The Basin of Attraction for Running Robots: Fractals, Multistep Trajectories, and the Choice of Control

Tom Cnops, Zhenyu Gan, and C. David Remy, *Member, IEEE*

***Abstract*— If the control authority of a running system is insufficient to reach a target state in a single step, i.e. if deadbeat control is not possible, then a stabilizing controller is faced with the decision on how to plan intermediate steps. In this work, we compare the performance of a simple greedy control policy (that computes deadbeat inputs and simply caps them) with the optimal performance found by an exhaustive search through decision space. The performance criterion used in this study is the basin of attraction: the set of all states from where the target state will be reached in a finite number of steps. Using the planar spring-loaded inverted pendulum (SLIP) as a model for a running robot, we compare the two control schemes and a fully passive behavior. To this end, we extended the passive slip model to include a controllable, yet limited variation of the touchdown angle and of the damping in the leg spring. We quantified the number of steps that it would take for the model to fall or converge from arbitrary initial states. The paper highlights how the passive stabilization, that is inherent to the SLIP model, greatly influences the dynamics of the controlled system. Furthermore, it reveals some new insights into the structure of basins of attractions of SLIP-like running models.**

I. Introduction

One of the staple models for analyzing legged locomotion is the Spring-Loaded Inverted Pendulum (SLIP). The SLIP model has been used successfully as a descriptive model [1] and as a control template [2], [3], [4]. As a result of this wide applicability, the problem of controlling the SLIP model has garnered widespread attention. Although it is known that a constant touch-down angle can result in stable gaits [5], [6], efforts to improve the model's behavior have included swing leg retraction to optimize the touchdown angle [7], dynamically changing the leg's spring stiffness [8] and adding a series linear actuator to the model [9], [10].

However, despite their variety, these controllers have in common that they all basically attempt deadbeat control: their goal is to eliminate the deviance from the target state after a single step. This approach works well in cases where the target state is not very far from the current state and can thus be reached in a single step. In this case the main concern is accurate tracking and tracking error is a valid performance metric. For many legged systems, constraints on the actuators such as bounded swing speed or leg power result in the existence of states from where a fall can be avoided, but from where a target state cannot be reached in a single step. As an example, imagine a sprinter who just crossed the finish line and needs many steps to slow down. Although a controller built around minimizing some distance metric to the target state does not necessarily fail if multiple steps are required, its performance might be suboptimal.

The authors are with the Robotics and Motion Laboratory (RAMlab), Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI (tcnops@umich.edu, ganzheny@umich.edu, cdremy@umich.edu)

An interesting question poses itself: how much can potentially be gained by improving the selection of intermediate steps along a multi-step trajectory? Since multi-step trajectories are inherently related to large deviations, local assessments of stability such as eigenvalues or gain margins cannot be used. Instead, the criterion used in this work to evaluate controller performance is the *basin of attraction*: the set of all states from where the controlled system will reach the target state within a finite number of steps. The basin of attraction can be subdivided in n-step controllability regions, defined as the set of all states from where the target state will be reached in n steps or less. This allows for comparison, both in terms of robustness (will a control policy prevent a fall or not) and efficiency (how many steps does it take to reach the target).

Another measure for robustness is the gait sensitivity norm [11], which can be calculated faster than the basin of attraction. The gait sensitivity norm has the disadvantage, however, that a careful consideration must be made as to which gait characteristics should be included in the analysis. It was for this reason that the basin of attraction was chosen as the metric in this work. Additionally, the basin of attraction yields a more accurate quantification of robustness, particularly when single impulse disturbances are considered, as the gait sensitivity norm is intended to evaluate periodic disturbances.

Due to the computational expense of calculating the basin of attraction, this work is constrained to what is arguably the simplest system that can exhibit running: the planar spring-loaded inverted pendulum (SLIP). To allow for active control of the motion, this otherwise passive system was extended to include a variable touchdown angle and damping ratio. A negative damping ratio is allowed in the model, to be able to inject energy into the system. This scheme was selected in favor of more elaborate schemes for reasons of simplicity and because the exact mechanism of energy control does not greatly influence our results. The control variables were bounded to a narrow range to represent physical limits in actuation magnitude. For this simplified model of a running robot, we characterized the complete basin of attraction and n-step controllability regions for a greedy control policy, similar to the ones described in literature, for example in [7], [8], [9], [10]. These results were then compared to a control policy that maximized both robustness and steps-to-target-efficiency, found by an exhaustive numerical search,

