

From Stable to Chaotic Juggling: Theory, Simulation, and Experiments

M. Bühler and D. E. Koditschek¹

Center for Systems Science
Yale University, Department of Electrical Engineering

Abstract

Robotic tasks involving intermittent robot-environment interactions give rise to return maps defining discrete dynamical systems that are, in general, strongly nonlinear. In our work on robotic juggling, we encounter return maps for which a global stability analysis has heretofore proven intractable. At the same time, local linear analysis has proven inadequate for any practical purposes.

In this paper we appeal to recent results of dynamical systems theory to derive strong predictions concerning the global properties of a simplified model of our planar juggling robot. In particular, we find that certain lower order local (linearized) stability properties determine the essential global (nonlinear) stability properties, and that successive increments in the controller gain settings give rise to a cascade of stable period doubling bifurcations that comprise a “universal route to chaos.” The theoretical predictions are first verified via simulation and subsequently corroborated by experimental data from the juggling robot.

1 Introduction

We have built the “plane-juggler,” a one degree of freedom robot that juggles two degree of freedom bodies — pucks falling (otherwise) freely on a frictionless plane inclined into the earth’s gravitational field [3]. We have developed the rudiments of a “geometric language” — a family of “mirror laws” that map puck states into desired robot states — capable of translating certain abstract goals (juggling one and two pucks, catching) into robot control laws whose closed loop behavior has been shown experimentally to accomplish these tasks in a stable and robust manner [2]. We have proven, as well, that for the task of juggling one puck the mirror law is correct by resort to a local stability analysis of linearized closed loop model [3, 1]. Unfortunately, this local analysis and its intrinsically weak conclusions have been of little use in predicting the physical consequences of different gain settings [1] and have convinced us of the practical necessity of a global stability analysis. However, even for the simpler line-juggler that gives rise to a *scalar* return map, a global stability analysis with standard mathematical techniques is exceedingly difficult. This paper marshals an array of theoretical tools for the global analy-

¹This work has been supported in part by PMI Motion Technologies, the North American Philips Laboratories, INMOS Corporation and the National Science Foundation under a Presidential Young Investigator Award held by the second author.

sis of a large class of scalar maps — among them our line-juggler map — that gives considerable promise of narrowing the gap between our analysis and practice in general. We capitalize on a “by-product” of the recently burgeoning study of bifurcations, chaos, and sensitive dependence on initial conditions in qualitative dynamical systems theory [8, 11, 5, 7, 6]. Specifically, we find that certain (seemingly unrestrictive) sufficient conditions for “reading off” from the derivative at a fixed point the essential global stability properties of an entire dynamical system are met by a simplified model of our closed loop system, the line-juggler. The strong predictions concerning the global properties are validated by a gratifying correspondence between theory and experiment. In addition, the analytical predictions seem to be relevant, as well, to the plane juggler as is strongly suggested by experimental data. Moreover, the coincidence of our systems’ stability mechanisms with these special cases may be shared by the underlying stability mechanism of Raibert’s hoppers [10, 9]. This coincidence, if physically intrinsic, holds great promise for advancing the science of robotics in intermittent dynamical environments.

2 The Line-Juggler Model

This section, devoted to the illustrative one degree of freedom case — the line-juggler — portrays more simply than can the two degree of freedom experimental system the modeling process as well as the underlying ideas in our new feedback control law, or “mirror algorithm.” First, we derive the open loop model for the line-juggler. The specification of a feedback control law will then give rise the closed loop model, a scalar return map of puck impact velocities that is analyzed in the following section. Both the modeling and the control law generalize in a straightforward fashion to the two degree of freedom juggler. For a detailed discussion of these issues and the complete derivation of the plane-juggler model and mirror algorithm refer to [1, 3, 2].

