (29)$$

⁷We use the “empty sum” convention that $\sum_{i=1}^0 \rho = 0$ for any $\rho \in \mathbb{R}$.

⁸ $W_0 = 0$ since $\mathbf{x}_0$ is known and randomness first enters through $\mathbf{d}_0$.




Fig. 2. Comparison for Prop. 1 with $B = 10$, $K = 100$, $\delta = 1$, and varying σ and λ . The Freedman-based bounds are shown in green when the conditions of Prop. 1 hold and blue when they do not. The Ville's-based bound is shown in red.

where (26) is due to multiplication by a value strictly less than zero, (27) is due to adding zero, (28) is due to $M_k \geq W_k$ as in (20), and (29) is due to $k \leq K$ and the nonnegativity of α , δ , and $\tilde{c}$. Thus, the unsafe event satisfies the containment:

$$\left\{ \min_{k \leq K} h(\mathbf{x}_k) < 0 \right\} \subseteq \left\{ \max_{k \leq K} M_k \geq \tilde{\alpha}^K\eta(\mathbf{x}_0) - \sum_{i=1}^K \tilde{\alpha}^{K-i}\frac{\tilde{c}}{\delta} \right\}$$

Since M_k satisfies $M_0 = 0$, $M_k - M_{k-1} \leq 1 \forall k \leq K$, and $\langle M \rangle_K \leq \frac{\sigma^2 K}{\delta^2}$, we can apply Thm. 2 (Freedman's Ineq.) with $\lambda = \frac{\tilde{\alpha}^K h(\mathbf{x}_0) - \sum_{i=1}^K \tilde{\alpha}^{K-i}\tilde{c}}{\delta}$ to achieve the probability bound⁹:

$$P_u(K, \mathbf{x}_0) \leq H\left(\frac{\tilde{\alpha}^K h(\mathbf{x}_0) - \sum_{i=1}^K \tilde{\alpha}^{K-i}\tilde{c}}{\delta}, \frac{\sigma\sqrt{K}}{\delta}\right).$$

(Step 6) If the system satisfies the DTCBF condition, then (17) holds with $(\tilde{\alpha} = \alpha, \tilde{c} = 0)$ so the desired bound is achieved with $\lambda = \alpha^K h(\mathbf{x}_0)/\delta$ and if the system satisfies the c-mart. condition then (17) holds with $(\tilde{\alpha} = 1, \tilde{c} = c)$ so the desired bound is achieved with $\lambda = h(\mathbf{x}_0)/\delta - Kc$. ■

B. Bound Tightness Comparison

We now relate the Freedman-based safety of Thm. 3 to the Ville's-based safety of Thm. 1. For systems that have an upper-bound h (10), a lower-bounded error (14), and a bounded conditional variance (15), we provide a range of values for σ , δ , K , B , and λ for which Thm. 3 is stronger. This Prop. provides a direct theoretical comparison (after changing notation) to the Ville's-based bounds in [10], [11], [14], [16].

Proposition 1: For some $\sigma, \delta, B > 0$, $\lambda \geq 0$ and $K \in \mathbb{N}_1$, consider the conditions

$$\lambda \delta \geq \sigma^2 K, \quad \lambda \leq B - \frac{\delta}{\varphi}, \quad (30)$$

where $\varphi = 2 \ln(2) - 1$. If these conditions hold, then

$$H\left(\frac{\lambda}{\delta}, \frac{\sigma\sqrt{K}}{\delta}\right) \leq 1 - \frac{\lambda}{B}. \quad (31)$$

Proof of this Proposition is provided in the Appendix.

Intuitively, conditions (30) stipulate that the conditional variance σ^2 and number of steps K must be limited by $\lambda\delta$, which is a function of the initial condition times the maximum single-step disturbance to $h(\mathbf{x}_k)$. Additionally, the initial condition must be less than the maximum safety bound B by an amount proportional to δ . The exact value of φ is a result of the first assumption ($\lambda\delta \geq \sigma^2 K$) and alternative values can be found by changing this assumption; for clarity, we leave exploration of these alternative assumptions to future work. The safety bounds for various λ and σ are shown in

⁹The proof can end after Step 5 and can be applied to any system satisfying (17). We specialize to DTCBFs and c-martingales for clarity.




