

Input-to-State Stability in Probability

Preston Culbertson¹, Ryan K. Cosner¹, Maegan Tucker¹, and Aaron D. Ames^{1,2}.

Abstract—Input-to-State Stability (ISS) is fundamental in mathematically quantifying how stability degrades in the presence of bounded disturbances. If a system is ISS, its trajectories will remain bounded, and will converge to a neighborhood of an equilibrium of the undisturbed system. This graceful degradation of stability in the presence of disturbances describes a variety of real-world control implementations. Despite its utility, this property requires the disturbance to be bounded and provides invariance and stability guarantees only with respect to this worst-case bound. In this work, we introduce the concept of “ISS in probability (ISSp)” which generalizes ISS to discrete-time systems subject to unbounded stochastic disturbances. Using tools from martingale theory, we provide Lyapunov conditions for a system to be exponentially ISSp, and connect ISSp to stochastic stability conditions found in literature. We exemplify the utility of this method through its application to a bipedal robot confronted with step heights sampled from a truncated Gaussian distribution.

I. INTRODUCTION

Control systems operating in practice are nearly always affected by disturbances, be they noise, modelling error, uncertain state estimates, or environmental interactions. This motivates the design of controllers which are robust to these uncertainties. Input-to-state stability (ISS) [1], [2] is a useful heuristic for the robustness of a control system. If a system is ISS, then, loosely, when the system is subjected to bounded disturbances, the system state will converge to some ball whose radius scales with the maximum disturbance norm; in the presence of zero disturbances, the system is asymptotically stable. ISS can be interpreted as guaranteeing the “graceful degradation” of asymptotic stability under bounded disturbances; bounded disturbance inputs still produce bounded state trajectories, and asymptotic stability is recovered as the input magnitude approaches zero.

However, as a robustness property, ISS suffers some drawbacks, particularly when reasoning about systems subject to stochastic disturbances. The central issue is that ISS reasons only about bounded disturbances, i.e., those whose norm is upper-bounded. However, many noise sources are more naturally modeled as continuous, unbounded random variables (e.g., systems subject to additive Gaussian noise); ISS-based tools cannot handle such cases. Further, the guarantees provided by ISS depend on the worst-case disturbance magnitude and are thus often highly conservative.

To remedy these limitations, in this paper we introduce a generalization of the ISS property for discrete-time systems subject to unbounded stochastic disturbances: *input-to-state stability in probability (ISSp)*. Intuitively, we say a system is ISSp if the typical ISS condition can hold with a probability

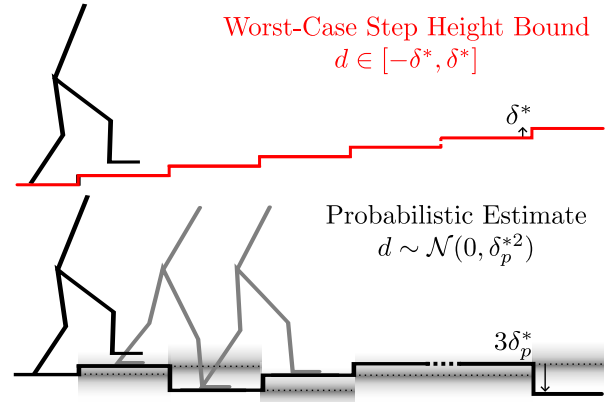


Fig. 1: This paper introduces input-to-state stability in probability (ISSp), which generalizes input-to-state stability (ISS) to systems with unbounded disturbances. We use this framework to study a seven-link walker traversing stochastic terrain. While ISSp only provides probabilistic guarantees, we find our framework yields more reasonable estimates for the tolerable set of step heights. **(Top)**: ISS-based guarantees must hold for any (bounded) disturbance signal; even for worst-case terrain (e.g., stairs) the walker must be able to remain stable. **(Bottom)**: ISSp reasons instead about how systems behave over finite horizons. While the stochastic step heights (shown as gradients) can sometimes be large, their distribution is concentrated near zero, and thus the walker has a high probability of remaining upright.

arbitrarily close to one for an arbitrary (but finite) horizon. Using tools from martingale theory, we provide Lyapunov conditions for the exponential form of ISSp. We also explore connections between ISS, ISSp, and more traditional stability notions for stochastic systems. We conclude with simulation studies of ISSp systems subject to unbounded disturbances, including a double-integrator subject to additive Gaussian noise and a bipedal robot walking on uncertain terrain as illustrated in Fig. 1. In particular, we show that our ISSp-based exit probability bound is indeed conservative for all examples, and show that we can provide non-trivial probabilistic stability guarantees for the biped over a larger disturbance set than the worst-case ISS bound [3].

There has been a large body of work on ISS. Beginning with the seminal papers by Sontag [2], [4] for continuous time and the extension of these results to discrete time [5], ISS has found utility in the fields of control theory [6] robotics [7], and of special note to this paper, robotic walking [8], [9], [3]. This paper leverages results on martingale theory [10] to extend (discrete-time) ISS to stochastic systems. In particular, the stochastic notions presented in this paper are similar to the set-invariance notions in [11], [12], but differ in that they add a notion of stochastic stability to reflect the convergence present in systems that are ISS. Theoretically, this work uses similar tools to those presented in [13], which uses a similar supermartingale to bound the finite-time exit probability of a system from a particular set. However, in this work, we provide a novel “ISS-like” interpretation of this supermartingale, and make explicit connections between the concepts of ISS and other stochastic stability notions found

This research was supported by the National Science Foundation (CPS Award #1932091), BP, and the Zeitlin Family Fund

¹ Authors are with the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125, USA; {pculbert, rkcosner, mtucker, ames}@caltech.edu

² Authors are with the Department of Control and Dynamical Systems, California Institute of Technology, Pasadena, CA 91125, USA.

in literature such as variable drift [14], recurrence [15], and boundedness of trajectories in probability [11].

II. BACKGROUND

Consider a discrete-time autonomous system,

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{d}_k) \quad (1)$$

with $k \in \mathbb{N}_{\geq 0}$, state $\mathbf{x}_k \in \mathcal{X} \subseteq \mathbb{R}^n$, equilibrium point $(\mathbf{x}^*, \mathbf{d}^*) = (\mathbf{0}, \mathbf{0})$, random disturbance $\mathbf{d}_k \in \mathbb{R}^d$, and continuous dynamics $\mathbf{f} : \mathcal{X} \times \mathbb{R}^d \rightarrow \mathbb{R}^n$. We assume each disturbance $\mathbf{d}_k \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}$ from some disturbance distribution \mathcal{D} .

A. Input-to-State Stability for Deterministic Systems

If the disturbance distribution for system (1) is bounded, then we can use the concept of Input-to-State Stability (ISS) to reason about the boundedness and convergence of the system's trajectories.

Definition 1 (Input-to-State Stability [5]). The system (1) is input-to-state stable (ISS) if there exist functions¹ $\beta \in \mathcal{K}\mathcal{L}$ and $\gamma \in \mathcal{K}$ such that, for each deterministic disturbance input $\mathbf{d}_k \in \mathbb{R}^m$ and each $\mathbf{x}_0 \in \mathbb{R}^n$, it holds that

$$\|\mathbf{x}_k\| \leq \beta(\|\mathbf{x}_0\|, k) + \gamma \left(\sup_{k \in \mathbb{N}_{\geq 0}} \|\mathbf{d}_k\| \right) \quad (\text{ISS})$$

for each $k \in \mathbb{N}_{\geq 0}$ and some $p \geq 1$.