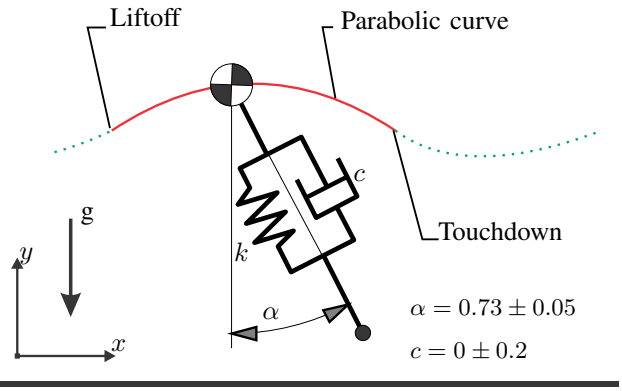

Fig. 1. The model of the spring-loaded inverted pendulum with damping that was used for the numerical work in this paper, alongside a possible trajectory of the center of mass. It is assumed that the system is running so that distinct stance and flight exist. The touchdown angle α and damping coefficient c can be chosen by the system's controller at the apex of flight, but are limited to a narrow range of possible values due to actuator limits.

explained in further detail in Section II.

II. Model

As shown in Figure 1, the SLIP model consists of a point mass atop a linear, massless spring in parallel with a purely viscous damper. The foot is modeled as a point at the bottom of the spring. We assume that whenever the foot is in contact with the ground there is no sliding with respect to the surface. The equations of motion of the SLIP model are shown below. These equations do not have a known closed-form solution, and only rough analytical approximations exist [6], hence this work relies on solving the equations numerically.

$$\ddot{\mathbf{r}}_{flight} = \begin{bmatrix} 0 \\ -g \end{bmatrix} \tag{1}$$

$$\ddot{\mathbf{r}}_{stance} = \frac{\frac{k}{m}(l_0 - l)}{l} \begin{bmatrix} x - x_s \\ y \end{bmatrix} - \frac{c}{m}\dot{\mathbf{r}} - \begin{bmatrix} 0 \\ g \end{bmatrix} \tag{2}$$

$$\textit{Touchdown condition: } \frac{y}{\cos\alpha} = l_0 \tag{3}$$

$$\textit{Liftoff condition: } l = \sqrt{(x - x_s)^2 + y^2} = l_0 \tag{4}$$

In the above equations $\mathbf{r}$ is the position of the center of mass given by the x and y coordinates, g is the acceleration due to gravity, k is the spring constant, m is the mass of the system l_0 is the rest length of the leg, l is the current length of the leg, x_s is the horizontal coordinate of the contact between the foot and the ground, c is the damping coefficient and α is angle between the leg and the vertical axis. c and α represent the control input to the system. As this work focuses on systems with limited actuation, these are constrained to a narrow interval.

The equations of motion, along with liftoff and touchdown conditions were implemented in a MATLAB framework [12]. As the leg is massless, no swing phase is modeled and we assume that the angle α is reached instantaneously after the foot left the ground. The control scheme consists of choosing the angle α and the damping coefficient c at each apex transition.

TABLE I

The parameters of the system after normalization. All parameters are normalized to total mass m, rest leg length l_0, and gravity g.

Symbol	Value	Unit	Description
System parameters			
m	1	$[m]$	Total mass
l_o	1	$[l_0]$	Rest leg length
g	1	$[g]$	Gravitational constant
k	25	$[\mathrm{m\,g/l_0}]$	Spring constant
Controlled variables			
c	0 ± 0.2	$[\mathrm{m}\sqrt{\mathrm{g/l_0}}]$	Damping coefficient
α	0.73 ± 0.05	[rad]	Touchdown angle
System states			
x	·	$[\mathrm{l_0}]$	Horizontal position
y	·	$[\mathrm{l_0}]$	Vertical position
$\dot{x}$	·	$[\sqrt{\mathrm{g\,l_0}}]$	Horizontal velocity
$\dot{y}$	·	$[\sqrt{\mathrm{g\,l_0}}]$	Vertical velocity
l	·	$[\mathrm{l_0}]$	Leg length
$time$	·	$[\sqrt{\mathrm{l_0/g}}]$	For derivatives

Step-to-step discretization and calculation method

The continuous SLIP model can effectively be made into a discrete-time system by only considering the Poincaré map formed by using apex transitions as an intersecting surface. As we assume running, the system will experience an apex transition at each step. The flight phase follows a ballistic trajectory that is uniquely defined by its apex height y_{apex} and apex horizontal velocity $\dot{x}_{apex}$, as the exact horizontal position is not relevant for our purposes and the vertical velocity $\dot{y}$ is always zero at the apex transition.