2.1 The open loop model

The simple model of the one degree of freedom line-juggler is displayed in Figure 1. A puck falls freely in the gravitational field toward a prismatic robot actuator. The robot’s and puck’s positions are denoted by r and b , respectively. The task — the vertical one-juggle — is to force the puck into a stable periodic trajectory with specified impact velocity (and thus apex position). Since the robot can only provide intermittent impacts to the puck it makes sense to examine the discrete map between puck states at those interactions as a function of the robot’s in-

puts. For now, we will ignore the robot's dynamics and assume it capable of applying arbitrary inputs to the puck during these recurring interactions. We can now examine the following control problem: Given a sequence of desired puck states — the task — find a sequence of robot control inputs that achieves it.




Figure 1: The Line-Juggler Model

First we construct the discrete model that relates two successive puck states $w = (b, \dot{b})$ just before impact as a function of the robot control inputs. This process consists of modeling the puck-robot impacts and the puck's flight dynamics.

For the impact model we make the common assumption that the elastic impact can be modeled accurately by a coefficient of restitution law and that the robot's velocity \dot{r} during impact remains unchanged. Assuming the puck and the robot are moving toward each other, $(\dot{b} - \dot{r}) < 0$, then the puck velocity just after impact \dot{b}' is related to the velocity just before impact \dot{b} via

$$\dot{b}' = -\alpha\dot{b} + (1 + \alpha)\dot{r} = c(\dot{b}, \dot{r}), \quad (1)$$

where $\alpha \in (0, 1)$ denotes the coefficient of restitution. Neglecting friction during flight, the puck's flight dynamics are given by

$$\begin{bmatrix} b(t) \\ \dot{b}(t) \end{bmatrix} = \begin{bmatrix} b' + \dot{b}'t - \frac{1}{2}\gamma t^2 \\ \dot{b}' - \gamma t \end{bmatrix}. \quad (2)$$

Here $w' = (b', \dot{b}')$ denotes the initial conditions, the puck states just after impact, and γ the gravitational constant. As the impacts are modeled to be instantaneous, the puck position during an impact remains unchanged, $b' = b$. If we now combine the impact model (1) with the flight model (2) and select the time of flight and the robot velocity at impact as our robot control inputs, we obtain the discrete map between successive puck impacts as a function of the two robot control inputs,

$$f(b, \dot{b}, u_1, u_2) = \begin{bmatrix} b + c(\dot{b}, u_2)u_1 - \frac{1}{2}\gamma u_1^2 \\ c(\dot{b}, u_2) - \gamma u_1 \end{bmatrix}. \quad (3)$$

2.2 The Mirror Law

The vertical one juggle task can now be specified as a sequence of desired puck states just before impact. Selecting $w^* = (b^*, \dot{b}^*)$ as the desired constant set point of (3) indicates that the impact should always occur at the position b^* with the velocity just before impact \dot{b}^* . If w^* is a fixed point of the closed loop dynamics, then the velocity just after impact must be $-\dot{b}^*$, and this "escape velocity" leads to a free flight puck trajectory whose apex occurs at the height $b_{apex} = b^* + \frac{\dot{b}^{*2}}{2\gamma}$. Thus, a constant w^* "encodes" a periodic puck trajectory which passes forever through a specified apex point, b_{apex} .

Successful control of the vertical one juggle task is achieved via a new class of feedback algorithms termed "mirror algorithms" [3]. Suppose the robot tracks exactly the continuous

"distorted mirror" trajectory of the puck,

$$r = -\kappa_{10}\dot{b},$$

where κ_{10} is a constant. In this case, impacts between the two do occur only when $(r, b) = (0, 0)$ with robot velocity

$$\dot{r} = -\kappa_{10}\dot{\dot{b}}. \quad (4)$$

For simplicity we will assume that the desired impact position is always selected to be $b^* = 0$. Any other impact position can be achieved by shifting the coordinate frame for robot and puck to that position. Now solving the fixed point condition $\dot{b}' = c(\dot{b}^*, \dot{r}(\dot{b}^*)) = -\dot{b}^*$ for κ_{10} using (1) and (4), yields a choice of that constant, $\kappa_{10} = (1 - \alpha)/(1 + \alpha)$ which ensures a return of the puck to the original height. Thus a properly tuned "distortion constant," κ_{10} will maintain a correct puck trajectory in its proper periodic course.