Intuitively, the bound on the state trajectory is a function of a sequence which converges to zero in time, $\beta(\|\mathbf{x}_0\|, k)$ and a term which grows with respect to the disturbance bound, $\gamma \left(\sup_{k \in \mathbb{N}_{\geq 0}} \|\mathbf{d}_k\| \right)$. If $\|\mathbf{d}_k\| = 0$ for all k , then ISS systems are asymptotically stable. Note that a similar inequality regarding an essentially bounded disturbance distribution \mathcal{D} :

$$\|\mathbf{x}_k\| \leq \beta(\|\mathbf{x}_0\|, k) + \gamma(\text{ess sup}\|\mathcal{D}\|), \quad (2)$$

can be employed to achieve ISS almost surely (Corollary 1) where ess sup is the essential supremum of the distribution \mathcal{D} , also written as the L^∞ -norm of \mathcal{D} .

We now introduce ISS-Lyapunov functions as tools for verifying a system's ISS property.

Definition 2 (ISS-Lyapunov Function [5]). A continuous function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is an ISS Lyapunov function for (1) if there exist $\kappa_1, \kappa_2, \kappa_3 \in \mathcal{K}_\infty$ and $\kappa_4 \in \mathcal{K}$ such that:

$$\kappa_1(\|\mathbf{x}\|) \leq V(\mathbf{x}) \leq \kappa_2(\|\mathbf{x}\|) \quad (3)$$

$$V(\mathbf{f}(\mathbf{x}, \mathbf{d})) - V(\mathbf{x}) \leq -\kappa_3(\|\mathbf{x}\|) + \kappa_4(\|\mathbf{d}\|) \quad (4)$$

for all $\mathbf{x} \in \mathbb{R}^n$ and all $\mathbf{d} \in \mathbb{R}^d$. Additionally, V is an exponential-ISS (E-ISS) Lyapunov function if there exist constants $a, b, c > 0$ and $\alpha \in (0, 1)$ such that $\kappa_1(r) = ar^c$, $\kappa_2(r) = br^c$, $\kappa_3(r) = \alpha r^c$.

The existence of an ISS-Lyapunov function can now be immediately used to verify that the system is ISS.

Theorem 1 ([5]). *If there exists an ISS Lyapunov Function for system (1), then system (1) is ISS.*

¹A continuous function $\gamma : [0, a) \rightarrow [0, \infty)$ for $a > 0$ is said to belong to class \mathcal{K} ($\gamma \in \mathcal{K}$) if it is strictly monotonically increasing and $\gamma(0) = 0$. If additionally $a = \infty$ and $\gamma(r) \rightarrow \infty$ as $r \rightarrow \infty$ then γ belongs to \mathcal{K}_∞ . A continuous function $\beta : [0, a) \times [0, \infty) \rightarrow [0, \infty)$ is said to belong to class $\mathcal{K}\mathcal{L}$ if for each fixed $s \geq 0$ the function $\beta(\cdot, s)$ is class \mathcal{K} and for each $r \geq 0$ the function $\beta(r, \cdot)$ is decreasing and $\beta(r, s) \rightarrow 0$ as $s \rightarrow \infty$.

B. Stochastic Preliminaries: L^p spaces and Martingales

Here we provide a brief discussion of random variables, martingales, and other tools that we will use to generalize ISS to the case of unbounded, stochastic disturbances. We will present this material at a level necessary to communicate these concepts clearly and accessibly. We refer readers to [16] for a precise measure-theoretic presentation of these ideas.

In this paper we consider disturbance signals which are sequences of random variables. A continuous random variable \mathbf{y} sampled from a distribution \mathcal{Y} (denoted $\mathbf{y} \sim \mathcal{Y}$) is a quantity that takes on values in \mathbb{R}^y according to a probability density $p(\mathbf{y}) \geq 0$, with $\mathbb{P}\{\mathbf{y} \in A\} \triangleq \int_A p(\mathbf{v})d\mathbf{v}$. By definition $\int_{\mathbb{R}^y} p(\mathbf{v})d\mathbf{v} = 1$ and the expectation of a random variable is given by $\mathbb{E}[\mathbf{y}] \triangleq \int_{\mathbb{R}^y} \mathbf{v}p(\mathbf{v})d\mathbf{v}$.

We now introduce L^p spaces of random variables.

Definition 3 (L^p Space [17]). A random variable $\mathbf{y} \sim \mathcal{Y}$ belongs to L^p (denoted as $\mathbf{y} \in L^p$), for $p > 0$, if

$$\|\mathbf{y}\|_{L^p} \triangleq \mathbb{E}[\|\mathbf{x}\|^p]^{\frac{1}{p}} < \infty. \quad (5)$$

We call $\|\cdot\|_{L^p}$ the p -norm of a random variable, which is finite for any random variable in L^p . Intuitively, for $0 < p \leq q$, $L^q \subseteq L^p$ [17, Thm. 8.2] since random variables in L^q have tails that decay faster than those in L^p ; additionally, L^∞ is the smallest L^p space and only contains random variables that are essentially bounded. Note that any norm $\|\cdot\|$ appearing without a subscript defines a typical norm on \mathbb{R}^n .

We can also reason about a random variable's conditional probability, i.e., its distribution given that another random variable has taken on a particular value. For two random variables X, Y the density of X given $Y = y$ is given by

$$p(x | y) = \frac{p(x, y)}{p(y)}$$

where $p(x, y)$ is the joint probability density of X, Y . The conditional expectation of X given $Y = y$ is $\mathbb{E}[X | Y]$.

The key tool used to reason about Lyapunov functions for our probabilistic notion of ISS 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 , $W : \mathcal{X} \times \mathbb{N}_{\geq 0} \rightarrow \mathbb{R}$, and suppose that $W(\mathbf{x}_k, k) \in L^1$ for $k \in \mathbb{N}_{\geq 0}$. The process $W_k \triangleq W(\mathbf{x}_k, k)$ is a supermartingale if:

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

where $\mathbf{x}_{0:k}$ indicates the random variables $\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k\}$. If, additionally, $W_k \geq 0$ for all $k \in \mathbb{N}_{\geq 0}$, 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 is *Ville's inequality*, which bounds the probability that a nonnegative supermartingale rises above a certain value:

Theorem 2 (Ville's Inequality [18]). *Let W_k be a nonnegative supermartingale. Then for all $\lambda \in \mathbb{R}_{>0}$,*

$$\mathbb{P} \left\{ \sup_{k \in \mathbb{N}_{\geq 0}} W_k > \lambda \right\} \leq \frac{\mathbb{E}[W_0]}{\lambda}. \quad (7)$$

Intuitively, Ville's inequality can be compared with Markov's inequality for nonnegative random variables; since the process \tilde{W}_k is nonincreasing in expectation, Ville's inequality allows us to reason about the probability the process instead reaches some value above λ .

III. STABILITY OF STOCHASTIC DISCRETE-TIME SYSTEMS

Traditional notions of stability may not necessarily apply to stochastic systems. For example, asymptotic stability to a point or forward invariance of a bounded set may be impossible in the presence of unbounded, stochastic disturbances. Thus more nuanced notions of stability are required [19]. In this section we provide an abridged discussion of existing stability notions for stochastic systems. Notably, we discuss recurrence, boundedness of trajectories in probability and input-to-state stability for distributions with bounded support.

A. Recurrence

Recurrence is an important notion of stability used in the analysis of Markov chains [15]. A bounded set $A \subset \mathcal{X}$ is recurrent if trajectories enter A in finite time and visit A infinitely often with probability 1 for all initial states $\mathbf{x}_0 \in \mathcal{X}$.

Definition 5 (Recurrence). For some bounded set $A \subset \mathcal{X}$ let the hitting time $\tau_A(\mathbf{x}) \triangleq \inf\{k \in \mathbb{N}_{\geq 0} \text{ s.t. } \mathbf{x}_k \in A, \mathbf{x}_0 = \mathbf{x}\}$. A set A is recurrent if for every $\mathbf{x} \in \mathcal{X}$, $\mathbb{P}\{\tau_A(\mathbf{x}) < \infty\} = 1$. We say a system (1) is recurrent if there exists a recurrent set A .

