Metastable Markov Chains

Cenk Oguz Saglam and Katie Byl

***Abstract*— In this paper, we discuss the dynamics of metastable systems. Such systems exhibit interesting long-living behaviors from which they are guaranteed to inevitably escape (e.g., eventually arriving at a distinct failure or success state). At the heart of this work, we emphasize (1) that for our goals, hybrid systems can be approximated as Markov Decision Processes, (2) that although corresponding Markov chains may include a very large number of discrete states, much of their dynamic behavior is well-characterized simply by the second-largest eigenvalue, which is directly analogous to a dominant pole for a discrete-time system and describes both the mean and higher-order modes of the escape statistics, and (3) that for many systems, one can accurately describe initial conditions as being rapidly forgotten, due to a significant separation in slow and fast decay rates. We present both theory and intuitive toy examples that illustrate our approach in analyzing such systems, toward enabling and encouraging other researchers to adopt similar methods.**

I. INTRODUCTION

There are large classes of natural and human-made systems which exist in a precarious state of stability, appearing to be locally stable for long periods of time until an external disturbance perturbs the system into a region of state space with a qualitatively different local behavior. Since these systems are guaranteed to exit these locally well-behaved regions with probability one given enough time, they cannot be classified as "stable", but it is also misleading to categorize them simply as "unstable". Such systems are *metastable*, with a toy example being a ball in a hollow on terrain where a sufficient disturbance would cause the ball to roll into another local minimum, as depicted in Figure 1. Physicists have explored this phenomenon in detail and have developed a number of tools for quantifying this behavior [1], [2], [3], [4]. Metastable processes have been observed in many other branches of science and engineering including familiar systems such as crystalline structures [5], flip-flops [6], and neuroscience [7].

More recently, the tools for quantifying metastable systems has been applied to walking robots to predict how a robot will perform over variable terrain for a given control policy [8], [10]. For such analyses, the walking robot, the environment, the system noise, and the control actions can be modeled together as a Markov chain. Assuming that the initial state of the robot lies within a so-called "metastable region" of state space, the eigenvalues of the state transition matrix of the Markov chain, specifically the largest eigenvalue not associated with the (absorbing) failed system state, can be used to predict the number of steps the robot can take before failing. While this approach is described in some detail in [8], we have found when building upon the basic concepts that many important aspects of the analytic approach are not immediately obvious nor (as yet) well-documented. In this paper, we attempt to clarify the overall approach and the utility of the results, using simple toy examples.

K. Byl and C. O. Saglam are with the Electrical and Computer Eng. Dept., Univ. of California, Santa Barbara, CA 93106 USA katiebyl at ece.ucsb.edu, saglam at ece.ucsb.edu

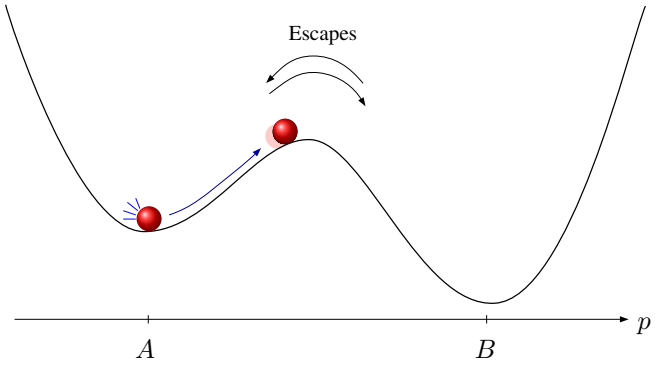

Fig. 1. Cartoon with multiple locally-stable equilibria under deterministic conditions. With sufficient noise in the model, the particle is guaranteed to transition from one local minimum to another (back and forth), although transitions may still be quite rare. From region A's perspective, "escape" corresponds to moving to region B. Figure inspired by [8] and [9].

To show the applicability of our methods to various problems we consider hybrid systems, which may exhibit either or both continuous and discontinuous dynamics. The requirement is the ability to model the full dynamics as a Markov chain. This can be done in many applications where a meaningful "step" definition can be proposed. As we justify in subsection II-B, we will model the escape of interest (e.g., moving from region A to B in Fig. 1 or falling of a walking robot) as an absorbing *halt* state.

In previous work, we have used the Mean First Passage Time (MFPT) to characterize the average number of Markov chain steps until reaching an absorbing failure state. In this paper, we present a more generalized concept – Value Before Escape (VBE) – and discuss both the mean and *variability* in a value of interest for a metastable system.

The rest of this paper is organized as follows. In Section II, we discuss the issue of approximating a hybrid system as a Markov chain, with only a finite number of discrete states. In Section III, we use the eigenvalues of the state transition matrix to calculate MFPT. Section IV discusses the relevance of employing a system-wide MFPT to characterize a metastable system and presents some toy examples that illustrate general principles. Section V introduces the VBE metric, which builds on MFPT. Sec. VI discusses future work. Finally, we conclude with a brief summary in Sec. VII.

II. OBTAINING MARKOV CHAIN

For systems in which an exact (discrete) Markov chain does not naturally exist, we can create one, using the meshing process described below.

