

Metastable Walking on Stochastically Rough Terrain

Katie Byl and Russ Tedrake

Abstract—Simplified models of limit-cycle walking on flat terrain have provided important insights into the nature of legged locomotion. Real walking robots (and humans), however, do not exhibit true limit cycle dynamics because terrain, even in a carefully designed laboratory setting, is inevitably non-flat. Walking systems on stochastically rough terrain may not satisfy strict conditions for limit-cycle stability but can still demonstrate impressively long-living periods of continuous walking. Here, we examine the dynamics of rimless-wheel and compass-gait walking on randomly generated rough terrain and employ tools from stochastic processes to describe the ‘stochastic stability’ of these gaits. This analysis generalizes our understanding of walking stability and may provide statistical tools for experimental limit cycle analysis on real walking systems.

I. INTRODUCTION

The science of legged locomotion is plagued with complexity. Many of the fundamental results for legged robots have come from detailed analytical and computational investigations of simplified models (e.g., [3, 5, 7, 10, 11]). These analyses reveal the limit cycle nature of ideal walking systems and employ Poincaré map analysis to assess the stability of these limit cycles. However, the very simplifications which have made these models tractable for analysis can limit their utility.

Experimental analyses of real machines based on these simple models [4] have revealed that real machines differ from these idealized dynamics in a number of important ways. Certainly the dynamics of impact and contact with the ground are more subtle than what is captured by the idealized models. But perhaps more fundamental is the inevitable stochasticity in the real system. More than just measurement noise, robots that walk are inherently prone to the stochastic influences of their environment by interacting with terrain which varies at each footstep. Even in a carefully designed laboratory setting, and especially for passive and minimally-actuated walking machines, the effects of this stochasticity can have a major effect on the long-term system dynamics. In practice, it is very difficult (and technically incorrect) to apply deterministic limit cycle stability analyses to our experimental walking machines - the real machines do not have true limit cycle dynamics.

In this paper, we extend the analysis of simplified walking models toward real machines by adding stochasticity into the our model. Although we have considered a number of sources of uncertainty, we will focus here on a compact and demonstrative model - where the geometry of the ground is drawn from a random distribution. Even with mild deviations in terrain from a nominal slope angle, the resulting trajectories of the machine are different on every step and for many noise distributions (e.g., Gaussian) the robot is guaranteed to eventually fall down (with probability one as $t \rightarrow \infty$). However, one

can still meaningfully quantify stochastic stability, in terms of expected time to failure, and maximization of this metric in turn provides a parameter for optimization in the design of control for a walking robot on moderately rough, unmodeled terrain.

II. BACKGROUND

Many stochastic dynamic systems exhibit behaviors which are impressively long-living, but which are also guaranteed to exit these behaviors (“fail”) with probability one given enough time. Such systems cannot be classified as “stable”, but it is also misleading and incomplete to classify them as “unstable”. Physicists have long used the term *metastable* to capture this interesting phenomenon and have developed a number of tools for quantifying this behavior [8, 9, 12, 15]. Many other branches of science and engineering have also borrowed the terminology to describe dynamic systems in a wide variety of fields. Familiar metastable systems include crystalline structures (e.g. diamonds), flip-flop circuits, radioactive elements, oscillatory wave patterns in the brain, and ferromagnetic materials, such as spin glass or magnetic tape film (which explains why a taped recording sitting in storage still inevitably fades over time).




Fig. 1. Cartoon of a particle subject to Brownian motion in a potential $U(x)$ with two metastable states, A and B .

The canonical example of metastability is a particle in a potential well subject to Brownian motion, as cartooned in Figure 1. These systems have local attractors which tend to keep the dynamics within a particular neighborhood in state space. In the limit as such systems become deterministic (no noise), these local attractors are fixed points, and the system is truly stable whenever the dynamics begin with an initial condition somewhere inside the basin of attraction of the fixed point. In contrast, stochasticity constantly pushes the dynamics about within this neighborhood, and for some systems and noise types, this turns a stable system into a

metastable one. Occasionally but repeatedly, such systems will deviate particularly far from a metastable attractor in state space (making “escape attempts”), and eventually, they will successfully exit (by which we mean entering a region where a different attractor is now a far more dominating influence).