Recurrence relates to the notion of stability for deterministic systems where trajectories remain within a set for all time, a property which is guaranteed for ISS systems. We refer the reader to [15] for a more thorough treatment of Markov chain stability, recurrence, and ergodic theory.

B. Boundedness in Probability

Another notion of stability for stochastic systems is the probability that the state remains in a bounded region. Since it is often impossible to keep trajectories of (1) bounded for all time [11], it is common to discuss these probabilities over some finite horizon $k \in \{0, \dots, K\}$ for some $K \in \mathbb{N}_{\geq 0}$.

Definition 6 (Bounded in Probability). The system (1) is bounded in probability for some $K \in \mathbb{N}_{\geq 0}$ if there exists an $M > 0$ and $\epsilon \in (0, 1)$ such that

$$\mathbb{P}\left\{\max_{k \leq K} \|\mathbf{x}_k\| \leq M\right\} \geq 1 - \epsilon. \quad (8)$$

This notion of stability is central to Harold Kushner's work on stochastic stability [10] which we draw on for this paper, and which formed the basis for recent martingale-based approaches to finite-time stability [11] and safety [12], [20], [13] for systems with unbounded uncertainty. This relates directly to the forward invariant region guaranteed to exist around the equilibrium point of ISS systems.

C. ISS for Bounded Disturbance Distributions

If the disturbance distribution \mathcal{D} for system (1) is only supported on a bounded set, then the essential supremum $\|\mathcal{D}\|_{L^\infty}$ is well defined; thus if a system satisfies the ISS condition (2), it is said to be stable in the ISS sense. Several authors have

worked to extend ISS to the setting of unbounded stochastic disturbances. [21] proposed an ISS condition for continuous-time systems with unbounded disturbances, but required the disturbance magnitude to be upper bounded by a class- \mathcal{K} of the state norm (thus, the disturbance vanishes at the equilibrium, a common but restrictive assumption). [22], [23] also study stochastic variants of ISS, but only require that the (ISS) condition hold for the expected trajectory (which does not guarantee boundedness of any trajectories).

IV. INPUT-TO-STATE STABILITY FOR UNBOUNDED RANDOM DISTURBANCES

In this paper, we seek to generalize the notion of input-to-state stability to systems that are subject to unbounded random disturbances. Specifically, two issues arise when the support of \mathcal{D} is unbounded: (i) the essential supremum $\|\mathcal{D}\|_{L^\infty}$ may not be well defined, rendering the ISS condition inapplicable, and (ii) the probability that \mathbf{x}_k remains in any bounded set for all $k \in \mathbb{N}$ is zero in general.

This second point is somewhat non-intuitive; however, consider a system with additive Gaussian noise, $\mathbf{x}_{k+1} = \hat{\mathbf{f}}(\mathbf{x}_k) + \mathbf{d}_k$, with $\mathbf{d}_k \sim \mathcal{N}(\mu, \Sigma)$ for some $\mu \in \mathbb{R}^n$ and $\Sigma = \Sigma^T > 0$. Then, since the tails of \mathbf{d} are unbounded, for any $B > 0$, $\mathbb{P}\{\|\mathbf{d}\| > 2B\} = \epsilon > 0$. This means, with probability ϵ , $\|\mathbf{f}(\mathbf{x}, \mathbf{d})\| \geq \|\mathbf{d}\| - \|\hat{\mathbf{f}}(\mathbf{x})\| > 2B - B = B$.² Thus, for any $K \in \mathbb{N}_{\geq 0}$,

$$\begin{aligned} \mathbb{P}\{\|\mathbf{x}_k\| < B, \forall k \leq K\} &\leq \mathbb{P}\{\|\mathbf{d}_k\| \leq 2B, \forall k \leq K\} \\ &= (1 - \epsilon)^K, \end{aligned}$$

since all \mathbf{d}_k are independent. Thus, as $K \rightarrow \infty$, the probability of the state remaining bounded goes to zero.

Thus, when generalizing ISS to the case of unbounded disturbances, we should expect weaker guarantees than those provided by the typical condition (ISS). With this in mind, we now define such a notion, *Input-to-State Stability in Probability (ISSp)*, which is well-defined for systems subject to unbounded noise.

Definition 7 (Input-to-State Stable in Probability). The system (1) is *input-to-state stable in probability (ISSp)* with respect to L^p if, for any $\epsilon \in (0, 1)$, $K \in \mathbb{N}_{\geq 0}$ and distribution $\mathcal{D} \in L^p$ such that $\|\mathcal{D}\|_{L^p}$, there exist functions $\beta \in \mathcal{KL}$, and $\gamma \in \mathcal{K}$ such that

$$\begin{aligned} \mathbb{P}\left\{\|\mathbf{x}_k\| \leq \beta(\|\mathbf{x}_0\|, k) + \gamma(\|\mathcal{D}\|_{L^p}), \forall k \leq K\right\} \\ \geq 1 - \epsilon. \end{aligned} \quad (9)$$

If this holds for $\beta(\|\mathbf{x}\|, k) = M\alpha^k \|\mathbf{x}\|$, for $M > 0, \alpha \in (0, 1)$, the system is exponentially input-to-state stable in probability (ISSp).

ISSp is a generalization of ISS to systems with (unbounded) stochastic disturbances. Intuitively, a system is ISSp if, for any disturbance in L^p , and for any finite horizon K , there exist β, γ such that the ISS condition (ISS) (with the L^∞ norm relaxed to the L^p) holds with a probability arbitrarily close to 1.

²We must have $\|\hat{\mathbf{f}}(\mathbf{x}_k)\| \leq B$ for $\|\mathbf{x}_k\| \leq B$; otherwise deterministic trajectories starting at \mathbf{x}_k would leave the set in one step.

As with ISS, we now relate ISSp to Lyapunov functions which can be used to verify this property.

Definition 8 (ISSp Lyapunov Function). A continuous function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ is an *ISSp Lyapunov Function* for the system (1) if there exist functions $\kappa_1, \kappa_2, \kappa_3 \in \mathcal{K}_{\infty}$ and $\kappa_4 \in \mathcal{K}$ such that,

$$\kappa_1(\|\mathbf{x}\|) \leq V(\mathbf{x}) \leq \kappa_2(\|\mathbf{x}\|) \quad (10)$$

$$\mathbb{E}[V(\mathbf{f}(\mathbf{x}, \mathbf{d}) - V(\mathbf{x}))] \leq -\kappa_3(V(\mathbf{x})) + \kappa_4(\|\mathcal{D}\|_{L^p}) \quad (11)$$

for all $\mathbf{x} \in \mathcal{X}$ and $\|\mathcal{D}\|_{L^p} < \infty$. Additionally, if there exist constants $a, b, c > 0$ and $\alpha \in (0, 1)$ such that $\kappa_1(r) = ar^c$, $\kappa_2(r) = br^c$, and $\kappa_3(r) = \alpha r$, then V is an *Exponential ISSp (E-ISSp) Lyapunov Function*

Remark. As in the typical ISS definition (ISS), since $\max\{a, b\} \leq a + b \leq \max\{2a, 2b\}$, for suitable choices of β, γ , the ISSp condition (9) is equivalent to

$$\mathbb{P}\left\{\|\mathbf{x}_k\| \leq \max\{\beta(\|\mathbf{x}_0\|, k), \gamma(\|\mathcal{D}\|_{L^p})\}, \forall k \leq K\right\} \geq 1 - \epsilon. \quad (12)$$

In this paper, for simplicity of exposition, we will consider exponential ISSp. Note that the results presented apply in the more general case, but the proofs become more complex.

V. LYAPUNOV CONDITIONS FOR E-ISSP

As for ISS, there exist Lyapunov conditions for E-ISSp. To this end, we will use tools from martingale theory (in particular, Ville's inequality) to demonstrate that the existence of a Lyapunov function satisfying a drift condition in expectation implies a system is E-ISSp.