Additionally, equation (2) shows that the system dynamics during stance depend only on the control inputs and the preceding ballistic trajectory. It is thus possible to replace the SLIP model by a Poincaré mapping from one apex transition to the next.

$$\mathbf{s_{k+1}} = \mathbf{P}(\mathbf{s_k},\ \mathbf{p},\ \mathbf{u}) \tag{5}$$

Where $\mathbf{s_i}$ is a vector containing the variables describing the apex (apex height y_{apex} and apex horizontal velocity $\dot{x}_{apex}$) after i steps, $\mathbf{p}$ contains the system parameters and $\mathbf{u}$ contains the control inputs, i.e. the selected touchdown angle α and damping coefficient c. The function $\mathbf{P}$ is the Poincaré map and is approximated by numerically solving equations (1) and (2) for a fine set of discrete values in the relevant ranges of $\mathbf{s}$ and $\mathbf{u}$ and interpolating. Once this is done, we can find all states $\mathbf{s^1}$ for which a control input $\mathbf{u}$ exists so that $\mathbf{P}(\mathbf{s^1},\ \mathbf{p},\ \mathbf{u})$ is the target state $\mathbf{s^0}$. We say the states $\mathbf{s^1}$ can reach the target in a single step. We can then look for all states $\mathbf{s^2}$ that cannot reach the target, but that can reach $\mathbf{s^1}$ within a single step. The states $\mathbf{s^2}$ can reach the target in two steps. After n-iterations we have found all states that can reach the target in n steps. The iteration ends if no more new states are discovered: all remaining states cannot reach any state from where the target can be reached.

III. Passive Dynamics

As the subject of this work is a constrained actuated system, it is useful to imagine how the system would behave in the limit of the largest possible constraints, as this could give us some insight in the underlying dynamics of the system. The limit of constraint for the actuated SLIP is the passive SLIP. In this case, the touchdown angle α is constrained to a single possible value and the damping coefficient c is set to zero. The passive SLIP is energy conservative: at each apex transition height and horizontal velocity must satisfy the following constraint:

$$E_{tot} = m\,g\,y_{apex} + \frac{1}{2}\,m\,\dot{x}_{apex}^2 = constant \qquad (6)$$

Although the existence of a basin of attraction in the passive SLIP has been demonstrated before for example in [13] and [5], we believe three lesser known aspects of the basin of attraction deserve to be highlighted. The basin of attraction for the passive SLIP is shown in Figure 2 for several touchdown angles, along with an indication of how many steps it takes to approach the attractor or to fall. It is immediately clear that, depending on the touchdown angle, the basins of attraction exhibit different characteristics. In common is the approximate location of the attracting periodic cycle in the basin and that at higher speeds, the attractor loses its stability and the basin of attraction disappears, although a very long 'tail' remains, formed by unstable periodic cycles. Both the stable and unstable periodic cycles are shown in green in the figure. In the case of low α values (i.e. the SLIP is more upright), long 'filaments' can be seen reaching upwards. These are not part of the basin of attraction, as they indicate states from where trajectories will go near the unstable periodic trajectory in the 'tail' of the basin of attraction and thus will take many steps before falling. Except for $\alpha = 1.08$, the upper-left of the basin of attraction is formed by an unstable periodic cycle and the stable periodic attractor lies below.

For sufficiently high α, the basin of attraction becomes convex. In the intermediate case however, for example for $\alpha = 0.73$, a gap can be seen in the basin of attraction. This has been demonstrated before in [5]. However, what we believe escaped previous notice, is that more gaps can be found near the border formed by the unstable periodic cycles. This results in a fractal structure for the basin of attraction. Although a fractal basin of attraction has been known for a walking biped, this fractal has a different structure [14]. The presence of this fractal structure implies that no matter how fine state space is sampled, the grid will always be too coarse to capture all features of the basin of attraction. Note that this fractal behavior does not present itself in any active version of the SLIP model, as adding a controller capable of changing either the energy level or touchdown angle of the runner will result in the destruction of arbitrarily fine features of the basin of attraction.