The ability to maintain the vertical one-juggle — fixed point condition — with such a simple mirror control law is an encouraging first step, but still impractical, as it is not stable. The second idea at work which will assure stability is borrowed from Marc Raibert [10], who also uses the total energy for controlling hopping robots. In the absence of friction, the desired steady state periodic puck trajectory is completely determined by its total vertical energy,

$$\eta(w) = \frac{1}{2}\dot{b}^2 + \gamma b.$$

This suggests the addition to the the original mirror trajectory,

$$r = -\kappa_1(w)b; \quad \kappa_1(w) \triangleq \kappa_{10} + \kappa_{11}[\eta(w^*) - \eta(w)], \quad (5)$$

of a term which "servos" around the desired steady state energy level. Thus, implementing a mirror algorithm is an exercise in robot trajectory tracking wherein the reference trajectory is a function of the puck's state.

2.3 The closed loop model

We will now show that all impacts between the one degree of freedom robot and the puck under the mirror law (5) must occur at the desired position $b = b^* = 0$. The impact conditions $r - b = 0, \dot{r} - \dot{b} > 0$ translate via equation (5) and its derivative into

$$[1 + \kappa_1(w)]b = 0, \quad (6)$$

$$(1 + \kappa_1(w))\dot{b} < 0. \quad (7)$$

Here we exploit $\dot{\kappa}_1 = -\kappa_{11}\dot{\eta} = 0$ as we ignore friction. As noted earlier, we restrict ourselves to puck impact velocities pointing down in the gravitational field toward the robot, $\dot{b} < 0$. Therefore, condition (7) is equivalent to $1 + \kappa_1(w) > 0$ and now the only solution to (6) is the desired puck robot impact position $b = 0$.

To construct the closed loop impact map we must now evaluate the effective robot control inputs at impact. The time of flight and the robot impact velocity is

$$u_1 = \frac{2\dot{b}'}{\gamma} \quad \text{and} \quad u_2 = \dot{r} = -\kappa_1\dot{b}.$$

Substituting these robot control inputs in (3) we obtain the scalar map of puck impact velocities just before impact at the

invariant impact position $\dot{b}^* = 0$,

$$f(\dot{b}) = \dot{b} \left(1 - \beta(\dot{b}^2 - \dot{b}^{*2}) \right), \quad (8)$$

where $\beta = \kappa_{11} \cdot (1 + \alpha)/2$.

In an effort to synthesize a more realistic closed loop map than (8) that will serve us in predicting experimental results we now include coulomb friction between puck and sliding surface. Furthermore, in order to prevent the puck from falling off the sliding plane, we incline the juggler in the gravitational field away from the vertical by an angle δ , which also has the effect of decreasing gravitational acceleration. Now the dynamics of the puck in isolation are described by

$$\ddot{b} = -\gamma \cos \delta - \text{sgn}(\dot{b}) \mu_{fric} \gamma \sin \delta.$$

Here μ_{fric} denotes the friction coefficient for dry friction. Proceeding now analogously to the frictionless case, we apply the same mirror law (5) and obtain the closed loop impact map corresponding to (8),

$$f(\dot{b}) = \zeta \dot{b} \left(1 - \beta(\dot{b}^2 - \dot{b}^{*2}) \right), \quad (9)$$

where β is defined as above and $\zeta = \zeta(\mu_{fric})$.