Theorem 3. *If there exists an E-ISSp Lyapunov function for system (1), then system (1) is E-ISSp.*

Proof. We begin by constructing a nonnegative supermartingale $W(\mathbf{x}_k, k)$ via a time-varying, affine transform of the Lyapunov function $V(\mathbf{x}_k)$. Rearranging the Lyapunov drift condition (11) with $\kappa_3(r) = \alpha r$ for some $\alpha \in (0, 1)$, we can see that $V(\mathbf{x}_k)$ almost resembles a supermartingale³,

$$\mathbb{E}[V(\mathbf{x}_{k+1}) | \mathbf{x}_k] \leq (1 - \alpha)V(\mathbf{x}_k) + \varphi \quad (13)$$

where we define $\varphi \triangleq \kappa_4(\|\mathcal{D}\|_{L^p}) \geq 0$. However, this is not exactly a supermartingale due to the $(1 - \alpha)$ scaling and the additive constant φ .

Thus, for a particular horizon $K \in \mathbb{N}_{\geq 0}$, we construct $W(\mathbf{x}_k, k)$ by undoing this scaling and translation. Letting $W_k \triangleq W(\mathbf{x}_k, k)$ for simplicity, this construction is:

$$W_k = \underbrace{\theta^k V(\mathbf{x}_k)}_{\text{rescale}} - \underbrace{\varphi \sum_{i=1}^k \theta^i}_{\text{translate}} + \underbrace{\varphi \sum_{i=1}^K \theta^i}_{\text{ensure } W_k \geq 0}, \quad (14)$$

with $\theta \triangleq \frac{1}{1-\alpha} > 0$ and the constant term $\varphi \sum_{i=1}^K \theta^i$ added to ensure $W_k \geq 0$.

³Note that $\mathbb{E}[V(\mathbf{x}_{k+1}) | \mathbf{x}_k] = \mathbb{E}[V(\mathbf{x}_{k+1}) | \mathbf{x}_{0:k}]$ since system (1) is Markovian.

Next we show W_k is a nonnegative supermartingale. We have $W_k \geq 0$ for any $\mathbf{x}_k \in \mathcal{X}$, since $V(\mathbf{x}_k) \geq 0$ by definition, and $\theta, \varphi \geq 0$. Further, we have

$$\mathbb{E}[W_{k+1} | \mathbf{x}_k] = \mathbb{E}\left[\theta^{k+1}V(\mathbf{x}_{k+1}) + \varphi \sum_{i=k+2}^K \theta^i\right] \quad (15)$$

$$\leq \theta^{k+1}((1 - \alpha)V(\mathbf{x}_k) + \varphi) + \varphi \sum_{i=k+2}^K \theta^i \quad (16)$$

$$= \theta^k V(\mathbf{x}_k) + \varphi \sum_{i=k+1}^K \theta^i = W_k, \quad (17)$$

where the inequality (16) follows from the drift condition (11) and (17) uses the fact that $\theta = \frac{1}{1-\alpha}$.

Since W_k is a nonnegative supermartingale, we can apply Ville's inequality (7) to bound the probability W_k that remains below any $\lambda > 0$ for all $k \leq K$. Specifically,

$$\mathbb{P}\left\{W(\mathbf{x}_k) \leq \lambda, \forall k \leq K\right\} \geq 1 - \frac{W(\mathbf{x}_0)}{\lambda}. \quad (18)$$

We also note that, using the geometric series identity $\sum_{i=1}^k \theta^{i-1} = \frac{\theta^k - 1}{\theta - 1}$, we can write W_k as

$$W_k = \theta^k V(\mathbf{x}_k) + \frac{\theta \varphi (\theta^K - \theta^k)}{\theta - 1}. \quad (19)$$

Examining the structure of W_k , if for all $k \leq K$ we have $W_k \leq \lambda$, rearranging this inequality results in

$$V(\mathbf{x}_k) \leq \left(\lambda - \frac{\theta^{K+1} \varphi}{\theta - 1}\right) \theta^{-k} + \frac{\theta}{\theta - 1} \varphi \quad (20)$$

$$\leq (\lambda - \varphi) \theta^{-k} + \frac{\theta}{\theta - 1} \varphi \quad (21)$$

$$\triangleq (M \|\mathbf{x}_0\|^c + \eta \varphi) \theta^{-k} + \frac{\theta}{\theta - 1} \varphi \quad (22)$$

$$\leq M \|\mathbf{x}_0\|^c \theta^{-k} + \eta \varphi + \frac{\theta}{\theta - 1} \varphi \quad (23)$$

Inequality (21) follows from $\theta > 1$ and $\varphi \geq 0$. Equality (22) follows from choosing $\lambda = M \|\mathbf{x}_0\|^c + (1 + \eta) \varphi$ for some $\eta \geq 0$ and $M, c > 0$. Inequality (23) is due to $\theta > 1$ and $\varphi, \eta \geq 0$.

Further, using the lower bound (10) on $V(\mathbf{x}_k)$ and the definition of θ , (23) becomes

$$a \|\mathbf{x}_k\|^c \leq M \|\mathbf{x}_0\|^c (1 - \alpha)^k + \left(\eta + \frac{1}{\alpha}\right) \varphi \quad (24)$$

for some $a > 0$ which, rearranging, and raising both sides to the power of $\frac{1}{c}$ (which preserves order since $c > 0$), yields

$$\|\mathbf{x}_k\| \leq \left(\frac{M}{a} \|\mathbf{x}_0\|^c (1 - \alpha)^k + \frac{(\eta + \frac{1}{\alpha}) \varphi}{a}\right)^{\frac{1}{c}} \quad (25)$$

$$\leq \left(\frac{M}{a}\right)^{\frac{1}{c}} \zeta \|\mathbf{x}_0\| (1 - \alpha)^{\frac{k}{c}} + \zeta \left(\frac{(\eta + \frac{1}{\alpha}) \varphi}{a}\right)^{\frac{1}{c}} \quad (26)$$

$$\triangleq \tilde{M} \tilde{\alpha}^k \|\mathbf{x}_0\| + \gamma_{\eta}(\|\mathcal{D}\|_{L^p}), \quad (27)$$

for $\tilde{M} \triangleq \zeta (M/a)^{\frac{1}{c}} > 0$, $\tilde{\alpha} \triangleq (1 - \alpha)^{\frac{1}{c}} \in (0, 1)$, some $\zeta > 0$ as needed, and $\gamma_{\eta}(r) \triangleq \zeta \left(\frac{\eta + \frac{1}{\alpha}}{a}\right)^{\frac{1}{c}} \kappa_4(r)^{\frac{1}{c}}$ which is a class

\mathcal{K} for all $\eta \geq 0$. The existence of ζ follows from Lemma A shown in the Appendix.

Thus, we now must ensure there exists a suitable choice of M, η such that the probability that this bound holds for all $k \in \{1, \dots, K\}$ is greater than $1 - \epsilon$. By Ville's inequality,

$$\begin{aligned} & \mathbb{P} \left\{ \|\mathbf{x}_k\| \leq \tilde{M}\tilde{\alpha}^k \|\mathbf{x}_0\| + \gamma_\eta(\|\mathcal{D}\|_{L^p}), \forall k \leq K \right\} \\ & \geq 1 - \frac{W_0}{\lambda} = 1 - \frac{V(\mathbf{x}_0) + \frac{\varphi}{\alpha}((1-\alpha)^{-K} - 1)}{M \|\mathbf{x}_0\|^c + (1+\eta)\varphi}, \end{aligned} \quad (28)$$

with $W_0 \geq 0$ by definition. Thus, as long as $\|\mathbf{x}_0\|^c, \|\mathcal{D}\|_{L^p}$ are not both zero, we can choose M, η large enough to have

$$\mathbb{P} \left\{ \|\mathbf{x}_k\| \leq \tilde{M}\tilde{\alpha}^k \|\mathbf{x}_0\| + \gamma_\eta(\|\mathcal{D}\|_{L^p}), k \leq K \right\} \geq 1 - \epsilon,$$

for any $\epsilon \in (0, 1)$, so the system must be E-ISSp. \square

Remark. The variables $M \geq 0$ and $\eta \geq 0$ are free parameters which can be varied to analyze the probability of convergence and boundedness respectively. We note that the bound in (29) may be very weak; stronger bounds can be achieved by removing the bounding steps in (21) and (23), but clarity was chosen over tightness for this proof.