Fig. 3. Probability that the system is unsafe: our bound from Cor. 1 (blue), ISSf bound (red). The x-axis is the level set expansion $-\epsilon$ and the y-axis is the failure probability (lower is better). The plots from left to right indicate safety for $K = 1, 100, 200, 300$, and 400 steps. Simulations where $\mathbb{E}[h(\mathbf{x}_k) | \mathcal{F}_{k-1}] = \alpha h(\mathbf{x}_k)$ and approximate probabilities from 1000 samples are shown for simulations where $h(\mathbf{x}_k)$ is sampled from 3 different conditional distributions: uniform (pink), truncated Gaussian (green), and a categorical (yellow) all which satisfy Cor. 1.

Fig. 2 where these conditions provide a *conservative* set of parameters over which this proposition holds.

C. Extending Input-to-State Safety

Since Thm. 3 assumes that h has lower-bounded errors (14), we can directly compare our method with Input-to-State Safety (ISSf) [2], which provides almost-sure safety guarantees.

In the context of our stochastic, discrete-time problem setting, the ISSf property can be reformulated as:

Proposition 2 (Input-to-State Safety): If the closed-loop system (2) satisfies the (DTCBF) condition and the bounded-jump condition (14) (a.s.) for some $\alpha \in [0, 1)$ and $\delta > 0$, then $h(\mathbf{x}_k) \geq \alpha^k h(\mathbf{x}_0) - \sum_{i=0}^{k-1} \alpha^i \delta$ for all $k \geq 0$ and $\mathcal{C}_\delta = \{\mathbf{x} \in \mathbb{R}^n \mid h(\mathbf{x}) \geq \frac{\delta}{1-\alpha}\}$ is safe (a.s.).

Proof: By combining the bounds (DTCBF) and (14):

$$h(\mathbf{x}_{k+1}) \geq \mathbb{E}[h(\mathbf{x}_{k+1}) \mid \mathcal{F}_k] - \delta \geq \alpha h(\mathbf{x}_k) - \delta \quad (\text{a.s.}) \quad (32)$$

Thus, for all $k \in \mathbb{Z}$, we have the lower-bound $h(\mathbf{x}_k) \geq \alpha^k h(\mathbf{x}_0) - \sum_{i=0}^{k-1} \alpha^i \delta$ (a.s.). Furthermore, for all time, $h(\mathbf{x}_k) \geq \frac{-\delta}{1-\alpha} \implies h(\mathbf{x}_{k+1}) \geq \frac{-\delta}{1-\alpha}$, so $\mathcal{C}_\delta$ is safe (a.s.). ■

To compare with ISSf, we wish to use Thm. 3 to bound the probability that our system leaves some expanded safe set $\mathcal{C}_\epsilon = \{\mathbf{x} \in \mathbb{R}^n \mid h(\mathbf{x}) \geq -\epsilon\}$ with $\epsilon \geq 0$ in finite time.

Corollary 1: If the hypotheses of Theorem 3 are satisfied and (2) satisfies the (DTCBF) for some $\alpha \in (0, 1)$, then for any value $\epsilon \geq 0$ and any $K \in \mathbb{N}_1$,

$$P \left\{ \min_{k \leq K} h(\mathbf{x}_k) < -\epsilon \right\} \leq H \left(\lambda, \frac{\sigma}{\delta} \left(\frac{1 - \alpha^{2K}}{1 - \alpha^2} \right)^{\frac{1}{2}} \right) \mathbb{1}_{\{-\epsilon \geq \alpha^K h(\mathbf{x}_0) - \sum_{i=0}^{K-1} \alpha^i \delta\}} \quad (33)$$

where $\lambda = \frac{\alpha^K}{\delta} (h(\mathbf{x}_0) + \epsilon)$.