A. Hybrid Model

Let x, γ, and ζ be the internal state, the randomness system experiences, and the control action respectively. To illustrate, for a walking robot, x is the robot's state, γ is random variable representing factors such as terrain variation or system noise, and ζ is the control action which may be a function of x and γ. We define vector $y = [x; \gamma; \zeta]$ to represent them all. Then our general hybrid model is represented as

$$\begin{aligned} \dot{y} &= f(y) \qquad y \in C \\ y^+ &= g(y) \qquad y \in D. \end{aligned} \tag{1}$$

C and D are known as flow and jump sets [11]. Note that this setting is compatible with less general cases like continuous and discrete systems with/without a control action or randomness.

B. Meshing for Markov Decision Process (MDP) Model

The first step is meshing to obtain a (discrete) Markov Decision Process (MDP) model. To do so, we need to choose a Poincaré-like section which does not necessarily decrease the dimension of the state. However, if the system has not escaped (from the region of interest) yet, it needs to keep passing through this section.

As a toy meshing example, consider Fig. 1. Our motivation is to quantify the likelihood over time that the ball will reach region B from region A. Looking from region A's perspective, escape corresponds to "next recurrence to state B", so the model sets region B as an absorbing halt state, even if it is not. In other words, the problem is constructed with the goal of finding "how long" the system will run before entering into state B (even if state B is not a "halt" state). In fact, all our metastable models assume at least one such (absorbing) state exists. Without loss of generality, we will refer to the absorbing halt state as *state 1* (x_1).

For the ball example, let p and v denote the horizontal position and velocity respectively, for which positive is to the right. The state is $x := [p; v]$. p_A is the position at point A, the mass of the ball is 1 kg, there is no control action, g_0 denotes the gravitational acceleration and the height difference between point A and the hill is $1/g_0$. Each time the ball passes through point A while not going to the left, someone will flick the ball to instantaneously change its velocity by γ, which will be normally distributed with zero mean and standard deviation σ. To begin with, let's assume the ball initially starts resting at point A. If σ is zero, then the ball will never exit the region around A. However, when $\gamma > 1.42$ the ball will escape to B with one hit. We are interested in stochastically changing γ values.

The way we approach such problems is first determining the Poincaré section as

$$S = \{x \in \mathbb{R}^2 : \ p = p_A, \ v \geq 0\}. \tag{2}$$

This section is chosen to be the sate just before the ball is flicked (before the impact). In fact, it is nothing but the jump set for our example. After defining this section we abuse the notation and refer to $x \in S$ simply by x. Then, the next state (intersecting S) is a function of the current state $x[n]$ and the randomness experienced during that step $\gamma[n]$, i.e.,

$$x[n+1] = h(x[n], \gamma[n]). \tag{3}$$

We then need to approximate this true dynamics with finitely many sets. We first choose a finite randomness set. In this example, let's consider $\gamma = \{-1, 0, 1\}$. Then, there are three states possible: Sate 1 (halt state representing escape to region B), State 2 ($v = 0$) and State 3 ($v = 1$). So, we have the system in Figure 2.

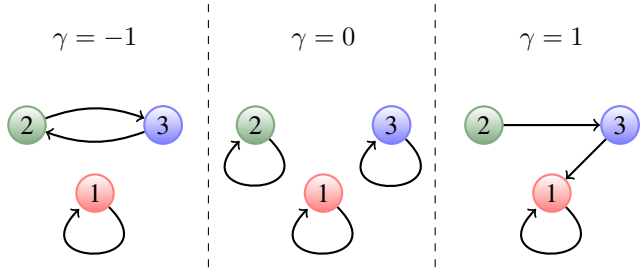

Fig. 2. Approximation of the full dynamics via Meshing. Result is a Markov Decision Process with no action and three possible randomness

Note that the approximation we had in Figure 2 will not be accurate enough due to the randomness set we chose. However, it was one of the simplest/most intuitive that could be done. Our goal was to explain the methodology with a toy example with hybrid dynamics. For accuracy, we need a much denser and (depending on the case) wider randomness range. For this toy example, we can further assume a deterministic (possibly non-linear) friction loss on the platform. This new model would be more realistic, and when $\sigma = 0$, the ball would converge to state 2 (a stable fixed point) if it started within region A with a small enough kinetic energy. In such more realistic models (which is of particular interest to the authors), there are infinitely many states to consider. Using a distance metric, we then need to appropriately choose finitely many states to represent them all. Please see [10] for an algorithm and example meshing of a 10D under-actuated bipedal walker.

C. Policy for ζ

Next, we describe the methodology for deriving the Markov chain model from a MDP. We will consider a more general case than the toy example of previous subsection. We model a Markov Decision Process as

$$x[n+1] = h(x[n], \gamma[n], \zeta[n]). \tag{4}$$

To illustrate, let's consider Figure 3, where $\gamma[n] \in \{\gamma_1, \gamma_2\}$ and $\zeta[n] \in \{\zeta_1, \zeta_2\}$, i.e., there are two available actions and two possible randomness. Also there are just 3 states, that is $x[n] \in \{x_1, x_2, x_3\}$.