VI. CONNECTIONS TO OTHER STABILITY NOTIONS

Here we discuss connections between ISS, ISSp, and other notions of stability for stochastic systems, as surveyed in Section III.

Corollary 1. If the system (1) is ISS, then it is ISSp with respect to L^∞ .

Proof. By definition, if a system is ISS, then for all $k \in \mathbb{N}_{\geq 0}$, there exist $\beta \in \mathcal{KL}, \gamma \in \mathcal{K}$ such that

$$\|\mathbf{x}_k\| \leq \beta(\|\mathbf{x}_0\|, k) + \gamma(\vartheta) \quad (29)$$

for all $\vartheta \geq \sup_{k \in \mathbb{N}_{\geq 0}} \|\mathbf{d}_k\|$.

Thus, since the L^∞ -norm (equivalently, the essential supremum) is finite for all random variables in L^∞ , for any $\mathbf{d} \sim \mathcal{D}$ with $\mathbf{d} \in L^\infty$, we have:

$$\mathbb{P} \left\{ \|\mathbf{x}_k\| \leq \beta(\|\mathbf{x}_0\|, k) + \gamma(\|\mathcal{D}\|_{L^\infty}), \forall k \in \mathbb{N} \right\} = 1.$$

Thus trivially, for any $K \in \mathbb{N}_{\geq 0}, \epsilon \in (0, 1)$, we have $\beta \in \mathcal{KL}, \gamma \in \mathcal{K}$ such that (9) holds for all distributions \mathcal{D} with $\mathbf{d} \in L^\infty$. Therefore, the system is ISSp w.r.t. L^∞ . \square

Corollary 1 provides a clear connection between ISS and ISSp: if a system is ISS, it by definition is ISSp for disturbances in L^∞ . Next, we discuss the relationship between ISS and ISSp w.r.t. L^2 which is a much larger class of unbounded random variables.

Corollary 2. If the system (1) is additive with respect to its disturbance and admits a twice-continuously differentiable, convex E-ISS Lyapunov function $V : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}$ such that $\sup_{\mathbf{x} \in \mathcal{X}} \|\nabla^2 V(\mathbf{x})\|_2 \leq \lambda_{\max}$ for some $\lambda_{\max} \geq 0$, then it is E-ISSp for $\mathbf{d} \in L_p$ with $\mathbb{E}[\mathbf{d}] = 0$ for $p \geq 2$.

Proof. The dynamics are additive with respect to the disturbance so system (1) can be rewritten as:

$$\mathbf{x}_k = \mathbf{f}(\mathbf{x}_k, \mathbf{d}) \triangleq \hat{\mathbf{f}}(\mathbf{x}_k) + \mathbf{d}. \quad (30)$$

The function V is an E-ISS Lyapunov function for (30) so it satisfies

$$V(\hat{\mathbf{f}}(\mathbf{x}) + \mathbf{d}) - V(\mathbf{x}) \leq -\alpha V(\mathbf{x}) + \kappa_4(\|\mathcal{D}\|_{L^\infty}) \quad (31)$$

for all $\mathbf{x} \in \mathcal{X}$, some $\alpha \in (0, 1)$, $\sigma \in \mathcal{K}$, and any $\mathbf{d} \in L^\infty$.

The expected value of the left side of this inequality is:

$$\begin{aligned} \mathbb{E} \left[V(\hat{\mathbf{f}}(\mathbf{x}) + \mathbf{d}) - V(\mathbf{x}) \right] &= \mathbb{E} \left[\mathbf{V}(\hat{\mathbf{f}}(\mathbf{x}) + \mathbf{d}) \right] - V(\mathbf{x}) \\ &\leq V(\hat{\mathbf{f}}(\mathbf{x}) + \mathbb{E}[\mathbf{d}]) - V(\mathbf{x}) + \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{d})) \end{aligned} \quad (32)$$

$$= V(\hat{\mathbf{f}}(\mathbf{x}) + 0) - V(\mathbf{x}) + \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{d})) \quad (33)$$

$$\leq -\alpha V(\mathbf{x}) + \frac{\lambda_{\max}}{2} \text{tr}(\text{cov}(\mathbf{d})) \quad (34)$$

where (32) accounts for Jensen's inequality as in [13, Lemma 1], (33) is due to the assumption that the $\mathbb{E}[\mathbf{d}] = 0$, and (34) is an application of the E-ISS bound (31).

Since bounded covariance implies boundedness in L^2 , if $\mathbf{d} \in L^2$ then V is an E-ISSp Lyapunov function for (30). Furthermore, since $L^2 \supseteq L_p$ for $p \geq 2$, V is an E-ISSp Lyapunov function for $\mathbf{d} \in L_p$ for all $p \geq 2$. \square

Next we discuss the relationship between ISSp and trajectories that are bounded in probability.

Corollary 3. If system (1) is ISSp w.r.t. L^p , then for any $\mathbf{d} \in L^p$, the system's trajectories are bounded in probability.

Proof. If the system is ISSp w.r.t. L^p , then for any $K \in \mathbb{N}_{\geq 0}$ and $\epsilon \in (0, 1)$, there exist $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$ such that

$$\mathbb{P} \left\{ \|\mathbf{x}_k\| \leq \beta(\|\mathbf{x}_0\|, k) + \gamma(\|\mathcal{D}\|_{L^p}), \forall k \leq K \right\} \geq 1 - \epsilon.$$

Then, since β is decreasing in k , we have that for $B_0 \triangleq \beta(\|\mathbf{x}_0\|, 0) + \gamma(\|\mathcal{D}\|_{L^p})$,

$$\mathbb{P} \{ \|\mathbf{x}_k\| \leq B_0, \forall k \leq K \} \geq 1 - \epsilon. \quad (35)$$

Thus the system trajectories are bounded in probability. \square

Like with traditional ISS, the ISSp condition 9 is equivalent to system trajectories remaining in a ball whose radius scales with the initial condition and the norm of the disturbance. Thus, if a system is ISSp, its trajectories (over a finite horizon) must be bounded in probability.

Finally, we look to discuss the relationship between ISSp and recurrence. To do this we, first restate an important result from drift analysis (see [14] for a detailed survey).

Theorem 4 (Variable Drift [14]). *Suppose there exists some function $V : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$, with γ -sublevel set $V_\gamma \triangleq \{\mathbf{x} \in \mathcal{X} \mid V(\mathbf{x}) \leq \gamma\}$ such that for all $\mathbf{x} \in \mathcal{X} \setminus V_\gamma$,*

$$\mathbb{E} [V(\mathbf{f}(\mathbf{x}, \mathbf{d})) - V(\mathbf{x})] \leq -h(V(\mathbf{x})), \quad (36)$$

for some increasing function $h : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$. Then, for any trajectory with initial state \mathbf{x}_0 , the hitting time $\tau_\gamma(\mathbf{x}_0) = \inf\{k \mid V(\mathbf{x}_k) \leq \gamma\}$ is bounded in expectation by

$$\mathbb{E} [\tau_\gamma(\mathbf{x}_0)] \leq \frac{\gamma}{h(\gamma)} + \int_\gamma^{V(\mathbf{x}_0)} \frac{1}{h(\sigma)} d\sigma. \quad (37)$$




Fig. 2: Simulation results for the double integrator over 1500 trials. **(Top):** Lyapunov function values $V(\mathbf{x}_k)$ plotted against the upper bound trajectory for various M . **(Middle):** Corresponding supermartingale values W_k for each trajectory. Note the geometric upper bound trajectories for $V(\mathbf{x})$ correspond one-to-one to the level sets of W_k . **(Bottom left):** Ground tracks for the double integrator position. **(Bottom right):** Monte Carlo estimation of the “success fraction” (fraction of trajectories with $V(\mathbf{x}_k) \leq \rho_k$) versus the guaranteed lower bound.