Proof: The (DTCBF) condition ensures that, for any $\epsilon \geq 0$:

$$\mathbb{E}[h(\mathbf{x}_{k+1}) + \epsilon \mid \mathcal{F}_k] \geq \alpha(h(\mathbf{x}_k) + \epsilon) + \epsilon(1 - \alpha) \quad (34)$$

$$\geq \alpha(h(\mathbf{x}_k) + \epsilon) \quad (35)$$

We apply the same proof as Thm. 3 starting at (18) with $(\eta(\mathbf{x}_k) = \frac{h(\mathbf{x}_k) + \epsilon}{\delta}, \tilde{\alpha} = \alpha, \tilde{c} = 0)$. Choosing $\lambda = \alpha^K \eta(\mathbf{x}_0)$ and bounding¹⁰ $\langle M \rangle_K \leq \sum_{i=1}^K \alpha^{2(K-i)} \frac{\sigma^2}{\delta^2} = \frac{\sigma^2(1 - \alpha^{2K})}{\delta^2(1 - \alpha^2)}$ as in (23) yields the desired bound without the indicator function by applying Thm. 2. The indicator function is a result of applying the lower bound on the safety value from Prop. 1, i.e., $h(\mathbf{x}_k) \geq \alpha^k h(\mathbf{x}_0) - \sum_{i=0}^{k-1} \alpha^i \delta$ (a.s.) for $k \in \mathbb{Z}$. ■

A comparison of Prop. 1 and Cor. 2 and Monte Carlo approximations for various ϵ and distributions (truncated normal, categorical) is shown in Fig. 3 where we can see that our method successfully upper-bounds the sampled safety probabilities with risk-sensitive guarantees that are much less conservative than the worst-case bounds provided by ISSf.

¹⁰This bound on $\langle M \rangle_K$ uses the finite geometric series identity and can also be applied for a tighter Thm. 3 and Prop. 1.

IV. CASE STUDY: BIPEDAL OBSTACLE AVOIDANCE

In this section we apply our method to a simplified model of a bipedal walking robot. In particular, the Hybrid Linear Inverted Pendulum (HLIP) model [29] approximates a bipedal robot as an inverted pendulum with a fixed center of mass (COM) height $z_0 \in \mathbb{R} > 0$. Its states are the planar position, relative COM-to-stance foot position, and COM velocity $\mathbf{p}, \mathbf{c}, \mathbf{v} \in \mathbb{R}^2$. The step-to-step dynamics are linear and the input is the relative foot placement, $\mathbf{u}_k \in \mathbb{R}^2$. The matrices $\mathbf{A} \in \mathbb{R}^{6 \times 6}$ and $\mathbf{B} \in \mathbb{R}^{6 \times 2}$ are determined by z_0 and gait parameters including the stance and swing phase periods. The HLIP model with an added disturbance matrix $\mathbf{D} \in \mathbb{R}^{6 \times 4}$ and disturbance $\mathbf{d} \in \mathbb{R}^4$ affecting position and velocity is:

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k + \mathbf{D}\mathbf{d}_k, \quad \mathbf{d}_k \sim \mathcal{D}.$$

where $\mathbf{x}_k = [\mathbf{p}_k^\top \ \mathbf{c}_k^\top \ \mathbf{v}_{k+1}^\top]^\top$. We augment the standard HLIP model and assume that $\mathbf{d}$ enters linearly and $\mathcal{D}$ is a 4-dimensional, 0-mean uniform distribution¹¹ with $\|\mathbf{d}\| \leq d_{\max}$.

We define safety for this system as avoiding a circular obstacle of radius $r > 0$ located at $(x, y) = \boldsymbol{\rho} \in \mathbb{R}^2$, so safety can be defined using the signed-distance function $h(\mathbf{x}) = \|\mathbf{p} - \boldsymbol{\rho}\|_2 - r$. Notably, this function has no upper bound and therefore the Ville's-based Thm. 1 does not apply.