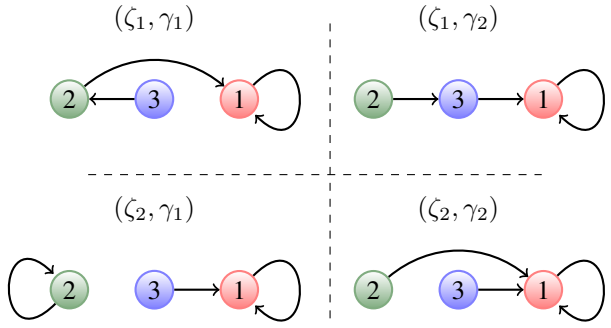

Fig. 3. Representation of a Markov Decision Process with two available actions and two possible randomness

A policy, π, is what determines which control action to take at each step. Optimal and robust policies can be obtained using dynamic programming tools [10], [12]. Let's say, we decide to use policy

$$\pi(x[n], \gamma[n]) = \begin{cases} \zeta_2 \text{ if } x[n] = x_2 \text{ and } \gamma[n] = \gamma_1 \\ \zeta_1 \text{ if } x[n] = x_2 \text{ and } \gamma[n] = \gamma_2 \\ \zeta_1 \text{ if } x[n] = x_3 \text{ and } \gamma[n] = \gamma_1 \\ \zeta_2 \text{ if } x[n] = x_3 \text{ and } \gamma[n] = \gamma_2. \end{cases} \tag{5}$$

Notice that we don't determine a control action for state x_1, since nothing can be done differently at halt state. When we use $\zeta[n] = \pi(x[n], \gamma[n])$, (4) becomes

$$x[n+1] = h(x[n], \gamma[n], \pi(x[n], \gamma[n])), \tag{6}$$

which is a function of the state and randomness only. The result is illustrated in Figure 4.

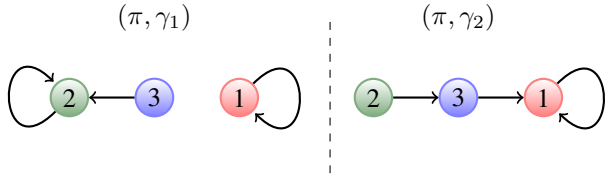

Fig. 4. MDP of Figure 3 after applying policy π defined in (5)

D. Distribution for γ

The last step before obtaining an absorbing Markov chain is to assume a distribution for randomness. For the toy example, say $P(\gamma_2) = 0.01$, i.e., with probability 0.01, $\gamma[n]$ will be γ_2. We then obtain Figure 5.

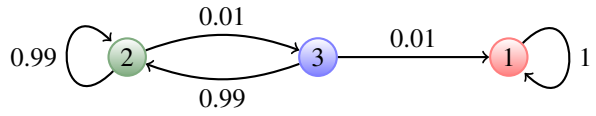

Fig. 5. MDP of Figure 3 after applying policy π defined in (5) and assuming distribution $P(\gamma_1) = 0.99$

More complicated systems will end up being surprisingly similar to Figure 5.

E. A Discrete Example

Consider tossing an unfair coin, for which the probability of having heads is 0.01. Even if the game goes forever, if we are interested in the number of flips before two heads in a row, then we should model this state (State 1) absorbing, i.e., the game ends when there are two heads in a row. Figure 5 represents this setting, where State 2 corresponds to a tails in the last flip. State 3 (tails-heads) would map to State 2 with probability 0.99 and to State 1 with probability 0.01. The game "ends" at x_1 (the halt state).

III. ABSORBING MARKOV CHAIN

The state distribution vector at step n is denoted by $p[n]$ and defined by

$$p_i[n] := Pr(x[n] = x_i). \tag{7}$$

So $p_i[n]$ corresponds to the probability of being at state x_i at step n. Since probability cannot be negative, $p[n]$ is a non-negative vector, and because the system has to be at a state at any step, $p[n]$ sums to 1. The state transition matrix, aka the Markov matrix or the stochastic matrix, has the following structure by definition:

$$T_{s\{ij\}} := Pr(x[n+1] = x_j \mid x[n] = x_i) \tag{8}$$

So, the element of T_s on the ith row and jth column gives the probability of transitioning from state x_i to state x_j. To illustrate, the Markov chain of Figure 5 is represented with

$$T_s = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0.99 & 0.01 \\ 0.01 & 0.99 & 0 \end{bmatrix} \tag{9}$$