III. METASTABLE LIMIT CYCLE ANALYSIS

The dynamics of walking systems are continuous, but they are punctuated by discrete impact events when a foot comes into contact with the ground. These impacts provide a natural time-discretization of a gait onto a Poincaré map. Therefore, we will consider walking systems governed by the discrete, closed-loop return-map dynamics:

$$\mathbf{x}[n+1] = \mathbf{f}(\mathbf{x}[n], \gamma[n]), \quad (1)$$

where $\mathbf{x}[n]$ denotes the state of the robot at step n and $\gamma[n]$ represents the slope of the ground, which is a random variable drawn independently from a distribution P_γ at each n . This model for stochastically rough terrain dramatically simplifies our presentation in this paper, but it also restricts our analysis to strictly forward walking¹. These state evolution equations represent a discrete-time, continuous-state Markov process (or infinite Markov chain). For computational purposes, we will also discretize the state space into a finite set of states, x_i . Defining the state distribution vector, $\mathbf{p}[n]$, as

$$p_i[n] = \Pr(\mathbf{X}[n] = x_i), \quad (2)$$

we can describe the state distribution (master) equation in the matrix form:

$$\mathbf{p}[n+1] = \mathbf{p}[n]\mathbf{T}, \quad T_{ij} = \Pr(\mathbf{X}[n+1] = x_j \mid \mathbf{X}[n] = x_i). \quad (3)$$

\mathbf{T} is the (stochastic) state-transition matrix; each row must sum to one. The n -step dynamics are revealed by the Chapman-Kolmogorov equation,

$$\mathbf{p}[n] = \mathbf{p}[0]\mathbf{T}^n.$$

We obtain the transition matrix numerically by integrating the governing differential equation forward from each mesh point, using barycentric interpolation [13] to represent the transition probabilities.

For walking, we will designate one special state, x_1 , as an absorbing state representing all configurations in which the robot has fallen down. Transitions to this state can come *from* many regions of the state space, but there are no transitions *away* from this state. Assuming that it is possible to get to this absorbing state (possibly in multiple steps) from any state, then this absorbing Markov chain will have a unique stationary distribution, with the entire probability mass in the absorbing state.

The dynamics of convergence to the absorbing state can be investigated using an eigenmode analysis [1]. Without loss of generality, let us order the eigenvalues, λ_i , in order of decreasing magnitude, and label the corresponding (left)

¹Including backward steps is straightforward, but requires the model to include spatio-temporal correlations in the slope angle

eigenvectors, \mathbf{v}_i , and characteristic times, $\tau_i = \frac{-1}{\log(\lambda_i)}$. The transition matrix from an absorbing Markov chain will have $\lambda_1 = 1$, with \mathbf{v}_1 representing the stationary distribution on the absorbing state. The magnitude of the remaining eigenvalues ($0 \leq |\lambda_i| < 1, \forall i > 1$) describe the transient dynamics and convergence rate (or mixing time) to this stationary distribution. Transient analysis on the walking models we investigate here will reveal a general phenomenon: λ_2 is very close to 1, and $\tau_2 \gg \tau_3$. This is characteristic of metastability: initial conditions (in eigenmodes 3 and higher) are forgotten quickly, and \mathbf{v}_2 describes the long-living (metastable) neighborhood of the dynamics. In metastable systems, it is useful to define the *metastable distribution*, ϕ , as the stationary distribution conditioned on having not entered the absorbing state:

$$\phi_i = \lim_{n \rightarrow \infty} \Pr(\mathbf{X}[n] = x_i \mid \mathbf{X}[n] \neq x_1).$$

This is easily computed by zeroing the first element of \mathbf{v}_2 and normalizing the vector to sum to one.

Individual trajectories in the metastable basin are characterized by random fluctuations around the attractor, with occasional “exits”, in which the system enters a region dominated by a different attractor. For walking systems this is equivalent to noisy, random fluctuations around the nominal limit cycle, with occasional transitions to the absorbing (fallen) state. The existence of successful escape attempts suggests a natural quantification of the relative stability of metastable attractors in terms of first-passage times. The *mean* first-passage time (MFPT) to the fallen absorbing state describes the time we should expect our robot to walk before falling down.

Let us define the mean first-passage time vector, \mathbf{m} , where m_i is the expected time to transition from the state x_i into the absorbing state. Fortunately, the mean first-passage time is particularly easy to compute, as it obeys the relation:

$$m_i = \begin{cases} 0 & i = 1 \\ 1 + \sum_{j>1} T_{ij}m_j & \text{otherwise} \end{cases}$$

(the expected first-passage time must be one more than the expected first-passage time after a single transition into a non-absorbing state). In matrix form, this yields the one-shot calculation:

$$\mathbf{m} = \begin{bmatrix} 0 \\ (\mathbf{I} - \hat{\mathbf{T}})^{-1} \mathbf{1} \end{bmatrix}, \quad (4)$$

where $\hat{\mathbf{T}}$ is \mathbf{T} with the first row and first column removed. \mathbf{m} quantifies the relative stability of each point in state space. One interesting characteristic of metastable systems is that the mean first-passage time around an attractor tends to be very flat; most system trajectories rapidly converge to the same metastable distribution (forgetting initial conditions) before escaping to the absorbing state. Therefore, it is also meaningful to define a system mean first-passage time, M , by computing the expected first-passage time over the entire metastable distribution,

$$M = \sum_i m_i \phi_i. \quad (5)$$

When $\tau_2 \gg \tau_3$, we have $M \approx \tau_2$, and when $\lambda_2 \approx 1$, we have

$$M \approx \tau_2 = \frac{-1}{\log(\lambda_2)} \approx \frac{1}{1 - \lambda_2}.$$

IV. NUMERICAL MODELING RESULTS

This section uses two simple, classic walking models to demonstrate use of the methodology presented in Section III and to illustrate some of the important characteristics typical for metastable walking systems more generally. The two systems presented here are the rimless wheel and the passive compass gait walker, each of which is illustrated in Figure 2.