Since $h(\mathbf{x})$ is not convex, we use a conservative halfspace convexification instead:

$$h(\mathbf{x}_{k+1}) \geq \widehat{\mathbf{e}}(\mathbf{p}_k)^\top (\mathbf{p}_{k+1} - \boldsymbol{\rho}) - r \triangleq \bar{h}(\mathbf{x}_{k+1}), \quad (36)$$

where $\widehat{\mathbf{e}}(\mathbf{p}) = \frac{(\mathbf{p} - \boldsymbol{\rho})}{\|\mathbf{p} - \boldsymbol{\rho}\|}$ and we apply the controller:

$$\begin{aligned} \mathbf{u}^* &= \min_{\mathbf{u} \in \mathbb{R}^2} \|\mathbf{u} - \mathbf{k}_{\text{nom}}(\mathbf{x}_k)\| \\ \text{s.t.} \quad &\mathbb{E}[\bar{h}(\mathbf{x}_{k+1}) \mid \mathcal{F}_k] \geq \alpha \bar{h}(\mathbf{x}_k) \end{aligned} \quad (37)$$

with $\alpha \in (0, 1]$ and where $\mathbf{k}_{\text{nom}}$ tracks a desired velocity.

We ran 5000 trials with 3 steps per second and compared against the theoretical bound from Thm. 3. Those values and planar pose trajectories can be seen in Fig. 1. Exact values and code for this and all other plots can be found at [27].

V. CONCLUSION

Despite the relative tightness guarantee of Prop. 1, the probability guarantees of our method are not necessarily tight, as can be seen in Fig. 3. Optimization of h without changing $\mathcal{C}$ as in [10] is a promising direction further tightening of our bound. Additionally, the case study shown in Section IV presents an immediate direction for future work which may involving a hardware demonstration of this method.

APPENDIX PROOF OF PROPOSITION 1

Proof: Define $\Delta(\lambda, B, \sigma, K, \delta) \triangleq 1 - \frac{\lambda}{B} - H\left(\frac{\lambda}{\delta}, \frac{\sigma\sqrt{K}}{\delta}\right)$.

¹¹See [27, Appx. H] for bounds for δ and σ given this problem structure.

If $\Delta(\lambda, B, \sigma, K, \delta) \geq 0$, then (31) must hold. We first show Δ is monotonically decreasing in σ^2 . Consider¹² $\frac{\partial \Delta}{\partial(\sigma^2)} = a(\lambda, \sigma, K, \delta)b(\lambda, \sigma, K, \delta)$ where

$$a(\lambda, \sigma, K, \delta) \triangleq \frac{-e^{\frac{\lambda}{\delta}}}{\delta^2 \sigma^2} \left(\frac{\sigma^2 K}{\lambda \delta + \sigma^2 K} \right)^{\frac{(\lambda \delta + \sigma^2 K)}{\delta^2}} < 0, \quad (38)$$

$$b(\lambda, \sigma, K, \delta) \triangleq \left(\sigma^2 K \ln \left(\frac{\sigma^2 K}{\lambda \delta + \sigma^2 K} \right) + \lambda \delta \right). \quad (39)$$

The function $a(\lambda, \sigma, K, \delta)$ is negative since $\delta, \sigma, K > 0$. For $b(\cdot)$, the logarithm bound $\ln(r) \geq 1 - 1/r$ ensures that:

$$b(\lambda, \sigma, K, \delta) \geq \sigma^2 K \left(1 - \frac{\lambda \delta + \sigma^2 K}{\sigma^2 K} \right) + \lambda \delta = 0. \quad (40)$$

Since $a < 0$ and $b \geq 0$, $\Delta(\lambda, B, \sigma, K, \delta)$ is monotonically decreasing with respect to σ^2 , so we can use the assumption $\sigma^2 K \leq \lambda \delta$ to lower bound Δ as:

$$\begin{aligned} \Delta(\lambda, B, \sigma, K, \delta) &\geq 1 - \frac{\lambda}{B} - \left(\frac{1}{2} \right)^{2\frac{\lambda}{\delta}} e^{\frac{\lambda}{\delta}} \\ &= 1 - \frac{\lambda}{B} - e^{(1-2\ln(2))\frac{\lambda}{\delta}} \triangleq 1 - \frac{\lambda}{B} - e^{-\varphi\frac{\lambda}{\delta}} \triangleq \Delta_1(\lambda, B, \delta) \end{aligned} \quad (41)$$

where $\varphi \triangleq 2\ln(2) - 1 > 0$.