We also wish to draw attention to the different dynamic behavior throughout the basin of attraction. In Figure 3, three trajectories are shown throughout the basin of attraction. In

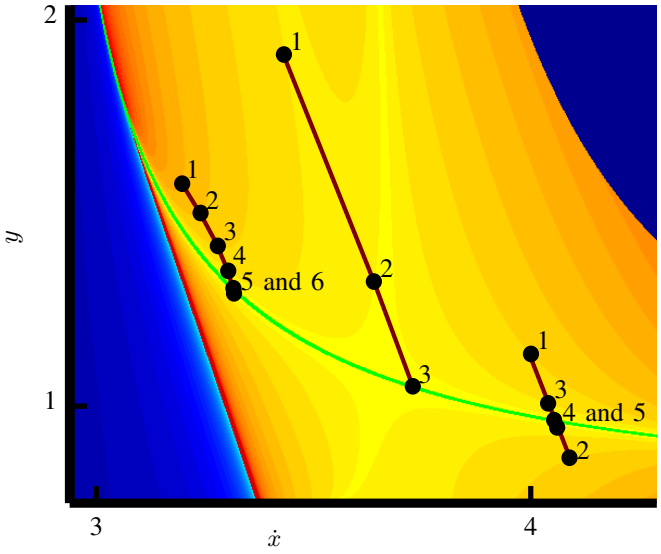

Fig. 3. Detail of Figure 2 for $\alpha = 0.73$. The numbered dots represent successive steps, while the three underlying curves indicate three energy levels.

the left-hand side of the region we see a slow progression towards stable, periodic cycle. It takes many steps to change the system's state by a relatively small amount. In the right-hand side of the figure we see the opposite. Each step considerably changes the system's state, resulting in oscillations around the periodic cycle. And centrally in the figure, the system progressed rapidly towards the periodic solution, without any oscillation. From a controller's perspective, the behavior in the left side of the figure is preferable if a certain unstable state should be tracked. As the system's state varies little on a step-to-step basis, the control inputs can cause relatively large changes in the system's state, or conversely, small control inputs would suffice to track the target. This is opposed to the behavior in the left side of the figure, where the controller would have to work against the dynamics. A considerable fraction of the controller's authority would be required to counter the system dynamics. However, if a stable state should be tracked, the behavior centrally in the figure is the most desirable, as disturbances would die out in the least possible number of steps without any control input.

As a side note, for $\alpha = 0.90$ or 1.08, the oscillations around the periodic cycle can actually be periodic themselves, i.e. a two-step, periodic cycle arises. This cycle is unstable however and exists only on the boundary of the basin of attraction.

It must also be kept in mind that for very small or very large values of α, no stable periodic cycles exist. Unstable periodic cycles continue to exist for almost all values of α, and thus for some states the system can take quite some steps before falling, similar to the structure formed by the 'filaments' described above.

The full basin of attraction for a fixed touchdown angle of about $\alpha_0 = 0.73$ is shown in Figure 2. This angle was chosen to highlight all features of the basin of attraction which do not all appear for arbitrary combinations of spring stiffness and touchdown angle. Although this figure is two-dimensional, in the absence of disturbances, the system