Fig. 2. The Rimless Wheel (left) and Compass Gait Walker (right) models.

A. Rimless Wheel

The rimless wheel (RW) model consists of a set of N massless, equally-spaced spokes about a point mass. Kinetic energy is added as it rolls downhill and is lost at each impulsive impact with the ground. For the right combination of constant slope and initial conditions, a particular RW will converge to a steady limit cycle behavior, rolling forever and approaching a particular velocity at any (Poincaré) “snapshot” in its motion (e.g., when the mass is vertically above a leg at $\theta = 0$ in Fig. 2). The motions of the rimless wheel on a constant slope have been studied in depth [3, 16].

In this section, we will examine the dynamics of the RW when the slope varies stochastically at each new impact. To do this, we discretize the continuous set of velocities, using a set of 250 values of ω , from 0.01 to 2.5 (rad/s). We also include an additional absorbing failure state, which is defined here to include all cases where the wheel did not have sufficient velocity roll past its apex on a particular step. Our wheel model has $N = 8$ spokes ($\alpha = \frac{\pi}{4}$). At each ground collision, we assume that the slope between ground contact points of the previous and new stance leg is drawn from an approximately² Gaussian distribution with a mean of $\bar{\gamma} = 8^\circ$. For clarity, we will study only wheels which begin at $\theta = 0$ with some initial, downhill velocity, ω_o , and we consider a wheel to have failed on a particular step if it does to reach an apex in travel: $\theta = 0$ with $\omega > 0$. (Clockwise rotations go downhill, as depicted Fig. 2, and have positive values of ω .) Note that the dynamic evolution of angular velocity over time does not depend on the choice of a particular magnitude of the point mass, and we will use spokes of unit length, $l = 1$ meter, throughout.

²To avoid simulating pathological cases, the distribution is always truncated to remain within $\pm 10^\circ$, or roughly 6σ , of the mean.

On a constant slope of $\gamma = 8^\circ$, any wheel which starts with $\omega_o > 0$ has a deterministic evolution over time and is guaranteed to converge to a fixed point of $\omega = 1.2097$ (rad/s). The return map defining the step-to-step transitions from ω_n to ω_{n+1} is given as:

$$\omega_{n+1} = \sqrt{\cos^2 \alpha \left(\omega_n^2 + \frac{2g}{L} (1 - \cos \beta_1) \right) - \frac{2g}{L} (1 - \cos \beta_2)}$$

where $\beta_1 = \frac{\alpha}{2} + \gamma$ and $\beta_2 = \frac{\alpha}{2} - \gamma$, with $\gamma > 0$ as the downhill slope. A plot of this return function is shown in Figure 3.

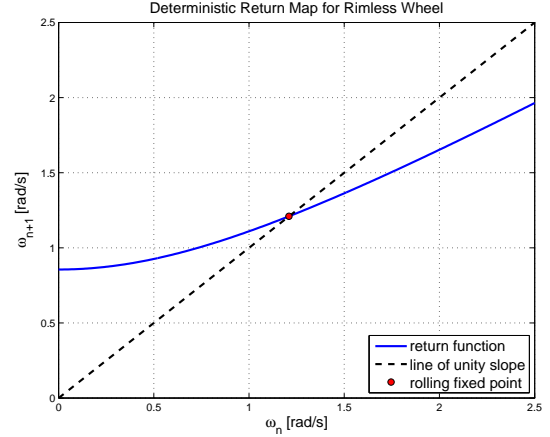


Fig. 3. Return map and fixed point for an 8-spoke rimless wheel on constant, downhill slope of 8° . Here, ω_n , is defined as angular velocity when the support spoke is exactly vertical.




Fig. 4. Return distribution and metastable “neighborhood” for an 8-spoke rimless wheel on downhill terrain with a mean step-to-step slope of 8 degrees and $\sigma = 1.5^\circ$. There is now a probability density function describing the transition from ω_n to ω_{n+1} .

When the slope between successive ground contacts is drawn from a stochastic distribution, the function given in Figure 3 is now replaced by a probabilistic description of the transitions, as illustrated in Figure 4. Given the current state is some particular ω_n , there is a corresponding probability density function (PDF) to describe what the next state, ω_{n+1} ,

will be. Figure 5 shows this set of PDF's clearly; it is a 3D plot of the same probabilistic return map shown from overhead in Figure 4. For our discretized model, each height value in Figure 5 is proportional to element T_{ij} of the transition matrix, where i is the state we are coming from (ω_n , on the x-axis) and j is the state we are going to (ω_{n+1} , on the y-axis); Figures 4 and 5 provide a graphical representation of the transition matrix describing this metastable dynamic system.