Next, we show that $\Delta_1(\lambda, B, \delta) \geq 0$ for¹³ $\lambda \in [0, B - \frac{\delta}{\varphi}]$. We prove this by showing that $\Delta_1(\lambda, B, \delta) \geq 0$ for $\lambda = \{0, B - \frac{\delta}{\varphi}\}$ and that Δ_1 is concave with respect to λ .

(1) Nonnegativity at $\lambda = 0$: $\Delta_1(0, B, \delta) = 0$.

(2) Nonnegativity at $\lambda = B - \frac{\delta}{\varphi}$:

$$\begin{aligned} \Delta_1\left(B - \frac{\delta}{\varphi}, B, \delta\right) &= \frac{\delta}{\varphi B} - e^{-\left(B - \frac{\delta}{\varphi}\right)\frac{\varphi}{\delta}} = \frac{\delta}{\varphi B} - e^{\left(1 - \frac{B\varphi}{\delta}\right)} \\ &\geq \frac{\delta}{\varphi B} - \frac{\delta}{\varphi B} = 0, \end{aligned} \quad (42)$$

where the inequality in line (42) is due to the previously used log inequality: $\ln(r) \geq 1 - \frac{1}{r} \iff r \geq e^{(1-\frac{1}{r})}$, which holds for $r = \frac{\delta}{B\varphi} > 0$ since $B, \delta, \varphi > 0$.

(3) Concavity for $\lambda \in [0, B - \frac{\delta}{\varphi}]$: Since $\frac{\varphi}{\delta} > 0$, the second derivative of Δ_1 with respect to λ is negative:

$$\frac{\partial^2 \Delta_1}{\partial \lambda^2} = -\left(\frac{\varphi}{\delta}\right)^2 e^{-\varphi\frac{\lambda}{\delta}} < 0. \quad (43)$$

Thus, Δ_1 is concave with respect to λ . Since, $\Delta_1(0, B, \delta) \geq 0$, $\Delta_1(B - \frac{\delta}{\varphi}, B, \delta) \geq 0$, and $\Delta_1(\lambda, B, \delta)$ is concave for all $\delta > 0$ and $B \geq \frac{\delta}{\varphi}$, it follows from the definition of concavity that $\Delta_1(\lambda, B, \delta) \geq 0$ for all $\lambda \in [0, B - \frac{\delta}{\varphi}]$.

Using this lower bound for $\Delta_1(\lambda, B)$, we have $\Delta(\lambda, B, \sigma, K, \delta) \geq \Delta_1(\lambda, B) \geq 0$ which implies the desired inequality (31). ■