Note that the impact maps (8) and (9) are only defined for positive velocities after impact. This restricts our previous domain $\dot{b} < 0$ for the vertical puck velocity just before impact for both cases without and with friction to the domain W defined by

$$\bar{b} < \dot{b} < 0 \quad \text{where} \quad \bar{b} \triangleq -[\dot{b}^{*2} + \frac{1}{\beta}]^{\frac{1}{2}}.$$

3 The Stability Properties of Unimodal Maps

We now sketch our theoretical tools derived from the Singer-Guckenheimer theory (for a complete derivation, see [4]) and describe their relevance to the present application. Bifurcation plots are generated on the computer as an illustration of the theoretical statements and for purposes of comparison with the experimental data presented in Section 4.

3.1 The Singer-Guckenheimer Theory

3.1.1 S-unimodal Maps

Singer and Guckenheimer stated their results for a very particular class of functions which preserve the unit interval, called *normal S-unimodal* maps. These functions increase strictly towards a unique maximum and strictly decrease over the remainder of the interval. Moreover, they have a negative *Schwarzian Derivative* [11] except at the maximum. Rather than seeking global asymptotic stability of a fixed point, we will content ourselves with *essential global asymptotic stability*. This property holds when the set of initial conditions that fails to converge to the fixed point has measure zero. Note that for all engineering purposes essential global asymptotic stability is indistinguishable from global asymptotic stability.

Singer showed that normal S-unimodal maps can have at most one attracting periodic orbit [11]. Guckenheimer showed that

the domain of attraction of such attracting orbits includes the entire unit interval with the possible exception of a zero measure set [8]. Thus, an asymptotically stable orbit of a normal S-unimodal map is essentially globally asymptotically stable. Although these strong results are stated in terms of the apparently restrictive class of normal S-unimodal maps, they extend as well to all differentiable conjugates. Namely, say that g is a *smooth S-unimodal* map if there is some normal S-unimodal map, f , to which g is differentially conjugate — i.e. there exists a smooth and smoothly invertible function, h such that $g = h \circ f \circ h^{-1}$. It is straightforward to show that an attracting orbit of a smooth S-unimodal map is essentially globally asymptotically stable [4].

Smooth S-unimodal maps form a sufficiently large family that this theory appears to have broad engineering applicability. For example, we demonstrate below that the line-juggler map falls within this class. Moreover, we have shown that simplified models of Raibert's hopping robots give rise to smooth S-unimodal maps as well [9].

3.1.2 Bifurcations of Unimodal Maps

If we now consider *one-parameter families* of unimodal maps, then certain strong and essentially universal (i.e. independent of the particular parametrized family) properties hold true. Namely we can expect predictable structural changes in the qualitative dynamics pertaining to almost identically related values of the parameter, entirely independent of the details of the particular family. Let g_μ be a unimodal map for each μ in some real interval, $\mathcal{M} \subset \mathbb{R}$. A particular value, μ_0 , is said to be a *bifurcation point* if there is no neighborhood of μ_0 in \mathcal{M} such that g_μ is conjugate to g_{μ_0} when μ is in that neighborhood. Intuitively, the qualitative behavior of the dynamics changes around a bifurcation point.

Now suppose that $\{g_\mu\}$ is a family of smooth S-unimodal maps. If there is an interval of values $\mathcal{M} \triangleq (\mu_0, \mu_\infty)$ of μ for which $g_{\mu_0}(c) = c$ and $g_{\mu_\infty}(c) = 1$, then we shall say that g_μ is a *full family* [7]. A full family must exhibit an accumulating cascade of period doubling bifurcations: i.e., from an asymptotically stable period one orbit until μ_1 , to an asymptotically stable period two orbit until μ_2 , an asymptotically stable period four orbit until μ_3 , and so on. Thus, unimodal families give rise to theoretically determined global bifurcation diagrams. A typical such bifurcation diagram is displayed in Figure 2 as taken straight out of [5].