Fig. 5. 3D view of the return distribution for the stochastic rimless wheel system. This is a smoothed rendering of the step-to-step transition matrix, T , with the probability density functions for some particular states (ω_n) overlaid as lines for greater clarity.

To generate the discrete transition matrix, we calculate $\omega_{n+1} = f(\omega_n, \gamma)$ for each of a discrete set of 601 possible γ values, in the range of ± 10 degrees from the mean. Each new state is then represented in the mesh using barycentric weighting interpolation [13], which (we note) inherently adds a small level of additional (unintended) noise to the modeled dynamic. In Figures 4 and 5, the noise has a standard deviation of $\sigma = 1.5^\circ$. Using MATLAB to take the 3 largest eigenvalues of the transpose of the transition matrix for this case, we find that the largest eigenvalue, λ_1 , is within 10^{-14} of being exactly unity, which is within the mathematical accuracy expected. This eigenvalue corresponds to the absorbing failure state, and the corresponding eigenvector sums to 1, with all values except the failure state having essentially zero weight³ in this vector (since all wheels will eventually be at this state, as $t \rightarrow \infty$). All other eigenvectors sum to zero (within numerical limits), since they must die away as $t \rightarrow \infty$. The second-largest eigenvalue is $\lambda_2 = 0.999998446$. Using the methods presented in Section III, this corresponds to a system-wide MFPT of about $1/0.000001554 = 643,600$ steps. Each initial condition has a particular MFPT, $m(\omega)$, which is obtained from Eq. 4 and plotted in Figure 6. Note that the value of the mean first-passage time is nearly flat throughout a large portion of state space. This is characteristic for metastable systems, which justifies the notion of a “system-wide” MFPT, $M \approx 1/(1 - \lambda_2)$, quantifying the overall stochastic stability

³All states except the failure state had a magnitude less than 10^{-10} , numerically.

of a particular dynamic system. For this particular case, there are no regions in state space (except the failure state) with MFPT significantly lower than the system-wide value, which is *not* typical more generally; the passive compass gait walker in Section IV-B is highly sensitive to initial conditions, for instance, although it too has regions of state space which share a nearly uniform MFPT value.




Fig. 6. Mean first-passage time as a function of the initial condition, ω_0 . Data are for a rimless wheel on stochastic terrain with mean slope of 8 deg and $\sigma = 1.5^\circ$. Points show the approximation obtained through eigen-analysis of the discretized system, and a smoothed line is overlaid. Note that MFPT is largely constant over a large portion of state space.

The eigenvector associated with λ_2 yields the PDF of the metastable dynamic process – the relative probability of being in any particular location in state space, given initial conditions have been forgotten and the walker has not yet failed. Figure 7 shows the resulting probability distribution functions for the rimless wheel for each of several levels of noise. Pictorially, each system-wide PDF for a metastable system is analogous to the fixed point for a stable, deterministic system. In the deterministic case, the probability of being exactly at the fixed point approaches unity as $t \rightarrow \infty$.




Fig. 7. Quasi-stationary probability density functions for the stochastic rimless wheel for each of several values of terrain noise, σ . Each distribution is estimated by renormalizing the eigenvector associated with the second-largest eigenvalue of the transpose of the transition matrix. Note that meshing inherently adds noise to the dynamic system; smoothed lines are drawn on top of the raw data (shown as points) from the scaled eigenvectors.

The third-largest eigenvalue of transition matrix, λ_3 , quantifies the characteristic time scale in which initial conditions are forgotten, as the dynamics evolve toward the metastable

distribution (or toward failure). For the case presented here ($\sigma = 1.5^\circ$), $\lambda_3 \approx 0.50009$, which means almost half of the PDF of the initial condition composed of this eigenvector is lost (“forgotten”) with each, successive step; an even larger fraction-per-step is lost for all remaining eigenvectors (with even smaller values of λ). Within a few steps, initial conditions for any wheel beginning in our range of analysis $0 < \omega_o \leq 2.5$ have therefore predominantly evolved into the metastable PDF (or have failed). If we multiply the metastable PDF, $\phi(\omega)$, by the transition matrix, we obtain a joint probability, $Pr(\omega_n, \omega_{n+1})$, of having just transitioned from ω_n to ω_{n+1} , given the wheel has not failed by step $n + 1$. This is shown both as a 3D plot in Figure 8 and as a set of overlaid contour lines in Figure 4.