Using this result, we can show that, if a system admits an E-ISSp Lyapunov function, then any Lyapunov sublevel set (above a particular value) must be recurrent.

Theorem 5. *If there exists an E-ISSp Lyapunov function w.r.t. L^p for system (1), then (1) is recurrent.*

Proof. Suppose there exists an E-ISSp Lyapunov function V for the system (1). Then, for $h(V(\mathbf{x}_k)) = \alpha V(\mathbf{x}_k) - \varphi$, with $\varphi \triangleq \kappa_4(\|\mathcal{D}\|_{L^p})$, we have

$$\mathbb{E}[V(\mathbf{x}_{k+1}) - V(\mathbf{x}_k) \mid \mathbf{x}_k] \leq -h(V(\mathbf{x}_k)). \quad (38)$$

For any $\gamma > \frac{\varphi}{\alpha}$, $h(V(\mathbf{x})) > 0$ for all $\mathbf{x} \in \mathcal{X} \setminus V_\gamma$; thus our system meets the variable drift condition (36).

Thus, consider some trajectory with an initial state $\mathbf{x}_0 \in \mathcal{X} \setminus V_\gamma$. Then, by Theorem 4, we have

$$\mathbb{E}[\tau_\gamma(\mathbf{x}_0)] \leq \frac{\gamma}{\alpha\gamma - \varphi} + \int_\gamma^{V(\mathbf{x}_0)} \frac{1}{\alpha\sigma - \varphi} d\sigma \quad (39)$$

$$\leq \frac{\gamma}{\alpha\gamma - \varphi} + \frac{1}{\alpha} \log\left(\frac{\alpha V(\mathbf{x}_0) - \varphi}{\alpha\gamma - \varphi}\right) < \infty. \quad (40)$$

Since $\mathbb{E}[\tau_\gamma(\mathbf{x}_0)] < \infty$, we must have $\mathbb{P}\{\tau_\gamma(\mathbf{x}_0) < \infty\} = 1$. Thus, for any $\gamma > \frac{\varphi}{\alpha}$, the sublevel set V_γ is recurrent. Since κ_1 is radially unbounded, V_γ must be bounded for all $\gamma \geq 0$, thus the system is recurrent. \square

VII. PRACTICAL EXAMPLE: LINEAR-QUADRATIC-GAUSSIAN CONTROL

We now consider the ISSp properties of some practical systems, and study the validity of our exit probability bounds via simulation. We begin by considering the case of linear time-invariant systems subject to additive, zero-mean Gaussian noise. The system dynamics are given by $\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k + \mathbf{d}_k$, where $\mathbf{u}_k \in \mathbb{R}^m$ is a control input to the system, and $\mathbf{d}_k \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \Sigma) = \mathcal{D}$ for some $\Sigma = \Sigma^T > 0$. Suppose our system implements an infinite-horizon LQR feedback policy, i.e., $\mathbf{u}_k = -\mathbf{K}\mathbf{x}_k$, where $\mathbf{K} = (\mathbf{R} + \mathbf{B}^T\mathbf{P}\mathbf{B})^{-1}(\mathbf{B}^T\mathbf{P}\mathbf{A})$ for \mathbf{P} satisfying the discrete-time algebraic Riccati equation. For this closed-loop system, $V(\mathbf{x}) = \mathbf{x}^T\mathbf{P}\mathbf{x}$ is an E-ISS Lyapunov function, with

$$\mathbb{E}[V(\mathbf{x}_{k+1}) - V(\mathbf{x}_k)] \leq -\alpha V(\mathbf{x}_k) + \sigma \|\mathcal{D}\|_{L^2}^2 \quad (41)$$

for $\alpha = \frac{\lambda_{\min}(\mathbf{Q})}{\lambda_{\max}(\mathbf{P})}$, $\sigma = \lambda_{\max}(\mathbf{P})$, and any $\mathcal{D} \in L^2$. In particular, if we pick $\lambda = MV(\mathbf{x}_0) + \frac{\sigma\|\mathcal{D}\|_{L^2}^2}{\alpha(1-\alpha)^K}$, we can bound the probability V_k rises above a time-varying trajectory:

$$\mathbb{P}\left\{V(\mathbf{x}_k) \leq \rho_k, \forall k \leq K\right\} \geq 1 - \frac{W_0}{\lambda}, \quad (42)$$

with $\rho_k = MV(\mathbf{x}_0)(1-\alpha)^k + \frac{\sigma\|\mathcal{D}\|_{L^2}^2}{\alpha}$, and W_0 defined as in (14). Specifically, we study how this bound varies numerically for a double integrator system in the plane (see [13] for a detailed dynamics derivation). Figure 2 plots the results of 1500 simulations of the double integrator. First, we plot the values of $V(\mathbf{x}_k)$ and W_k across multiple choices of M . We can see that the trajectories $W_k = \lambda$ correspond exactly to $V(\mathbf{x}_k) = \rho_k$; thus, intuitively, the event $W_k > \lambda$ is exactly the event where V_k rises above a shifted geometric sequence ρ_k .

We also plot the trajectories of the system in the plane, along with Lyapunov level sets V_ρ for $\rho = \max_k \rho_k$, evaluated when the system’s velocity is zero. We choose the maximum value of ρ_k since if $V_k \leq \rho_k$ for all $k \leq K$, we must have $V_k \leq \max_k \rho_k$. Interestingly, one can show the probability bound (42) is equivalent to the exit probability bound provided by Kushner [10], up to a choice of scaling to construct W_k (we explore this point further in Section VIII).

Finally, we plot our bound (42) on the “success probability” $\mathbb{P}\{V_k \leq \rho_k, \forall k \leq K\}$ versus the fraction of trajectories that remained under ρ_k for various choices of M . While we note that our bound is, indeed, a lower bound on the success probability, it is quite a weak lower bound; thus, finding Lyapunov functions and martingales that yield stronger bounds is an interesting direction for future work.

VIII. PRACTICAL EXAMPLE: SEVEN-LINK WALKER

Consider the seven-link walker as shown in Fig. 1. As detailed in [3], walking can be distilled down to the discrete-time dynamical system described by the Poincaré return map:

$$\mathcal{P} : B_\rho(\mathbf{x}^*) \times [d^-, d^+] \rightarrow S_{[d^-, d^+]} := \bigcup_{d \in [d^-, d^+]} S_d, \quad (43)$$

$$\mathbf{x}_{k+1} = \mathcal{P}(\mathbf{x}_k, d_k), \quad d_k \in [d^-, d^+],$$

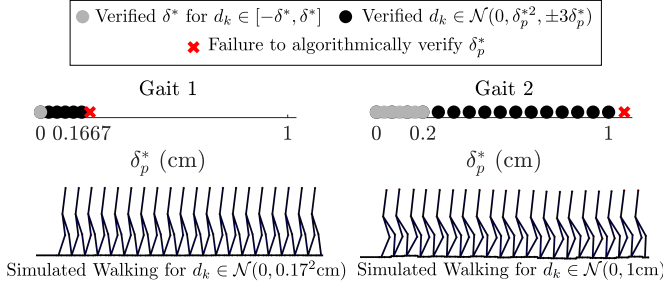


Fig. 3: Algorithmic results of ISSp compared to ISS. As shown, the stochastic ISSp condition yields more realistic predictions of the tolerable step heights for two gaits (the same gaits as those compared in [3]).

for some sequence of step heights $d_k \in [d^-, d^+] \subset \mathbb{R}$, $k \in \mathbb{N}_{\geq 0}$ and $S_d \subset \mathbb{R}^n$ denoting the uncertain guard condition:

$$S_d = \{\mathbf{x} \in \mathbb{R}^n \mid h(\mathbf{x}) = d, \dot{h}(\mathbf{x}) < 0\}, \quad (44)$$

where $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is typically selected to denote the vertical height of the swing foot relative to the stance foot. Note that the partial function nature of \mathcal{P} implies that solutions may not exist for all time, i.e., the solution \mathbf{x}_k might leave the ball $B_\rho(\mathbf{x}^*)$ on which \mathcal{P} is well-defined.