Unimodal period doublings have a universal structure. Denote as μ_n those points in the bifurcation diagram where there is a bifurcation from length 2^{n-1} to 2^n . Then the ratios $\frac{\mu_n - \mu_{n-1}}{\mu_{n+1} - \mu_n}$ converge to some universal number $\delta = 4.66920\dots$ *regardless of the family or the details of the parametrization* [5]. We will show later that our line-juggler satisfies these conditions and indeed, from the simulated bifurcations diagram we can verify this universal number.

3.1.3 Practical Summary

In the fortunate case of encountering a smooth S-unimodal return map, f , the job of determining the essential global limit behavior of its associated dynamics reduces to simple algebra and calculus. After finding the fixed points of the N^{th} iterate




Figure 2: Bifurcation Diagram for $f_\mu = 1 - \mu x^2$ as Shown in Collet and Eckmann

of f (algebra), we compute the magnitude of its first derivative (calculus) at a fixed point. If that magnitude is less than unity then we may expect *all* experiments performed upon the corresponding physical apparatus (with gains set to the appropriate values) to result in periodic steady state behavior exhibiting no more than N distinct states.

If, moreover, we encounter a one-parameter family of smooth S-unimodal maps, and the family is full, then appropriate adjustments of the parameter will afford any conceivable stable periodic behavior. Eventually, when the period, N , gets to be a sufficiently large number, our ability to distinguish periodic from “chaotic” steady state behavior will be compromised by the imprecision of our measurements.

3.2 Applications to the Line-Juggler Return Map

We now explore the implications of the preceding theory for our particular dynamics. In order to be directly applicable to our physical apparatus, we will use the models that include friction in the following sections. The case without friction can readily be recovered by setting $\zeta = 1$.

The original map for successive impact velocities (9) has a fixed point at

$$\dot{b}_{fp} = -\sqrt{\dot{b}^{*2} + \frac{1-1/\zeta}{\beta}} \quad (10)$$

and a unique minimum at $\dot{b}_c = \frac{1}{\sqrt{3}}\dot{b}^*$ with $f(\dot{b}_c) = \zeta^2 \dot{b}_c^3$. Moreover, it is not hard to see that f is smooth S-unimodal [4]. Thus from the preceding results we immediately have

Theorem 1 ([4]) *The mirror algorithm for the line-juggler results in a successful vertical one juggle which is essentially globally asymptotically stable as long as*

$$0 < \beta < \frac{2/\zeta - 1}{\dot{b}^{*2}}.$$

It is easy to verify that g_μ is also a *full family*: For $\mu \in$

$(\mu_0, \mu_\infty) = (\frac{3}{2}, \frac{3}{2}\sqrt{3})$, we obtain $g_{\mu_0}(c) = c$ and $g_{\mu_\infty}(c) = 1$. Now we know that after the fixed point \dot{b}_{fp} becomes unstable, the map f will exhibit period doublings that will eventually lead to chaotic behavior, as predicted before. This is confirmed in Figure 3, which show the bifurcation diagram (obtained via simulation) for our specific line-juggler map (with friction). The ratio $\frac{\mu_n - \mu_{n-1}}{\mu_{n+1} - \mu_n}$ is evaluated for the first three bifurcations, $n = 2$, directly from the two figures, as 4.6 with an accuracy of about 0.3 due to the large β stepsize. This value is close to the expected limiting value for $n \rightarrow \infty$ of $\delta = 4.66920\dots$

Given the settings $\mu_{fric} = 0.16$, $\alpha = 0.7$, $\zeta = 0.9115$, and $\dot{b}^* = -125$, one can predict from $f'(\dot{b}_{fp}) = -1$ the first β -bifurcation value $7.6 \cdot 10^{-5}$. The corresponding value of $\kappa_{11} = 2\beta/(1 + \alpha) = 8.9 \cdot 10^{-5}$ is confirmed accurately from Figure 3.