Similar to non-negativity of $p[n]$, we have $T_{s\{ij\}} \geq 0$. And because any state will transition (possibly to the halt state or the starting state itself) after each step, each row sums to one. For the rest of the paper, we assume the number of states is $\ell > 1$. So, the state transition matrix is ℓ by ℓ. The state transition matrix gives the next state distribution, given the current one:

$$p[n+1] = T_s' \, p[n] = (T_s')^{n+1} p[0]. \tag{10}$$

Let λ be an eigenvalue of T_s. Then, there exists a non-zero vector v such that

$$T_s v = \lambda v. \tag{11}$$

As shown in [13], we next show that every eigenvalue λ of T_s satisfies $|\lambda| \leq 1$. Let k be such that $|v_j| \leq |v_k|$ for all $1 \leq j \leq \ell$. Equating the k-th components in equation (11) gives

$$\sum_j T_{s\{kj\}} v_j = \lambda v_k. \tag{12}$$

We then have

$$\begin{aligned} |\lambda v_k| = |\lambda| \, |v_k| &= \left| \sum_j T_{s\{kj\}} v_j \right| \\ &\leq \sum_j T_{s\{kj\}} |v_j| \leq \sum_j T_{s\{kj\}} |v_k| = |v_k|, \end{aligned} \tag{13}$$

where we used $T_{s\{kj\}} \geq 0$ and $\sum_j T_{s\{kj\}} = 1$. $|\lambda||v_k| \leq |v_k|$ implies $|\lambda| \leq 1$.

For the rest of the paper, we will prefer working with the transpose of T_s to make the following sections easier to follow. The transpose operation will be represented with a prime ($'$) symbol. Since T_s is square, T_s' has the same eigenvalues as T_s. Due to the nature of transpose operation and the structure of T_s, each column of T_s' sums to one.

Remember that x_1 is an absorbing state, which represents the end of game, no matter how the system escaped (e.g., the robot slipped or hit the wall). Then, T_s' has the following form:

$$T_s' = \begin{bmatrix} 1_{1\times 1} & T_1 \\ 0 & \hat{T} \end{bmatrix}_{\ell \times \ell} \tag{14}$$

Note that $\lambda = 1$ and $v = [1\ 0\ ...\ 0]'$ satisfies the equation

$$T_s' v = \lambda v. \tag{15}$$

To distinguish (possibly non-distinct) eigenvalues, we will note them by λ_j, where $1 \leq j \leq \ell$. Without loss of generality, we will let $\lambda_1 = 1$ and the associated basis vector be $v_1 = [1\ 0\ ...\ 0]'$.

Existence of Jordan normal form for any square matrix is fundamental to Linear Algebra. Consider a Jordan normal form of $\hat{T}$ given by

$$\hat{T} = \hat{V}\hat{J}\hat{V}^{-1}, \tag{16}$$

Then, as we will verify, a Jordan normal form of T_s' is given by

$$T_s' = VJV^{-1}, \tag{17}$$

$$\text{where } J = \begin{bmatrix} 1 & 0 \\ 0 & \hat{J} \end{bmatrix}, \tag{18}$$

$$\text{and } V = \begin{bmatrix} 1 & -[1\ ...\ 1]\hat{V} \\ 0 & \hat{V} \end{bmatrix} \tag{19}$$

Note that sum of each column of V equals zero, except the first one. Furthermore, these columns form a basis in $\mathbb{R}^\ell$. Equation (17) can be verified as follows. The inverse of V is

$$V^{-1} = \begin{bmatrix} 1 & [1\ ...\ 1] \\ 0 & \hat{V}^{-1} \end{bmatrix}. \tag{20}$$

Then, the right hand side of (17) can be calculated as

$$\begin{bmatrix} 1 & [1\ ...\ 1] - [1\ ...\ 1]\hat{V}\hat{J}\hat{V}^{-1} \\ 0 & \hat{V}\hat{J}\hat{V}^{-1} \end{bmatrix} = \begin{bmatrix} 1 & [1\ ...\ 1](I - \hat{T}) \\ 0 & \hat{T} \end{bmatrix} \tag{21}$$

Equation (17) is thus verified because $T_1 + [1\ ...\ 1]\hat{T} = [1\ ...\ 1]$ (columns of T_s' sum to one).

The spectrum of $\hat{T}$, denoted by $\sigma(\hat{T})$, is the set of distinct eigenvalues of $\hat{T}$. The spectral radius of $\hat{T}$ is given by

$$\rho(\hat{T}) = \max_{\lambda \in \sigma(\hat{T})} |\lambda|. \tag{22}$$

Let the spectral radius be $r = \rho(\hat{T})$. Remembering that $\hat{T}$ is a non-negative square matrix, we present the following two facts, which are proven in [14].

1) $r \in \sigma(\hat{T})$ (i.e., r is an eigenvalue of $\hat{T}$).

2) $Az = rz$ for some $z \in \mathcal{N} = \{v |\ v \geq 0 \text{ with } v \neq 0\}$

We will let $\lambda_2 := r$ and v_2 will refer to the associated column in V. Then, $v_2 = [-\|z\|_1\ z']'$. Now, let us consider the following state distribution

$$\phi := \begin{bmatrix} 0 & \frac{z'}{\|z\|_1} \end{bmatrix}' = v_1 + \frac{1}{\|z\|_1} v_2 \tag{23}$$

We will call ϕ the *metastable distribution*. Note that it is a valid initial state distribution since it sums to one and each element is non-negative. Back to our toy example, we have

$$\lambda_2 = 0.9999 \text{ and } \phi = \begin{bmatrix} 0 \\ 0.9901 \\ 0.0099 \end{bmatrix} \tag{24}$$

This metastable distribution represents a simple probability distribution: The state is x_2 with probability 0.9901, and it is x_3 with probability 0.0099. The first element of ϕ, i.e., the probability of being at x_1, is zero, by definition (of not yet having escaped).