REFERENCES

- [1] A. D. Ames et al., "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, Aug. 2017. [Online]. Available: <http://arxiv.org/abs/1609.06408>
- [2] S. Kolathaya and A. D. Ames, "Input-to-state safety with control barrier functions," *IEEE Control Syst. Lett.*, vol. 3, no. 1, pp. 108–113, Jan. 2019.
- [3] S. Bansal et al., "Hamilton–Jacobi reachability: A brief overview and recent advances," in *Proc. IEEE 56th Annu. Conf. Decis. Control (CDC)*, 2017, pp. 2242–2253.
- [4] F. Borrelli et al., *Predictive Control for Linear and Hybrid Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2017.
- [5] Code. 2023. "Repository this work." [Online]. Available: <https://github.com/rkcosner/freedman.git>
- [6] A. Abate et al., "Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, 2008.
- [7] M. Fränzle et al., "Measurability and safety verification for stochastic hybrid systems," in *Proc. 14th Conf. Hybrid Syst. Comput. Control*, 2011, pp. 43–52.
- [8] M. P. Chapman et al., "Risk-sensitive safety analysis using conditional value-at-risk," *IEEE Trans. Autom. Control*, vol. 67, no. 12, pp. 6521–6536, Dec. 2022.
- [9] L. Lindemann et al., "STL robustness risk over discrete-time stochastic processes," in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, 2021, pp. 1329–1335.
- [10] J. Steinhardt and R. Tedrake, "Finite-time regional verification of stochastic non-linear systems," *Int. J. Robot. Res.*, vol. 31, no. 7, pp. 901–923, Jun. 2012. [Online]. Available: <https://doi.org/10.1177/0278364912444146>
- [11] C. Santoyo et al., "A barrier function approach to finite-time stochastic system verification and control," *Automatica*, vol. 125, Mar. 2021, Art. no. 109439. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0005109820306415>
- [12] O. So et al., "Almost-sure safety guarantees of stochastic zero-control barrier functions do not hold," 2023, *arXiv:2312.02430*.
- [13] M. Black et al., "Safety under uncertainty: Tight bounds with risk-aware control barrier functions," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2023, pp. 12686–12692. doi: [10.1109/ICRA48891.2023.10161379](https://doi.org/10.1109/ICRA48891.2023.10161379).
- [14] H. Kushner, *Stochastic Stability and Control*. New York, NY, USA: Academic, 1967.
- [15] S. Prajna et al., "Stochastic safety verification using barrier certificates," in *Proc. 43rd IEEE Conf. Decis. Control (CDC)*, vol. 1, 2004, pp. 929–934.
- [16] R. Cosner et al., "Robust safety under stochastic uncertainty with discrete-time control barrier functions," in *Proc. Robot. Sci. Syst.*, 2023, pp. 1–8.
- [17] C. Santoyo et al., "Verification and control for finite-time safety of stochastic systems via barrier functions," in *Proc. IEEE Conf. Control Technol. Appl. (CCTA)*, 2019, pp. 712–717.
- [18] F. B. Mathiesen et al., "Inner approximations of stochastic programs for data-driven stochastic barrier function design," in *Proc. 62nd IEEE Conf. Decis. Control (CDC)*, 2023, pp. 3073–3080.
- [19] R. K. Cosner et al., "Generative modeling of residuals for real-time risk-sensitive safety with discrete-time control barrier functions," May 2023, *arxiv:abs/2311.05802*.
- [20] J. Ville. "Etude critique de la notion de collectif." 1939. [Online]. Available: http://www.numdam.org/item/THESE_1939_218_1_0.pdf
- [21] D. A. Freedman, "On tail probabilities for martingales," *Ann. Probab.*, vol. 3, no. 1, pp. 100–118, 1975.
- [22] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *Int. J. Robot. Res.*, vol. 4, no. 1, pp. 90–98, 1986.
- [23] A. Agrawal and K. Sreenath, "Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation," Jul. 2017. [Online]. Available: <http://www.roboticsproceedings.org/rss13/p73.pdf>
- [24] J. Breeden et al., "Control barrier functions in sampled-data systems," *IEEE Control Syst. Lett.*, vol. 6, pp. 367–372, 2022.
- [25] P. Culbertson et al., "Input-to-state stability in probability," in *Proc. IEEE Conf. Decis. Control (CDC)*, 2023, pp. 367–372.
- [26] G. Grimmett and D. Stirzaker, *Probability and Random Processes*. Oxford, U.K.: Oxford Univ. Press, Jul. 2020.
- [27] R. K. Cosner et al., "Extended version of this paper." 2023. [Online]. Available: <https://arxiv.org/pdf/2403.05745.pdf>
- [28] X. Fan et al., "Hoeffding's inequality for supermartingales," *Stochastic Processes Appl.*, vol. 122, no. 10, pp. 3545–3559, Oct. 2012.
- [29] X. Xiong and A. Ames, "3-D underactuated bipedal walking via H-LIP based gait synthesis and stepping stabilization," *IEEE Trans. Robot.*, vol. 38, no. 4, pp. 2405–2425, Aug. 2022.

¹²The derivation of $\frac{\partial \Delta}{\partial(\sigma^2)}$ is provided in our extended manuscript [27].

¹³This interval is non-empty since $\lambda \geq 0$ and $B \geq \lambda + \frac{\delta}{\varphi}$ implies $B \geq \frac{\delta}{\varphi}$.