Figure 3: Line-Juggler with Friction: Simulated Bifurcation Diagram

4 Experimental Results

In this section we present experimental data to validate the models developed in Section 2 and the theoretical predictions by comparisons with simulated data from Section 3.2. In Section 4.1 devoted to the line-juggler, we will use the theoretical insights presented in Section 3.1 for scalar return maps to predict the dynamical behavior of the physical apparatus. The correspondence between simulated and experimental data is gratifying. We are able to predict and verify experimentally the transients of the stable fixed points as well as higher period stability properties, specifically bifurcations to stable period two and period four orbits. In Section 4.2 we back with experiments our speculations [4] about the applicability of the theory to the two degree of freedom juggler.

4.1 Line-Juggler Experiments

All our past experiments were based on the plane-juggler, which allows for planar puck motion which is controlled by impacts with a revolute motor, as depicted in Figure 4. The one-juggle is implemented by restricting the puck motion to a vertical line. The resulting closed loop impact map, when using the proper mirror algorithm for the motor, is identical to the one-juggle

map (9) [4].




Figure 4: The Line-Juggler Implementation

We have shown in the previous section that the local dynamical behavior is essentially global. The data in Figure 5 confirm that the transients can be predicted by recourse to local linear analysis of the scalar impact map. Evaluating the derivative of (9) at the fixed point (10) for the four gain settings $\kappa_{11} = 3/5/7/9 \cdot 10^{-5}$ shown in the figure, we predict locally an overdamped, critically damped, underdamped and an unstable response, respectively. This behavior is confirmed even from large initial conditions ("globally") on the juggling apparatus. When inspecting the transient for the last gain setting $\kappa_{11} = 9 \cdot 10^{-5}$ closely we see that it maintains a small oscillation, the predicted onset of instability. The fixed point in the presence of friction (10) depends on the gain setting κ_{11} . Figure 5 confirms as well our ability to predict the steady state values for the vertical impact velocities with less than 3% error.

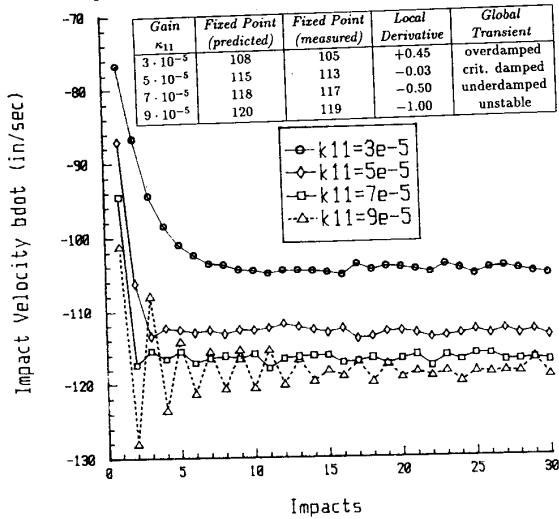


Figure 5: Line-Juggler Transients: Experimental Data

When increasing the gain beyond the predicted first κ_{11} bifurcation value of $8.99 \cdot 10^{-5}$, we expect to see a stable period two orbit. Visually, this shows the puck oscillating between two different juggling heights. Indeed, for $\kappa_{11} = 11 \cdot 10^{-5}$, Figure 6 shows the divergence from the unstable fixed point of $f \circ f$.

in/sec) towards a stable period two — a stable fixed point of $f \circ f$.




Figure 6: A Transient Toward Period Two: Experimental Data

The strength of our analytical predictions is demonstrated further when increasing the gain to $\kappa_{11} = 12.25 \cdot 10^{-5}$: an observable period four trajectory emerges, as shown in Figure 7. As the gain increases, the higher period orbits become more and more sensitive to perturbations. This together with the slower transient recovery from these perturbations causes the period four orbit to appear and disappear at unregular intervals, and explains the fact that we have not seen higher period orbits than period four.