Taking a step when ϕ is the initial condition, we obtain

$$T_s'\phi = T_s'\left(v_1 + \frac{1}{\|z\|_1} v_2\right) = v_1 + \frac{\lambda_2}{\|z\|_1} v_2 \tag{25}$$

Naturally, the resulting distribution is also non-negative and sums to one. In addition, the first element is

$$1 + \frac{\lambda_2}{\|z\|_1}(-\|z\|_1) = 1 - \lambda_2 \tag{26}$$

This means the system escaped with probability $1 - \lambda_2$. Furthermore, given the system did not escape, we also have ϕ as the final probability distribution for the states. This can be seen by zeroing the first element of $T_s'\phi$ and scaling to sum to one. Then we have the following result: When ϕ is the initial state distribution, the probability of escaping is $1-\lambda_2$, the probability of staying in the same distribution (ϕ) is λ_2. In our toy example, these can be seen by calculating

$$T_s'\phi = \begin{bmatrix} 0.0001 \\ 0.9900 \\ 0.0099 \end{bmatrix} \text{ and } \frac{T_s'\phi - (1-\lambda_2)\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}}{\left\| T_s'\phi - (1-\lambda_2)\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \right\|_1} = \phi. \tag{27}$$

Note that $\lambda_2 = 0.9999$ is close to one. This will be the case (actually λ_2 will be much higher in many cases) when we are interested in absorbing Markov chains with "rare escapes".

We then calculate the average number of steps before escape, given the initial condition was ϕ, i.e. $p[0] = \phi$. This corresponds to the term Mean First Passage Time (MFPT) in [8]. The higher MFPT is, the more stable a system is said to be. There are two cases depending on the probability of taking a step: If $\lambda_2 = 1$, then the probability of escape is zero. In this case the system will take infinitely many steps without escaping to the halt state. The other case ($\lambda_2 < 1$) is relatively more complicated and interesting as explained

next. Note that from this point on we'll focus on the case $\lambda_2 < 1$.

Given the probability of taking a step without escaping is $\lambda_2 < 1$, the probability of taking n steps only, equivalently escaping at the nth step is simply

$$Pr(x[n] = x_1,\ x[n-1] \neq x_1) = \lambda_2^{n-1}(1-\lambda_2). \quad (28)$$

Realize that as $n \to \infty$, the right hand side goes to zero, i.e., the system will eventually escape. Note that we also count the step which ended up escaping as a step. This can be verified considering escaping at the first step (taking 1 step only). When $n = 1$ is substituted, we get $1-\lambda_2$ as expected. Then, the average number of steps can be then calculated as

$$\begin{aligned} MFPT &= E[FPT] \\ &= \sum_{n=1}^{\infty} n\ Pr(x[n] = x_1,\ x[n-1] \neq x_1) \\ &= \sum_{n=1}^{\infty} n\lambda_2^{n-1}(1-\lambda_2) = \frac{1}{1-\lambda_2}, \end{aligned} \quad (29)$$

where we used the fact that $\lambda_2 < 1$. As a result, MFPT can then be calculated using

$$M = \begin{cases} \infty & \lambda_2 = 1 \\ \frac{1}{1-\lambda_2} & \lambda_2 < 1 \end{cases} \quad (30)$$

The MFPT of our toy example is

$$\frac{1}{1-\lambda_2} = \frac{1}{1-0.9999} = 10^4 \quad (31)$$

So, in the toy example if we start with the initial state distribution ϕ given in (24), then the system will take 10^4 steps on average.

The standard deviation of FPT can be calculated by

$$\begin{aligned} E[FPT^2] &= \sum_{n=1}^{\infty} n^2\ Pr(x[n] = x_1,\ x[n-1] \neq x_1) \\ &= \sum_{n=1}^{\infty} n^2\lambda_2^{n-1}(1-\lambda_2) = \frac{1+\lambda_2}{(1-\lambda_2)^2} \end{aligned}$$

$$\implies \sqrt{E[FPT^2] - (E[FPT])^2} = M\sqrt{\lambda_2} \quad (32)$$

This corresponds to $M\sqrt{\lambda_2} = 10^4\sqrt{0.9999} = 9.9995 \times 10^3$ for our toy example. λ_2 being close to one results with a standard deviation close to mean.

We are also interested in obtaining the MFPT vector, m, which gives the MFPT for each state.

$$m_i := \begin{cases} 0 & i = 1 \\ 1 + \sum_j T_{s\{ij\}} m_j & otherwise \end{cases} \quad (33)$$

The equation above says it will take zero steps to go to the halt state if the robot escaped already. Otherwise, the number of steps until halt is 1 less after a step is taken. For the toy example, this means $m_1 = 0$, $m_2 = 1 + 0m_1 + 0.99m_2 + 0.01m_3$, and $m_3 = 1 + 0.01m_1 + 0.99m_2 + 0m_3$. These equations can be solved to get

$$m = \begin{bmatrix} 0 \\ 1.99 \times 10^4 \\ 200 \end{bmatrix} \quad (34)$$

This means if we start at state x_2, we expect to take 1.99×10^4 steps before escape.