Fig. 8. 3D view of the metastable “neighborhood” of state-to-state transitions, (ω_n, ω_{n+1}) . If a rimless wheel starts from some arbitrary initial condition and has not fallen after several steps, this contour map represents the joint probability density function of being in state ω_n now and transitioning to ω_{n+1} . The contour lines drawn are identical to those overlaid in Figure 4. They correspond to the neighborhood of likely (ω_n, ω_{n+1}) pairings, analogous to the unique fixed point of the deterministic case.

This particular system has a beautiful simplicity which allows us to extract some additional insight from the conditional probability in Figure 8. Because of the definition of ω_n as being the velocity when the mass is at its apex in a given step, the value of $\omega_{n+1} = 0$ represents the boundary to the absorbing failure state in this example. If we visualize the contours of the conditional probability as they extend toward $\omega_{n+1} = 0$ in Figure 4, we see that most failures do not occur because we transition from a very slow state (ω_n close to zero) to failure but are more typically due to sudden transitions from more dominant states in the metastable distribution to failure.

Finally, when this methodology is used to analyze the rimless wheel for each of a variety of noise levels (σ), the dependence of system-wide MFPT on σ goes as shown in Figure 9. For very low levels of noise, MATLAB does not find a meaningful solution (due to numerical limits). As the level of noise increases, the MFPT decreases smoothly but precipitously. (Note that the y-axis is plotted on a logarithmic scale.) The stochastic stability of each particular system can be quantified and compared by calculating this estimate of MFPT which comes from λ_2 of the transition matrix.




Fig. 9. Mean first-passage time (MFPT) for the rimless wheel, as a function of terrain variation, σ . Estimates above 10^{14} correspond to eigenvalues on the order of $1 - 10^{-14}$ and are beyond the calculation capabilities of MATLAB.

B. Passive Compass Gait Walker

The second metastable dynamic system we analyze in this paper is a passive compass gait (CG) walker. This system consists of two, massless legs with concentrated masses at the intersection of the legs (“the hip”) and partway along each leg, and it has been studied in detail by several authors, e.g., [5, 7, 14]. Referring to Figure 2, the parameters used for our metastable passive walker are $m = 5$, $m_h = 1.5$, $a = .7$, and $b = .3$. Given an appropriate combination of initial conditions, physical parameters and constant terrain slope, this ideal model will walk downhill forever.

When each step-to-step terrain slope is instead selected from a stochastic distribution (near-Gaussian, as in Section IV-A), evolution of the dynamics becomes stochastic, too, and we can analyze the stochastic stability by creating a step-to-step transition matrix, as described in detail for the rimless wheel. The resulting system-wide MFPT as a function of terrain noise, $M(\sigma)$, is shown in Figure 10. Note that it is similar in shape to the dependence shown in Figure 9.

To analyze this system, our discretized mesh is defined using the state immediately after each leg-ground collision. The state of the walker is defined completely by the two leg angles and their velocities. On a constant slope, these four states are reduced to three states, since a particular combination of slope and inter-leg angle will exactly define the orientation of both the stance and swing leg during impact. Although the slope is varying (rather than constant) on stochastic terrain, we still use only three states to define our mesh. To do so, we simulate the deterministic dynamics (including impacts) a short distance forward or backward in time to find the robot state at the Poincaré section where the slope of the line connecting the “feet” of the legs is equivalent to our desired, nominal slope. Because the dynamics between collisions are entirely deterministic, these two states are mathematically equivalent for the stochastic analysis. If such a state does not exist for a particular collision (which occurs only very rarely),




Fig. 10. Mean first-passage time as a function of terrain variation. Results for analysis of a compass gait walker using a discretized (meshed) approximation of the transitions. Average slope is 4 degrees, with the standard deviation in slope shown on the x-axis. Noise is a truncated Gaussian distribution, limited to between 0 and 8 degrees for all cases.

we treat this as a member of the absorbing failure state. This approximation allows us to reduce the dimensionality from 4 states to 3, which improves numerical accuracy significantly. Specifically, it has allowed us to mesh finely enough to capture near-infinite MFPT for low-noise systems, while using four states did not. The three states we use in meshing are: (1) absolute angular velocity of the stance leg, X_3 , (2) relative velocity of the swing leg, X_4 , and (3) the inter-leg angle, α .

Figure 11 shows a slice of the basin of attraction for this compass gait on a constant slope (top), along with regions in state space with nearly-constant MFPT (bottom two) for two different magnitude of noise (σ) in terrain. Each slice is taken at the same inter-leg angle, $\alpha \approx 25.2^\circ$. In the deterministic case, the basin of attraction defines the set of all states with infinite first-passage time: all walkers beginning with an initial condition in this set will converge toward the fixed point with probability 1. For stochastic systems which result in metastable dynamics, there is an analogous region which defines initial conditions having MFPT very close to the system-wide value, M . Interestingly, the deterministic and stochastic basin shapes are quite similar here; we expect this may often be the case for systems such as this with discrete jumps in state space.