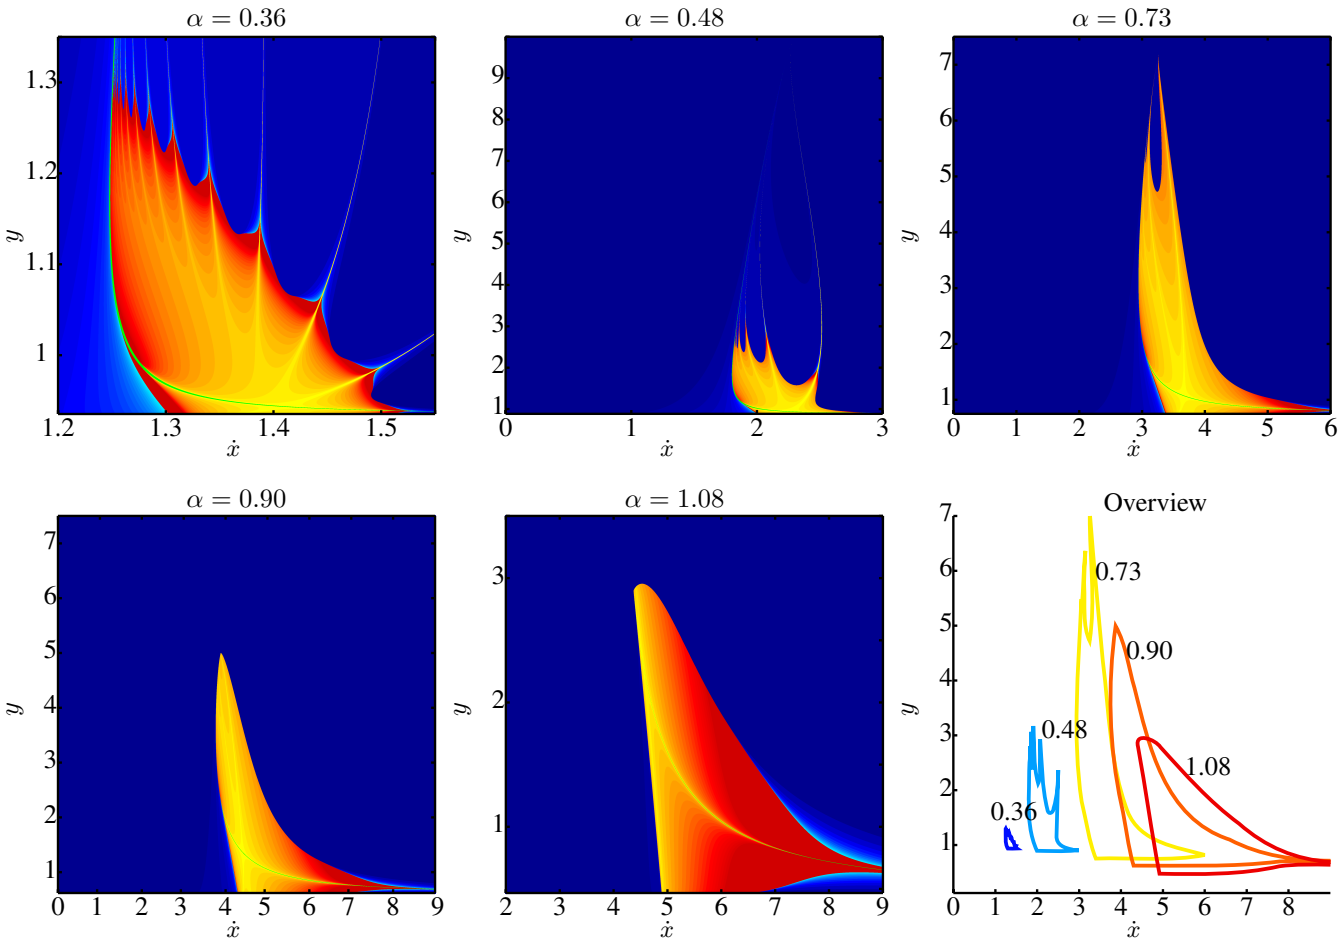

Fig. 2. Each point on this graph represents a ballistic trajectory, defined by its horizontal velocity (x-axis) and apex height (y-axis). The colors indicate the evolution of the system in time. Blue indicates a fall, with darker blue falling after fewer steps. Yellow and red indicate that the system will evolve to a stable periodic cycle. Redder tints mean that it will take more steps to approach the periodic cycle. The green lines connect all stable and unstable periodic cycles.

can only move along one-dimensional parabolas of constant energy, as indicated by equation (6).

IV. Active Stabilization

Introducing damping allows the system to change energy levels. In principle, if the touchdown angle and the damping coefficient were completely unbounded, the system would have two independent control inputs in a two-dimensional state space and be perfectly controllable. In a realistic system there are practical limits to actuators. In order to represent this, we limit both the touchdown angle and damping coefficient. The range for c was chosen to be -0.2 to 0.2. This means that the system can gain or loose roughly ten percent of its energy in a single step (although this of course greatly depends on the exact hopping speed and height). We intentionally choose relatively small control inputs compared to some real systems in order to increase the number of steps in the tracking trajectories and to magnify the difference between the different control schemes. The range for α was chosen to be 0.73 ± 0.05. The central value for α was chosen because the corresponding basin of attraction shows characteristics of both small and large touchdown angles, as shown in Figure 2, while the range was limited so that a runner would need to take several steps before coming to a complete stop.