By using (14), it is straightforward to obtain

$$m = \begin{bmatrix} 0 \\ (I - \hat{T}')^{-1}1 \end{bmatrix} \quad (35)$$

To be able to calculate (35), we need $(I-\hat{T}')$ to be invertible. This is equivalent to having $\lambda_2 < 1$, which is the hidden assumption we made while defining the MFPT vector. The system-wide MFPT calculated in (29) can be also obtained by

$$M = m'\phi = \frac{1}{\|z\|_1}[1\ ...\ 1](I - \hat{T})^{-1}z \quad (36)$$

This makes sense because each state has its own MFPT, and MFPT of the metastable distribution is just linear combination of each state's MFPT weighted according to ϕ. We show the equivalence next.

$$\begin{aligned} \hat{M} &= \frac{1}{\|z\|_1}[1\ ...\ 1](I - \hat{T})^{-1}z \\ &= \frac{1}{\|z\|_1}[1\ ...\ 1](I - \hat{T})^{-1}(I - \hat{T} + \hat{T})z \\ &= \frac{1}{\|z\|_1}[1\ ...\ 1](I + (I - \hat{T})^{-1}\hat{T})z \\ &= \frac{1}{\|z\|_1}[1\ ...\ 1]z + \frac{1}{\|z\|_1}[1\ ...\ 1](I - \hat{T})^{-1}\hat{T}z \\ &= 1 + \lambda_2\frac{1}{\|z\|_1}[1\ ...\ 1](I - \hat{T})^{-1}z \\ &= 1 + \lambda_2\hat{M} \quad \implies \hat{M} = 1/(1-\lambda_2) = M \end{aligned} \quad (37)$$

Note that M is upper bounded by the largest element in m. In fact, any initial state distribution, $p[0]$, will have an MFPT that is a convex combination of the m_i values.

IV. Mean First Passage Time (MFPT)

We would like to answer why we should be using MFPT of metastable distribution. First of all, it is a lower bound for average steps taken from at least one of the states, because it is a convex combination of m_is. So, there are state(s) at least as stable. Secondly, it is a good measure of overall stability. Often systems quickly converge to their metastable distributions, where MFPT becomes the true value. Thirdly, system-wide MFPT also has advantages over calculating the MFPT vector. In case T_s is very large, finding the second largest eigenvalue is relatively very easy, whereas finding the inverse to calculate MFPT vector costs more time. Also, a scalar representing the stability is much more easier to understand than a possibly huge vector.

Remember (10). The following is taken from [15]. Without loss of generality assume J has the following form

$$J = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & J_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & J_k \end{bmatrix}_{\ell \times \ell} \tag{38}$$

where

$$J_i = \begin{bmatrix} \lambda_i & 1 & & \\ & \ddots & \ddots & \\ & & \lambda_i & 1 \\ & & & \lambda_i \end{bmatrix}_{\ell_i \times \ell_i} \tag{39}$$

with $1 > \lambda_2 \geq |\lambda_3|$.. (eigenvalues ordered descending). Then,

$$(T_s')^n = V^{-1} J^n V = V^{-1} \begin{bmatrix} 1 & 0 & \dots & 0 \\ 0 & J_2^n & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & J_k^n \end{bmatrix} V, \tag{40}$$

where

$$J_i^n = \begin{bmatrix} \lambda_i^n & n\lambda_i^{n-1} & & \frac{n!\lambda_i^{n-\ell_i-1}}{(n-\ell_i-1)!(n_i-1)!} \\ & \ddots & \ddots & \\ & & \lambda_i^n & n\lambda_i^{n-1} \\ & & & \lambda_i^n \end{bmatrix}_{\ell_i \times \ell_i}, \tag{41}$$

and ℓ_i is the size of i^{th} Jordan block. Note that as a special case of this equation we have $J_1 = 1$. While writing equation above we presume that any entry with negative factorials is zero. This equation can be verified by induction on n [15].

If the system is metastable, then we have λ_2 close to 1, otherwise it wouldn't take many steps before halting. All terms except J_1 die as n increases, but in case λ_2 is close to 1, J_2 dies slowly. If in addition $\ell_2 = 1$ and λ_3 is small, all J_is except the first two vanish quickly, which has been the case for walking robots we studied [10]. We will illustrate some cases where this is not the case.

When (T_s) is diagonalizable, it is easier to see the same facts. Let the initial distribution be

$$p[0] = c_1 v_1 + c_2 v_2 + ...c_\ell v_\ell. \tag{42}$$

Note that $c_1 = 1$ to have $\|p[0]\|_1 = 1$. Then,

$$p[n] = (T_s')^n p[0] = v_1 + c_2 \lambda_2^n v_2 + ... + c_\ell \lambda_\ell^n v_\ell. \tag{43}$$

In the light of this section we see that the metastable distribution is also given by

$$\phi_i = \lim_{n \to \infty} Pr(X[n] = x_i \ | X[n] \neq x_1]), \tag{44}$$

when the limit exists.