The image at the top of Figure 12 shows the deterministic basin of attraction for this CG walker more clearly. This plot was generated by sampling carefully over the state space and simulating the dynamics. The plot at the top of Figure 11 intentionally uses the same mesh discretization used for the stochastic system, to provide a better head-to-head comparison of the change in shape due to the addition of terrain noise (as opposed to the noise of the discretization itself). The second image in Figure 12 shows the deterministic basin of attraction for a different set of physical parameters ($m = m_h$; $a = b = .5$) on the same, constant slope of 4° . This basin looks qualitatively more delicate and the resulting performance of this walker on stochastic terrain is in fact much worse (e.g.,




Fig. 11. Basin of attraction (top) for deterministic CG walker and map of MFPT for low-noise ($\sigma = 0.5^\circ$, lower left) and high-noise ($\sigma = 1.0^\circ$, lower right) examples. To aide in visual comparison, all 3 plots use the same mesh. The “near-constant MFPT basin” for each stochastic system is essentially a low-pass filtered version of the deterministic basin of attraction, and its shape does not change significantly, even when the magnitude of the MFPT itself varies greatly (e.g., 180,000 steps [left] vs 390 [right]). This region represents a boundary on the volume in state space from which a walker is likely to be pulled into the metastable distribution.

MFPT of about 20 steps when $\sigma = 0.5^\circ$, where we find $M = 180,000$ for the other walker).

Just as in the case of the rimless wheel, the fixed point (for our deterministic compass gait system) is now replaced (in the stochastic case) by a probability density function, defining the likelihood of being in any particular state (conditioned on not having fallen) as $t \rightarrow \infty$. Figure 13 shows 2D contour plot sections of the PDF obtained from the eigen-analysis of the stochastic compass gait. The outermost contour defines a boundary containing 0.999 of the probability distribution in state space. The distribution spreads over more of state space as the level of noise increases, in a manner analogous to the widening of the distribution with noise seen in Figure 7.

Finally, we note that the relationship in state space between the PDF of the metastable dynamics, shown in Figure 13, and the region of nearly-uniform mean first-passage time, M , shown at the bottom of Figure 11, hints at where successful “escape attempts” are most likely to occur over time. Figure 14 overlays these two regions across a different dimensional slice of the 3D space for the $\sigma = .5^\circ$ and $\sigma = 1.0^\circ$ cases. As the tails of the metastable PDF (shown in yellow) approach the boundary of the uniform-MFPT basin (shown in blue), there is a higher probability of failing on any given step during the metastable process, resulting in turn in a less stochastically stable system (i.e., with a lower system-wide value of M).

V. DISCUSSION

This section briefly discusses the use of the stochastic methods presented toward designing controllers for walking systems and also provides a few further observations on the properties of metastable systems which result in multiple attractors (e.g., period- n gaits).

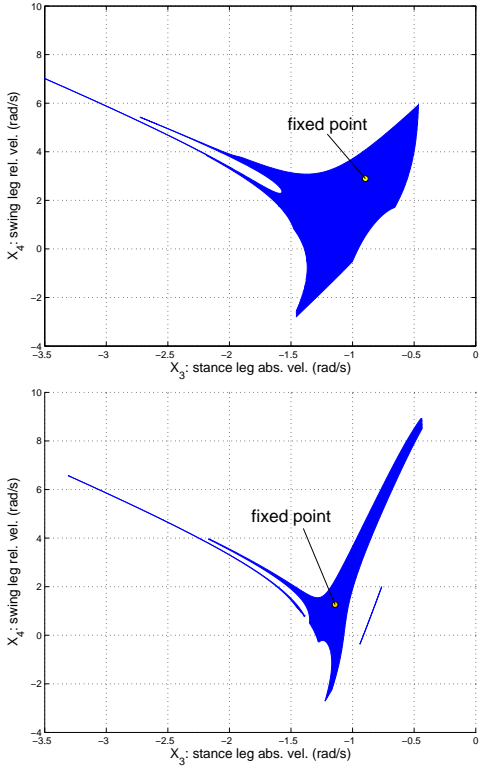


Fig. 12. Basins of attraction (blue region) and fixed point for two compass gait walkers, each on a constant slope of 4° . Walker with basin at top is more stable and uses the parameters defined for the stochastic system described throughout Section IV-B; for the other walker, $m = m_h$ and $a = b = .5$. MFPT is infinite inside the shaded region and is small (1-4 steps) outside of it. This image shows only a slice of the 3D basin, taken at the inter-leg angle of the fixed point for each respective walker. The fixed point is at $X_3 = -.89$ (rad/s), $X_4 = 2.89$ (rad/s), $\alpha = 25.2^\circ$ for the first walker, and it is at $X_3 = -1.14$ (rad/s), $X_4 = 1.26$ (rad/s), $\alpha = 33.4^\circ$ for the lower one. The deterministic basin of attraction for the second walker is narrower in shape, and this walker is significantly less stable on stochastic terrain.