Guaranteed Robustness to Uncertain Terrain. Prior work considered input-to-state stability of (43) with bounded step heights described as the set $\mathbb{D} \triangleq [-\delta, \delta] \subset \mathbb{R}$ with $\delta > 0$ [3]. Moreover, leveraging this discrete-time representation for bounded disturbances [3] introduced an ISS perspective on bipedal locomotion. Explicitly, a periodic walking gait with a nominal fixed point $\mathbf{x}^* = \mathcal{P}(\mathbf{x}^*, 0)$ is defined as δ -robust for a given $\delta > 0$ if for the discrete-time dynamical system (43), with any $d_k \in [-\delta, \delta]$, there exists some forward invariant set $\mathcal{W} \subset B_\rho(\mathbf{x}^*)$ such that for all $\mathbf{x}_0 \in \mathcal{W}$, the system is ISS. Moreover, this definition of robustness was shown to be verifiable through an ISS Lyapunov function.

Specifically, Theorem 2 of [3] states that if the Lyapunov condition as in (4) is satisfied, then the periodic gait is ISS. To verify this, a candidate Lyapunov function can be synthesized by approximating the exponentially stable discrete-time system using the linearization of the Poincaré return map for $d_k = 0$:

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k := D\mathcal{P}(0, 0)\mathbf{x}_k.$$

Then, the Lyapunov matrix $\mathbf{P} = \mathbf{P}^T > 0$ is obtained by solving the discrete-time Lyapunov equation $(\mathbf{A}^T\mathbf{P}\mathbf{A} - \mathbf{P} = -\mathbf{Q})$ for $\mathbf{Q} = \mathbf{Q}^T > 0$ which provides a discrete-time Lyapunov function $V(\mathbf{x}) = \mathbf{x}^T\mathbf{P}\mathbf{x}$.

Probabilistic Robustness to Uncertain Terrain. To obtain more reasonable estimates of the maximum step heights that a given periodic gait can withstand, in this work we will instead consider step heights drawn from some distribution $d_k \sim \mathcal{D}$ and apply the ISSp methodology. Specifically, we take $\mathcal{D} := \mathcal{N}(0, \delta_p^2)$ such that $\delta_p > 0$ now represents the standard deviation of the distribution. Note that due to the partial nature of the \mathcal{P} , we will truncate \mathcal{D} at $3\delta_p$ to ensure that there exists some δ_p such that $\mathbb{E}[\Delta V] < +\infty$, we will denote this truncated Gaussian as $\mathcal{N}(\cdot, \cdot, \pm a)$ where a denotes the truncation interval [16].

Consider the ISSp Lyapunov condition from Def. (8):

$$\mathbb{E}[V(\mathcal{P}(\mathbf{x}, d)) - V(\mathbf{x}) \mid \mathbf{x}] \leq -\tilde{\alpha}V(\mathbf{x}) + \tilde{\sigma}, \quad (45)$$




Fig. 4: The probabilistic bound, evaluated for Gait 2, quickly decreases as δ increases. However, the simulation results show that the walking is able to remain periodic for all values determined to be δ -robust in Opt. 47, highlighting that ISSp yields more reasonable estimates of δ compared to the strict ISS condition presented in [3].

with $\tilde{\alpha} = 2k\lambda_{\max}(\mathbf{P})$, and $\tilde{\sigma} = k(\chi\delta_p)^c$, where $k \in (0, 1)$ is a user-defined variable dictating the convergence of the Lyapunov condition and $\chi > 0$ is used to scale the set over which the expected exponential decay condition holds. Note that this condition can be equivalently expressed in the form:

$$\begin{aligned} \|\mathbf{x} - \mathbf{x}^*\| \geq \chi\delta_p &\implies \\ \mathbb{E}[V(\mathcal{P}(\mathbf{x}, d)) - V(\mathbf{x}) \mid \mathbf{x}] &\leq -k\|\mathbf{x} - \mathbf{x}^*\|^c. \end{aligned} \quad (46)$$

This ISSp Lyapunov function can be utilized with the algorithmic approach introduced in [3] to solving the following optimization problem:

$$\begin{aligned} (\delta_p^*, \chi_p^*) &= \operatorname{argmax}_{\delta_p, \chi_p > 0} \delta_p & (47) \\ \text{s.t. } &\mathbb{E}_{d \sim \mathcal{D}}[V(\mathcal{P}(\mathbf{x}, d)) - V(\mathbf{x}) \mid \mathbf{x}] \leq -k\|\mathbf{x}\|^2 \\ &\forall \|\mathbf{x}\| = \chi_p\delta_p, \quad d \sim \mathcal{N}(0, \delta_p^2, \pm 3\delta_p), \end{aligned}$$

As shown in Fig. 3, the algorithmic approach to the updated optimization (47) results in more reasonable estimates of the maximum tolerable step height for each of the two gaits considered in [3].

Probabilistic Guarantees for ISSp. While relaxing the Lyapunov condition to the one in (47) yields more realistic estimates of δ^* , this relaxed condition no longer satisfies the assumptions needed to be provably ISS. Instead, we can use probabilistic bounds to assert that the system is ISSp.

To do this, we first need to approximate a reasonable estimate of the Lyapunov level set that bounds the evolution of the system after K steps. Rearranging the Lyapunov condition (45) for the largest χ_p^* and δ_p^* identified by Opt. (47), and using the fact that $\mathcal{W} \triangleq \{\mathbf{x} \mid V(\mathbf{x}) \leq \kappa_2(\chi\delta)^c\}$ in Theorem 2 of [3], we obtain the Lyapunov bound:

$$V(\mathbf{x}_K) \leq \tilde{\rho} \triangleq (1 - \alpha)^K \lambda_{\max}(P)(\chi_p^*\delta_p^*)^2 + k(\chi_p^*\delta_p^*)^2.$$

As discussed in Sec. VII, the probabilistic bound associated with remaining within this Lyapunov level set can be obtained from Kushner [10]. Importantly, when your Lyapunov level set $\tilde{\rho}$ is less than $\tilde{\sigma}/\alpha$, the bound (42) is extremely conservative. Following Kushner, it is possible to find a better choice of W_k that yields a better probability bound. Specifically, one can use the bound:

$$\begin{aligned} &\mathbb{P}(V(\mathbf{x}_K) \leq \tilde{\rho}), \quad \forall k \leq K \\ &\geq \begin{cases} \frac{\tilde{\rho} - V(\mathbf{x}_0)}{\tilde{\rho}} \left(\frac{\tilde{\rho} - \tilde{\sigma}}{\tilde{\rho}}\right)^K, & \tilde{\rho} \geq \frac{\tilde{\sigma}}{\alpha} \\ 1 - \frac{V(\mathbf{x}_0)(1-\alpha)^K + \tilde{\sigma} \sum_{i=1}^K (1-\alpha)^{i-1}}{\tilde{\rho}}, & \text{otherwise.} \end{cases} \end{aligned} \quad (48)$$

with $\tilde{\sigma}$ determined for each δ_p via Monte Carlo sampling.

In Fig. 4, the probabilistic bound is illustrated for the values (χ_p^*, δ_p^*) obtained using (47) for Gait 2. To verify the probabilistic bound, Monte Carlo sampling was implemented to estimate the true probability that the system remains within $\tilde{\rho}$ after $K = 10$ steps. We simulate the system for this horizon and report both the fraction of trajectories remaining stable, as well as the fraction of trajectories remaining in the Lyapunov sublevel set $\tilde{\rho}$.