Next, we look at some possible cases with toy examples.

Example 1 - Distinct Eigenvalues

Our toy example had $\lambda_2 > |\lambda_3|$. In particular, it had $J_2 = \lambda_2 = 0.9999$, $J_3 = \lambda_3 = -0.0099$. This is the perhaps the easiest, yet most common case: Distinct eigenvalues, which implies diagonalizability, i.e., $J_i = \lambda_i$.

Example 2 - size 2 Jordan block

$$T_s' = \begin{bmatrix} 1 & 0.1 & 0.05 \\ 0 & 0.9 & 0.05 \\ 0 & 0 & 0.9 \end{bmatrix} \implies J_2 = \begin{bmatrix} 0.9 & 1 \\ 0 & 0.9 \end{bmatrix} \tag{45}$$

$$m = \begin{bmatrix} 0 \\ 10 \\ 15 \end{bmatrix}, \ M = 10, \quad \phi = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \tag{46}$$

This is also a good example to see that MFPT is a lower bound for the most stable state.

Example 3 - $\lambda_2 = |\lambda_3|$

$$T_s' = \begin{bmatrix} 1 & 0.1 & 0.1 \\ 0 & 0 & 0.9 \\ 0 & 0.9 & 0 \end{bmatrix} \implies J_2 = -J_3 = 0.9 \tag{47}$$

$$m = \begin{bmatrix} 0 \\ 10 \\ 10 \end{bmatrix}, \ M = 10, \quad \phi = \begin{bmatrix} 0 \\ 0.5 \\ 0.5 \end{bmatrix} \tag{48}$$

Note that this example forms a cycle with two states, given robot will not fall. (x_2 maps to x_3, and vice versa)

Example 5 - $\lambda_2 = \lambda_3$ with Jordan blocks of size 1

$$T_s' = \begin{bmatrix} 1 & 0.1 & 0.1 \\ 0 & 0.9 & 0 \\ 0 & 0 & 0.9 \end{bmatrix} \implies J_2 = J_3 = 0.9 \tag{49}$$

$$m = \begin{bmatrix} 0 \\ 10 \\ 10 \end{bmatrix}, \ M = 10, \quad \phi = \begin{bmatrix} 0 \\ a \\ 1-a \end{bmatrix}, \ a \in [0, 1] \tag{50}$$

In this case, state 2 and 3 are isolated. So, they can be studied independently.

V. Value Before Escape (VBE)

Often, one wishes to include multiple objectives in a single value function, for example penalizing failure events (e.g., falling down) while also rewarding fast speed and low energy use. Here, we introduce a new term called Value Before Escape (VBE) and discuss its probability distribution for a metastable system. In this section, the "value" will simply be the number of steps before failure (see Sec. VI for future extension). A cumulative value achieved for a given trial of a metastable system will be drawn from what is essentially an exponential probability density function, so long as the system begins with a "reasonable" (non-fatal) initial condition. This is because initial conditions are rapidly forgotten, due to a separation in eigenvalues of the transition matrix, so that the dynamics before absorption are strongly dominated by a single, dominant, discrete-time pole (λ_2). Thus, for a particular "confidence level", pr, one can rapidly estimate both a lower and upper Value Before Escape, LVBE and UVBE, respectively; i.e., one would observe a value above LVBE with probability pr, and one would observe a value below UVBE with probability pr.

Let us define the probability of taking more than LVBE steps as λ_2^{LVBE}. Then, the lower bound on number of steps taken with probability pr can be calculated by

$$\text{LVBE}(pr) = log_{\lambda_2} pr. \tag{51}$$

Note that LVBE(1)=0, so we can only guarantee taking a single step, which leads to the halt state. The probability of taking less than UVBE steps is $1-\lambda_2^{\text{UVBE}-1}$. Then, the upper bound on number of steps taken with probability pr is

$$\text{UVBE}(pr) = log_{\lambda_2}(1-pr) + 1. \tag{52}$$

Note that $\lim_{pr\to 1} \text{UVBE}(pr) = \infty$, so it may take infinitely many steps before converging to the halt state (theoretically). We can then limit the VBE for a given probability

$$\text{LVBE} \leq \text{VBE} \leq \text{UVBE}. \tag{53}$$

To have LVBE<UVBE we require $pr > \lambda_2/(1+\lambda_2)$, which is approximately 0.5 for $\lambda_2 \approx 1$. To illustrate the advantage of looking to VBE, let's go back to our toy example. The probability of taking more than $M = 10^4$ steps is only 36.79 %. On the other hand, we have

$$\begin{aligned} 1.05 \times 10^3 \leq \text{VBE}(0.9) \leq 2.3 \times 10^4 \\ 100.5 \leq \text{VBE}(0.99) \leq 4.6 \times 10^4. \end{aligned} \tag{54}$$