A. Impacts on Control Design

One of the primary goals of a controller is to enhance the dynamic stability of a system. For walking systems, we propose throughout this paper that this should be defined as increasing the *stochastic stability*. We would like time-to-failure to be long, and we would like a system to converge toward the metastable distribution from a large set of initial conditions. The tools provided here can be used in optimizing controllers with either or both of these two aims in mind.

As an example, consider an active compass gait walker, with a torque source at the hip but with the ankles still unactuated at the ground contact. Putting this walker on a *repeating* terrain, as depicted in Figure 15, allows us to mesh across the entire state space of possible post-collision poses. By designing a low-level PD controller to regulate inter-leg angle, we can discretize the action space on a single once-per-step policy decision. The optimal high-level policy (to select desired inter-leg angle) for the system can now be solved via value iteration. Preliminary results for such a control methodology allow this underactuated compass gait model to walk continuously over

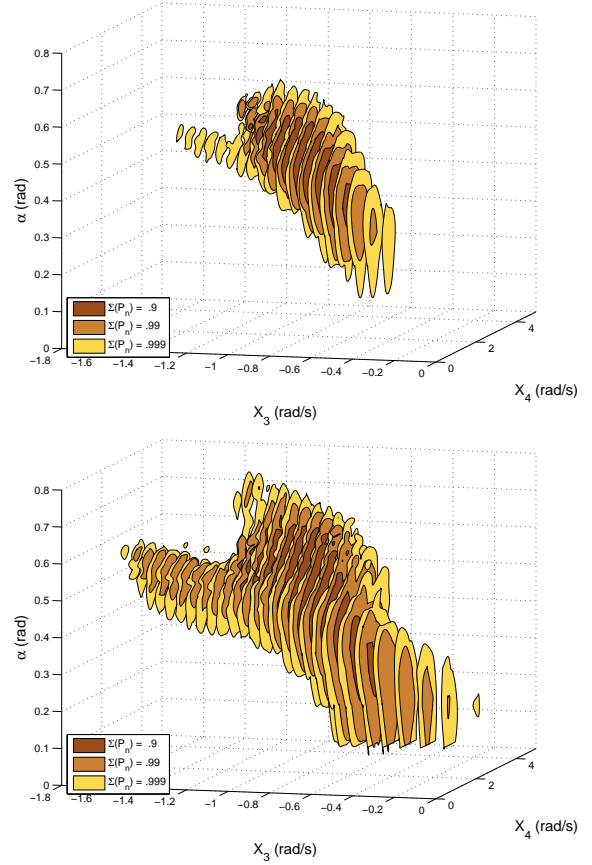


Fig. 13. On stochastic terrain, there is no fixed point for the compass gait walker. Instead, there are metastable “neighborhoods” of state space which are visited most often. As time goes to infinity, if a walker has not fallen, it will most likely be in a this region. The contours shown here are analogous to the PDF magnitude contours in Figure 7; they are drawn to enclose regions capturing 90%, 99%, and 99.9% of walkers at any snapshot during metastable walking. Top picture corresponds to $\sigma = 0.5^\circ$. Larger noise ($\sigma = 1.0^\circ$, bottom) results in larger excursions in state space, as expected.

impressively rough terrain [2].

B. Multiple stable limit cycles

Metastable dynamic systems sometimes have an inherent periodicity. We expect this may be the case on a slightly steeper slope, for instance, where compass gait models experience period-doubling bifurcations [7]. Another case where periodicity arises is for wrapping terrain, such as the terrain for the controlled walker in Figure 15. Wrapping is a realistic model for many in-laboratory walking robots, as they are often confined to walk on a boom – repeatedly covering the same terrain again and again. In our simulation of a hip-actuated CG walker on wrapping terrain, we observe that a repeating, n -step cycle results in multiple eigenvalues, λ_2 through λ_{n+1} , all with magnitude just under unity. They are complex eigenvalues, as are the corresponding eigenvectors. The top left image in Figure 15 shows such a set of eigenvalues, all lying just within the unit circle. The next-smallest set of eigenvalues are all significantly smaller in this example. The complex eigenvalues and eigenvectors mathematically capture an inherent period-




Fig. 14. Metastable system: Contours of the stochastic “basin of attraction” are shown where MFPT is $0.5M$, $0.9M$ and $0.99M$ (blue) versus contours where the integral of the PDF accounts for $.9$, $.99$, and $.999$ of the total metastable distribution (yellow). The metastable dynamics tend to keep the system well inside the “yellow” neighborhood. As the tails of this region extend out of the blue region, the system dynamics become less stochastically stable (lower M). The axis into the page represents the angle of the swing leg relative velocity, X_4 , and a slice is taken at $X_4 = 2.33\text{rad/s}$. Terrain variation for the top plot is $\sigma = 0.5$ degrees (with $M \approx 180,000$ steps). For the noisier system at bottom ($\sigma = 1.0$ degrees), M is only 20 steps or so.