A. Optimal Control

A mentioned above, a trivial strategy was used to find optimal control inputs. By simulating every possible input at every possible system state (after discretization), trajectories are found that can reach the target in the least number of steps for all system states. Contour lines of the number of steps required to reach the target are shown in Figure 4. Note that the discretization error accumulates for each step, and thus the contours become increasingly inaccurate. However, the error-estimation for 12-step trajectories (twelve times the discretization spacing) is less than the thickness of the lines in the figure. The target state was selected to be $(\dot{x}_{apex}, y_{apex}) = (3, 2.5)$, which corresponds to the periodic cycle with the lowest energy for the case $\alpha = 0.73$.

It can be observed that the sequence of steps is sometimes complex in the sense that there appears to be no clear way of determining the necessary choices in advance, which means that a realistic controller will most likely choose suboptimal

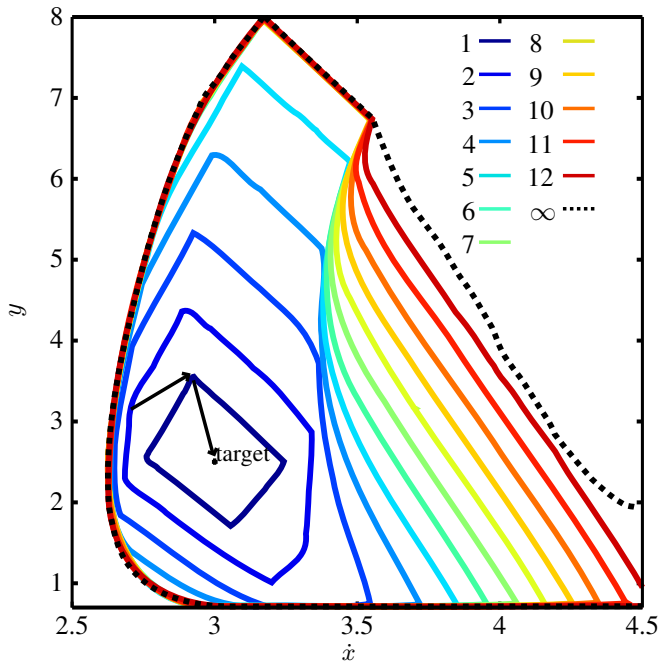

Fig. 4. The performance of the optimal control scheme. The black dot indicates the target state. The solid lines are the first twelve contours of the number of steps required to reach the target. The dotted line envelops the region from where the target can be reached in a finite number of steps. The two black arrows indicate a trajectory towards the target state.

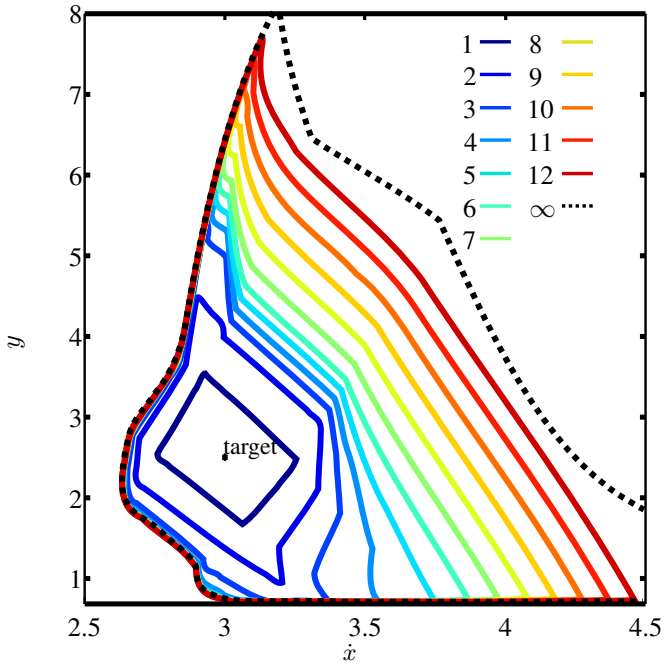

Fig. 5. The performance of the greedy control scheme. The black dot indicates the target state. The solid lines are the first twelve contours of the number of steps required to reach the target. The dotted line envelops the region from where the target can be reached in a finite number of steps.

paths. For example to reach the target from the initial state $(\dot{x}_{apex}, y_{apex}) = (2.73, 3.71)$ two steps are required, both use the smallest allowed touchdown angle, but the first step requires $c = -0.2$, while the second requires $c = 0.2$. These two steps are indicated on Figure 4 with the black arrows.