These numbers are not coincidence. As $\lambda_2 \to 1$, the probability of taking M steps is

$$\lim_{\lambda_2 \to 1} \lambda_2^{1/(1-\lambda_2)} = \frac{1}{e} \approx 0.3679. \tag{55}$$

In fact, 0.3679 is an upper limit for the probability of taking $1/(1-\lambda_2)$ steps for any λ_2. In addition, when λ_2 and pr are close to one, we have

$$\text{LVBE}(pr) = \frac{ln(pr)}{ln(\lambda_2)} \approx \frac{1-pr}{1-\lambda_2} = (1-pr)M \tag{56}$$

$$\text{UVBE}(pr) = \frac{ln(1-pr)}{ln(\lambda_2)} + 1 \approx -ln(1-pr)M \tag{57}$$

Note that LVBE is of interest for walking systems, to give a conservative bound on steps to failure, while UVBE would be helpful in modeling epedemics, to get a conservative time when everyone will be healthy (i.e., recurrence to an "all-healthy" system state).

VI. FUTURE WORK

The MFPT measures a discrete number of steps. In our future work, we plan to continue work on Mean First Passage Value (MFPV), where the value may combine total time to failure, speed, energy spent, distance traveled and other factors, combined to achieve some desired trade-off between these performance metrics. This corresponds to updating (33), where "the reward" is 1 if a step is taken from a non-halt state and can be handled with dynamic programming, as suggested in [10] and applied in [16], [17].

VII. CONCLUSIONS

The goal of this work is to explain the theory and intuition underlying the use of eigenanalysis of a Markov chain representation of a metastable system toward understanding (and eventually optimizing) long-term behavior, when absolute guarantees of stability cannot be made. The paper is intended to encourage and facilitate use of these methods in analyzing a wide range of metastable, hybrid systems. We have focused here on absorbing Markov chains. Using toy examples, we show how to obtain a Markov chain from a hybrid system, to obtain a stochastic transition matrix. We show that the second largest eigenvalue of this matrix sets the convergence rate to the absorbing state, and that the expected number of steps before halt, given the initial distribution is the so-called metastable distribution, provides a useful (single value) system-wide metric. We also introduce Value Before Escape, and describe the calculation of lower and upper confidence bounds on its value.

VIII. ACKNOWLEDGMENT

This work was supported by the Institute for Collaborative Biotechnologies through grant W911NF-09-0001 from the U.S. Army Research Office. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

REFERENCES

[1] P. Talkner, P. Hanggi, E. Freidkin, and D. Troutmann, "Discrete dynamics and metastability, mean first passage times and escape rates," *Journal of Statistical Physics*, 1987.

[2] P. Hanggi, P. Talkner, and M. Borkovec, "Reation-rate theory: Fifty years after kramers," *Reviews of Modern Physics*, 1990.

[3] R. Muller, P. Talkner, and P. Reinmann, "Rates and mean first passage times," *Physics A*, 1997.

[4] N. Kampen, *Stochasitc Processes in Physics and Chemistry*. Elsevier, Amsterdam, 2007.

[5] A. E. Larsen and D. G. Grier, "Like-charge attractions in metastable colloidal crystallites," *Nature*, 1997.

[6] H. J. Veendrick, "The behavior of flip-flops used as synchronizers and prediction of their failure rate," *IEEE J. of Solid-State Circuits*, 1980.

[7] A. Fingelkurts and A. Finglekurts, "Making complexity simpler: Multivariability and metastability in the brain." *Int. J. of Neuroscience*, 2004.

[8] K. Byl and R. Tedrake, "Metastable walking machines," *Int. J. of Robotics Research*, vol. 28, pp. 1040–1064, Aug 2009.

[9] M. Benallegue and J.-P. Laumond, "Metastability for high-dimensional walking systems on stochastically rough terrain," in *Proc. Robotics: Science and Systems (RSS)*, 2013.

[10] C. O. Saglam and K. Byl, "Robust policies via meshing for metastable rough terrain walking," in *Proc. Robotics: Science and Systems (RSS)*, 2014.

[11] R. Goebel, R. Sanfelice, and A. Teel, *Hybrid Dynamical Systems: Modeling, Stability, and Robustness*. Princeton U. Press, 2012.

[12] R. Bellman, "A Markovian decision process," *Indiana University Mathematics Journal*, vol. 6, pp. 679–684, 1957.

[13] K. Matthews. (1995) Number theory web. [Online]. Available: http://www.numbertheory.org/courses/MP274/markov.pdf

[14] C. D. Meyer, Ed., *Matrix Analysis and Applied Linear Algebra*. Philadelphia, PA, USA: Soc. for Ind. and App. Mathematics, 2000.

[15] J. P. Hespanha, *Linear Systems Theory*. Princeton, New Jersey: Princeton Press, Sep. 2009.

[16] C. O. Saglam and K. Byl, "Quantifying the trade-offs between stability versus energy use for underactuated biped walking," in *Proc. of IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*, 2014.

[17] P. Agharkar, R. Patel, and F. Bullo, "Robotic surveillance and Markov chains with minimal first passage time," in *Proc. of IEEE Int. Conf. on Decision and Control (CDC)*, Dec. 2014, to appear.