This most satisfying correspondence between analytical predictions and experimental data is summarized in Figure 8: the experimentally acquired bifurcation diagram for the line-juggler. It coincides with good accuracy with the model prediction in Figure 3 up to the second bifurcation. This plot was acquired in an




Figure 7: A Period Four Sequence: Experimental Data




Figure 8: The Line-Juggler Bifurcation Diagram: Experimental Data

analogous fashion to the simulated plot. The puck was dropped at a height corresponding to $\dot{b} = -100$ in/sec. We discarded the first 20 impacts to assure steady state juggling and logged the following 50 impacts. For the last four gain settings we logged the following 100 impacts since the spread of impact velocities increased. This was repeated for the κ_{11} gain range in $0.25 \cdot 10^{-5}$ increments.

4.2 Planar Juggler Experiments

We now remove the guiding wires in Figure 4 to obtain the original plane-juggler. The one-juggle mirror algorithm can be generalized to accomplish the vertical one-juggle task on the plane-




Figure 9: The Bifurcation Diagram for the Unconstrained (2dof) Juggler: Experimental Data

juggler [4]. However, the complexity of the new closed loop system admits only a local linear proof of correctness. After showing that the plane-juggler dynamics at a fixed horizontal position are identical to the line-juggler, we argue that even in the presence of horizontal position errors the line-juggler analysis should serve us for qualitative predictions. This can be verified in Figure 9 which displays an experimentally measured bifurcation diagram for the planar juggler. Again, the puck is dropped with an initial height which results in an initial vertical impact velocity of about -100 in/sec. After the first 20 impacts have passed, we log the next 100 impacts shown in the plot. We start with an initial value of $6 \cdot 10^{-5}$ and record a run until $\kappa_{11} = 12.5 \cdot 10^{-5}$.

This bifurcation plot is almost identical to that of the line-juggler in Figure 8. Due to the larger perturbations present, the spread of the data is considerable larger. However, we can still identify a bifurcation to a stable period two orbit after the value of $\kappa_{11} = 10 \cdot 10^{-5}$, which is close to the value $\kappa_{11} = 8.99 \cdot 10^{-5}$ predicted from the theory of the simpler system.

Acknowledgements

We would like to thank John Guckenheimer for his kind tutorial efforts on our behalf.

References

- [1] M. Bühler, D. E. Koditschek, and P. J. Kindlmann. A Simple Juggling Robot: Theory and Experimentation. In V. Hayward and O. Khatib, editors, *International Symposium on Experimental Robots*. Springer-Verlag, Montréal, Canada, Jun 1989.
- [2] M. Bühler, D. E. Koditschek, and P. J. Kindlmann. Planning and Control of Robotic Juggling Tasks. In H. Miura, editor, *Fifth International Symposium on Robotics Research*, pages 270–281, Tokyo, Japan, 1989. MIT Press.
- [3] M. Bühler, D. E. Koditschek, and P. J. Kindlmann. A family of robot control strategies for intermittent dynamical environments. *IEEE Control Systems Magazine*, 10(2):16–22, Feb 1990.
- [4] Martin Bühler. *Robotic Tasks with Intermittent Dynamics*. PhD thesis, Yale University, May 1990.
- [5] P. Collet and J. P. Eckmann. *Iterated Maps on the Interval as Dynamical Systems*. Birkhäuser, Boston, 1980.
- [6] Robert L. Devaney. *Introduction to Chaotic Dynamical Systems*. Addison Wesley, Reading, MA, 1987.
- [7] J. Guckenheimer and P. Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, New York, 1983.
- [8] John Guckenheimer. Sensitive dependence to initial conditions for one dimensional maps. *Communications in Mathematical Physics*, (70):133–160, 1979.
- [9] D. E. Koditschek and M. Bühler. Analysis of a simplified hopping robot. *The International Journal of Robotics Research*, (to appear).
- [10] Marc H. Raibert. *Legged Robots That Balance*. MIT Press, Cambridge, MA, 1986.
- [11] David Singer. Stable orbits and bifurcations of maps of the interval. *SIAM J. Applied Mathematics*, 35(2):260–267, Sep 1978.