B. Greedy method

A much more feasible approach than exhaustively trying all controller inputs to find an optimal trajectory is based on deadbeat control. In most cases, deadbeat control would be possible if there were no actuator limits and the control inputs required for deadbeat control can be approximated reasonably well. A greedy controller could then simply find those control inputs that respect the actuator limits and are closest to the hypothetical deadbeat control inputs.

The number of steps required to reach the fixed point with this scheme is shown in Figure 5. It is interesting to note that according to our calculations all initial conditions either result in a fall or reach the target state in a finite number of steps. However, it is possible that some periodic trajectories exist on the boundary of the basin of attraction, but that these are unstable, so that in practice these trajectories will also converge to the target or result in a fall.

When comparing the greedy controller in Figure 5 with an exhaustive search of all solutions in Figure 4 we do notice that the greedy control scheme clearly misses some opportunities to avoid a fall. Compared to the optimal controller, its basin of attraction is 84.3% as large, and compared to the passive case its basin of attraction is 195.4% as large.

Lastly, let us look at an overlay of the basins of attraction. In Figure 6 the basins of attraction for the greedy control scheme, the optimal control scheme and the passive system

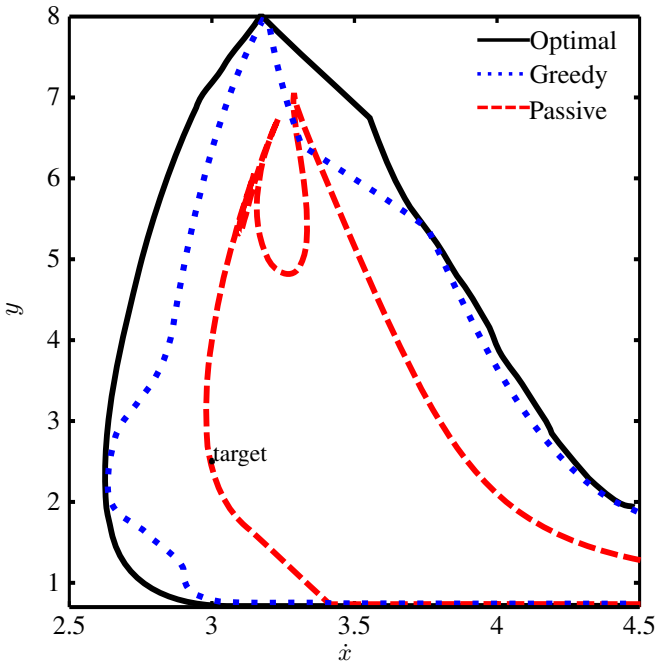

Fig. 6. A comparison of the basin of attraction for the optimal and greedy controllers and for the passive case.

are shown. At some points, the boundary for the greedy case touches the boundary of the optimal case. This means that in these instances the greedy controller has optimal performance. On the other hand, in a select few cases, the greedy controller also exhibits worse performance than the passive SLIP, although its basin of attraction is almost twice as large. The basins of attraction for the greedy and optimal controller have roughly the same shape as that for the passive

SLIP, even if they are much larger. This underlines the importance of the passive dynamics of the system and means that the distance between a point and the edge of the basin of attraction could be used as a rough metric of the control authority needed to stabilize that point.

V. Conclusion

Despite being a very simple system, the passive SLIP has a rich structure in its basin of attraction. It can contain fractal structures, strong nonlinear dynamics and two-step cycles. Additionally, the unstable periodic cycles can greatly influence how many steps it takes until the system falls. Trajectories that pass close to these cycles can potentially take many steps before falling. The importance of these passive dynamics for the actuated system cannot be overstated. For example, the long 'tail' in the basin of attraction for the actuated system can directly be linked to the unstable periodic cycles in these region. Although actuation does, of course, improve robustness, the shape of the passive basin of attraction is recognizable in the actuated systems. A state near the basin of attraction of the passive system will thus likely require less actuation to avoid a fall.