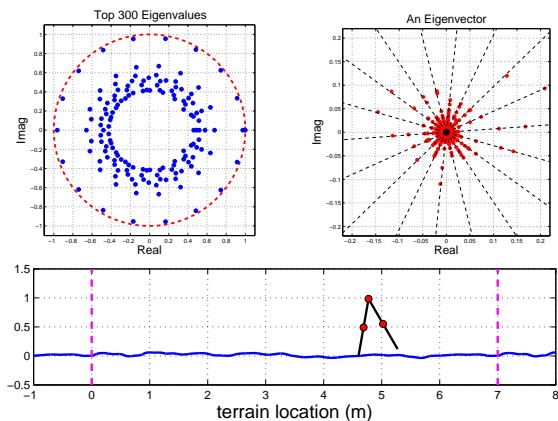


Fig. 15. Controlled compass gait walker, with torque at the hip. To solve for an optimal policy using value iteration, the terrain wraps every 7 meters. The optimization maximizes the MFPT from any given state. An eigenanalysis reveals a complex set of eigenvalues (top), spaced evenly about (but strictly inside of) the unit circle. Corresponding eigenvectors are also complex.

icity, in which the probability density function changes over time in a cyclical manner.

VI. CONCLUSIONS

The goal of this paper has been to motivate the use of stochastic analysis in studying and (ultimately) enhancing the stability of walking systems. Robots that walk are inherently more prone to the stochastic influences of their environment than traditional (e.g., factory) robots. Locomotory systems capable of interacting with the real world must deal with significant uncertainty and must perform well with limited energy budgets and despite limited control authority.

The stochastic dynamics of walking on rough terrain fit nicely into the well-developed study of metastability. The simplified models studied here elucidate the essential picture of a metastable limit cycle dynamics which makes occasional escape attempts to the fallen-down state. Metrics for stochastic stability, such as the mean first-passage time, may be potent metrics for quantifying both the relative stability across state-space and the overall system stability for real walking systems.

REFERENCES

- [1] Katie Byl and Russ Tedrake. Stability of passive dynamic walking on uneven terrain. In Art Kuo, editor, *Proceedings of Dynamic Walking 2006*, May 2006.
- [2] Katie Byl and Russ Tedrake. Approximate optimal control of the compass gait on rough terrain. In *Proceedings IEEE International Conference on Robotics and Automation (ICRA)*, 2008.
- [3] Michael J. Coleman, Anindya Chatterjee, and Andy Ruina. Motions of a rimless spoked wheel: a simple 3d system with impacts. *Dynamics and Stability of Systems*, 12(3):139–160, 1997.
- [4] Steven H. Collins, Andy Ruina, Russ Tedrake, and Martijn Wisse. Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307:1082–1085, February 18 2005.
- [5] M Garcia, A Chatterjee, A Ruina, and M Coleman. The simplest walking model: Stability, complexity, and scaling. *Journal of Biomechanical Engineering – Transactions of the ASME*, 120(2):281–288, Apr 1998.
- [6] C.W. Gardiner. *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*. Springer-Verlag, third edition, 2004.
- [7] Ambarish Goswami, Benoit Thuilot, and Bernard Espiau. Compass-like biped robot part I : Stability and bifurcation of passive gaits. Technical Report RR-2996, INRIA, October 1996.
- [8] Peter Hanggi, Peter Talkner, and Michal Borkovec. Reaction-rate theory: fifty years after kramers. *Reviews of Modern Physics*, 62(2):251–342, Apr 1990.
- [9] N.G. Van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, third edition, 2007.
- [10] Daniel E. Koditschek and Martin Buehler. Analysis of a simplified hopping robot. *International Journal of Robotics Research*, 10(6):587–605, Dec 1991.
- [11] Tad McGeer. Passive dynamic walking. *International Journal of Robotics Research*, 9(2):62–82, April 1990.
- [12] Reinhard Muller, Peter Talkner, and Peter Reimann. Rates and mean first passage times. *Physica A*, 247(247):338–356, Jun 1997.
- [13] Remi Munos and Andrew Moore. Barycentric interpolators for continuous space and time reinforcement learning. In M. S. Kearns, S. A. Solla, and D. A. Cohn, editors, *Advances in Neural Information Processing Systems*, volume 11, pages 1024–1030. NIPS, MIT Press, 1998.
- [14] Mark W. Spong and Gagandeep Bhatia. Further results on control of the compass gait biped. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 1933–1938, 2003.
- [15] P. Talkner, P. Hanggi, E. Freidkin, and D. Trautmann. Discrete dynamics and metastability: Mean first passage times and escape rates. *J. of Stat. Phys.*, 48(1/2):231–254, 1987.
- [16] Russell L Tedrake. *Applied Optimal Control for Dynamically Stable Legged Locomotion*. PhD thesis, Massachusetts Institute of Technology, 2004.