IX. CONCLUSION

In this paper, we introduced the notion of input-to-state stability in probability (ISSp), which generalizes ISS for discrete-time systems with unbounded stochastic disturbances. We provided Lyapunov conditions for exponential ISSp, drew connections between ISS, ISSp, and traditional stability notions for stochastic systems, and provided simulation studies of ISSp systems (including an LQG system and bipedal walker) where we provide practical, probabilistic stability guarantees for systems subject to random disturbances.

This work opens numerous directions for future work. In particular, while the martingale-based tools used in this paper require only a simple expected drift condition (4) on the Lyapunov function, martingale-based probability bounds are typically considered quite weak. Thus, for practitioners, an important question is how to choose Lyapunov function V , and the supermartingale W , to obtain the tightest bound possible (see also sum-of-squares-based approaches [11], [20]). It is also likely that tighter bounds may be obtained by exploiting particular structure in the disturbance distribution and dynamics.

Another direction of interest is to investigate how stability notions degrade when the system is subject to state uncertainty (as compared to the process noise considered in this work). Like for ISSp, we would expect these stability guarantees to hold only in probability; it remains an open question whether a similar property to the “smooth degradation” observed under ISS and ISSp can be found for systems with uncertain state.

Lastly, the application of ISSp to bipedal locomotion motivates its use for studying the robustness of periodic walking gaits to uncertain terrain. Future work in this area includes applying the notion of ISSp to the gait synthesis framework to systematically generate nominal walking trajectories that have probabilistic guarantees of robustness for reasonable estimates of uncertain terrain.

REFERENCES

- [1] E. Sontag, “Smooth Stabilization Implies Coprime Factorization,” *Automatic Control, IEEE Transactions on*, vol. 34, pp. 435–443, May 1989.
- [2] E. D. Sontag and Y. Wang, “On characterizations of the input-to-state stability property,” *Systems & Control Letters*, vol. 24, pp. 351–359, Apr. 1995.
- [3] M. Tucker and A. D. Ames, “An input-to-state stability perspective on robust locomotion,” *IEEE Control Systems Letters*, 2023.
- [4] E. Sontag, “On the Input-to-State Stability Property,” *European Journal of Control*, vol. 1, Nov. 1995.
- [5] Z.-P. Jiang, E. Sontag, and Y. Wang, “Input-to-state stability for discrete-time nonlinear systems,” *IFAC Proceedings Volumes*, vol. 32, pp. 2403–2408, July 1999.
- [6] D. Nesic and A. R. Teel, “Input-to-state stability of networked control systems.”

- [7] D. Angeli, E. Sontag, and Y. Wang, “A characterization of integral input-to-state stability,” *IEEE Transactions on Automatic Control*, vol. 45, pp. 1082–1097, June 2000. Conference Name: IEEE Transactions on Automatic Control.
- [8] W.-L. Ma, S. Kolathaya, E. R. Ambrose, C. M. Hubicki, and A. D. Ames, “Bipedal robotic running with DURUS-2D: Bridging the gap between theory and experiment,” in *Intl. Conference on Hybrid Systems: computation and control*, pp. 265–274, 2017.
- [9] S. Kolathaya, J. Reher, A. Hereid, and A. D. Ames, “Input to state stabilizing control lyapunov functions for robust bipedal robotic locomotion,” in *2018 American Control Conference*, pp. 2224–2230, IEEE, 2018.
- [10] H. Kushner, *Stochastic Stability and Control*. Mathematics in Science and Engineering, Academic Press, 1967.
- [11] J. Steinhardt and R. Tedrake, “Finite-time regional verification of stochastic non-linear systems,” *The International Journal of Robotics Research*, vol. 31, pp. 901–923, June 2012. Publisher: SAGE Publications Ltd STM.
- [12] C. Santoyo, M. Dutreix, and S. Coogan, “Verification and Control for Finite-Time Safety of Stochastic Systems via Barrier Functions,” May 2019. arXiv:1905.12077 [cs].
- [13] R. K. Cosner, P. Culbertson, A. J. Taylor, and A. D. Ames, “Robust Safety under Stochastic Uncertainty with Discrete-Time Control Barrier Functions,” in *Robotics: Science and Systems*, July 2023.
- [14] J. Lengler, “Drift Analysis,” in *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization* (B. Doerr and F. Neumann, eds.), Natural Computing Series, pp. 89–131, Cham: Springer International Publishing, 2020.
- [15] S. P. Meyn and R. L. Tweedie, *Markov Chains and Stochastic Stability*. London: Springer, 1993.
- [16] G. Grimmett and D. Stirzaker, *Probability and Random Processes*. Oxford University Press, July 2020.
- [17] R. L. Wheeden and A. Zygmund, *Measure and Integral: an Introduction to Real Analysis*, vol. 308. CRC Press, 2015.
- [18] J. Ville, “Etude critique de la notion de collectif,” *Bull. Amer. Math. Soc.*, vol. 45, no. 11, p. 824, 1939.
- [19] H. J. Kushner, “Some Basic Ideas in Stochastic Stability,” *Annals of Economics and Social Measurement*, vol. 3, no. 1, pp. 85–90, 1974.
- [20] C. Santoyo, M. Dutreix, and S. Coogan, “A barrier function approach to finite-time stochastic system verification and control,” *Automatica*, vol. 125, p. 109439, Mar. 2021.
- [21] J. Tsiniias, “Stochastic input-to-state stability and applications to global feedback stabilization,” *International Journal of Control*, vol. 71, pp. 907–930, Jan. 1998.
- [22] Y. Tang, X. Wu, P. Shi, and F. Qian, “Input-to-state stability for nonlinear systems with stochastic impulses,” *Automatica*, vol. 113, p. 108766, Mar. 2020.
- [23] R. D. McAllister and J. B. Rawlings, “Stochastic Exponential Stability of Nonlinear Stochastic Model Predictive Control,” in *2021 60th IEEE Conference on Decision and Control (CDC)*, pp. 880–885, Dec. 2021.
- [24] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge, UK ; New York: Cambridge University Press, 2004.

APPENDIX

A. Lemma for Lyapunov Conditions

We use the following lemma to prove the sufficiency of the ISSp Lyapunov conditions in Theorem 3.

Lemma 1. For $x_1, x_2, p > 0$, there exists $\zeta > 0$ such that $(x_1^p + x_2^p)^{\frac{1}{p}} \leq \zeta(x_1 + x_2)$.

Proof. Case 1: Suppose $p \geq 1$. Then $(|x_1|^p + |x_2|^p)^{\frac{1}{p}} \triangleq \|\mathbf{x}\|_p$ defines the ℓ_p norm for $\mathbf{x} \triangleq [x_1, x_2]^T$ on \mathbb{R}^2 . Since ℓ_p norms are equivalent [24], there exists $\zeta > 0$ such that $\|\mathbf{x}\|_p \leq \zeta \|\mathbf{x}\|_1 = \zeta(|x_1| + |x_2|)$. The result follows since $x_1, x_2 > 0$.

Case 2: Suppose $0 < p < 1$. Then, x^p is a concave function [24], and $\left(\frac{x_1^p + x_2^p}{2}\right) \leq \left(\frac{x_1 + x_2}{2}\right)^p$. Since $x^{\frac{1}{p}}$ is an increasing function, exponentiating both sides preserves ordering, $\left(\frac{x_1^p + x_2^p}{2}\right)^{\frac{1}{p}} \leq \frac{x_1 + x_2}{2}$, and thus we have $(x_1^p + x_2^p)^{\frac{1}{p}} \leq 2^{\frac{1}{p}-1}(x_1 + x_2) \triangleq \zeta(x_1 + x_2)$, with $\zeta > 0$ as needed. \square