The relative robustness of the greedy controller is also a remarkable result. It succeeds in reaching the target state almost everywhere this is possible: the surface of its basin of attraction is only 15% smaller than for the optimal case. On the other hand, its steps-to-target-efficiency is considerably worse. For most trajectories longer than a few steps, the greedy controller requires at least one more step to reach the target than in the optimal case. This means that for applications where a fast return to the target state is necessary, a control scheme similar to the greedy controller has room for improvement if multi-step trajectories are to be expected. Although it was not explicitly included in any calculations in this work, it seems worthwhile to spend a few words discussing the influence of noise and other disturbances. Since, compared to the optimal controller, the greedy controller needs more steps to converge to the target state, a system governed by the greedy controller would spend more time away from the target state and would be more susceptible to the influence of disturbances. It is thus likely, that disturbances would be tolerated better by the optimal controller. However, it is hard to say how much influence this would have in practice, since high levels of noise would also prevent trajectories of more than a couple steps to ever reach the target. For the remaining short trajectories, the performance of the optimal and greedy control schemes is much closer, and the advantage mentioned above would not be very pronounced.

Lastly, one should consider that states can be perfectly stable, yet 'far' from the tracking point, in the sense that many steps are required to reach it. For example, under the optimal control scheme, some states can only reach the target in twelve steps, even though they are stable under the passive scheme. This implies that there is a lot to gain from having a controller that does not strongly track a target state but instead also considers the viability kernel, as such a controller would not spend long times or large amounts of energy needlessly changing the system.

References

[1] R. Blickhan and R. Full, "Similarity in multilegged locomotion: bouncing like a monopode," *Journal of Comparative Physiology A*, vol. 173, no. 5, pp. 509–517, 1993.

[2] M. Hutter, C. D. Remy, M. Höpflinger, and R. Siegwart, "Slip running with an articulated robotic leg," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pp. 4934–4939, IEEE, 2010.

[3] R. Altendorfer, N. Moore, H. Komsuoglu, M. Buehler, H. B. Brown Jr, D. McMordie, U. Saranli, R. Full, and D. E. Koditschek, "Rhex: a biologically inspired hexapod runner," *Autonomous Robots*, vol. 11, no. 3, pp. 207–213, 2001.

[4] I. Poulakakis and J. Grizzle, "Monopedal running control: Slip embedding and virtual constraint controllers. ieee," in *RSJ International Conference on Intelligent Robots and Systems*, pp. 323–330.

[5] R. Ghigliazza, R. Altendorfer, P. Holmes, and D. Koditschek, "A simply stabilized running model," *SIAM Journal on Applied Dynamical Systems*, vol. 2, no. 2, pp. 187–218, 2003.

[6] H. Geyer, A. Seyfarth, and R. Blickhan, "Spring-mass running: simple approximate solution and application to gait stability," *Journal of theoretical biology*, vol. 232, no. 3, pp. 315–328, 2005.

[7] J. D. Karssen, M. Haberland, M. Wisse, and S. Kim, "The optimal swing-leg retraction rate for running," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pp. 4000–4006, IEEE, 2011.

[8] M. Ernst, H. Geyer, and R. Blickhan, "Spring-legged locomotion on uneven ground: a control approach to keep the running speed constant," in *International Conference on Climbing and Walking Robots (CLAWAR)*, pp. 639–644, World Scientific, 2009.

[9] G. Piovan and K. Byl, "Reachability-based control for the active slip model," *Submitted to International Journal of Robotics Research*, 2014.

[10] H. R. Vejdani and J. W. Hurst, "Swing leg control for actuated spring-mass robots," in *International Conference on Climbing and Walking Robots (CLAWAR)*, pp. 536–542, 2012.

[11] D. G. Hobbelen and M. Wisse, "A disturbance rejection measure for limit cycle walkers: The gait sensitivity norm," *Robotics, IEEE Transactions on*, vol. 23, no. 6, pp. 1213–1224, 2007.

[12] C. Remy, K. Buffinton, and R. Siegwart, "A matlab framework for efficient gait creation," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pp. 190–196, Sept 2011.

[13] A. Seyfarth, H. Geyer, M. Günther, and R. Blickhan, "A movement criterion for running," *Journal of biomechanics*, vol. 35, no. 5, pp. 649–655, 2002.

[14] A. Schwab and M. Wisse, "Basin of attraction of the simplest walking model," in *Proceedings of the ASME Design Engineering Technical Conference*, vol. 6, pp. 531–539, 2001.