

A Comparison of Regulation and Entrainment in Two Robot Juggling Strategies

Preprint of a paper to be included in
Proceedings of the Sixth International Symposium on Robotics Research
Hidden Valley, PA
October 1993

D. E. Koditschek *

Artificial Intelligence Laboratory
EECS Department, College of Engineering
University of Michigan, Ann Arbor

Abstract

This paper presents a stability analysis of a simple Shannon juggler and contrasts its properties with those of the mirror juggler. It discusses the need to develop a systems theory for strongly coupled tunable oscillators.

1 Introduction

This paper is concerned with the analysis of two very different robot juggling methodologies — one proposed by Bühler [2]; the other advanced by Atkeson and Schaal [14]. I will refer to the first method as the *mirror* scheme and the second as the *Shannon* scheme. The salient empirical features of these contrasting approaches may be summarized as follows:

Regulation: both strategies result in various laboratory mechanisms that can bat one or more balls into a stable periodic orbit

Entrainment: both strategies can achieve stable phase relationships between the actuated and unactuated degrees of freedom.

Feedback: The mirror strategy employs a continuous stream of sensory information about a ball's flight to develop a robot response.

The Shannon strategy provides the robot no information about the ball at all.

Domain of Attraction: The mirror scheme has a very large tolerance for poor conditions at startup and can often recover from adversarial perturbations in the ball's flight pattern.

The Shannon scheme is much more sensitive to startup conditions and will generally fail if the ball's flight pattern is disturbed.

The relatively greater robustness (ability to handle disturbances) of the mirror jugglers comes at relatively greater operational cost (dependence upon a continuous sensor stream) in comparison to the Shannon jugglers. One wishes to weigh the relative improvements in capability achieved against the additional operational cost. One wishes for a means of hybridization — a “mixing strategy” through which sensory costs can be adjusted to develop a desired level of robustness.

This paper pursues an analytical framework within which I hope such tradeoffs can eventually be analyzed and hybridization schemes achieved. But it falls far short of such achievements. Using standard techniques of dynamical systems theory I present here what I believe to be the first stability analysis of a Shannon juggler. Using a “Task Encoding” formalism developed in some of my earlier work [9], I am able to display the mechanical feedback policy implicit in the Shannon juggler's motion. This can be directly compared with the mechanical feedback policy induced by the mirror juggler's sensor based strategy. The latter works according to a nice regulation theory for tunable oscillators that we have described in previous papers [10, 1]. I do not yet, however, have much

*Supported in part by the National Science Foundation under grant IRI-9123266, and in part by the Special Semester on Spatial and Temporal Dynamics of the Université de Montréal Centre de Recherches Mathématiques.

to say about hybridization because I do not completely understand the entrainment mechanism implicit in the Shannon machines. A systems theory for coupled oscillators requires both regulation and entrainment techniques.

The paper is organized as follows. After a brief digression on coupled oscillators that fills out the last remarks, I introduce the notation and models relevant to the juggling analysis. I next review our previous analysis of the mirror jugglers and their regulation properties. I finally offer a preliminary analysis of a Shannon juggler. This underscores the desirability of developing entrainment analysis and synthesis techniques for robotics applications.

2 Missing: A Systems Theory for Coupled Oscillators

A nonlinear systems theory worthy of practice would provide techniques for coupling together tunable oscillators.¹ This requires methods of regulation and methods of entrainment. I believe we have some practicable regulation theories, and I sketch one below. I don't yet know of any entrainment theory relevant to robotics.

2.1 Regulation of Simple Oscillators

Figure 1 depicts a tunable oscillator — a one degree of freedom system with an attracting cycle. Tuning a gain parameter results in a range of qualitatively different behaviors that consistently emerge from arbitrary initial conditions. Said another way, every state space plane defined by a fixed value of the tuning parameter quickly contracts onto the selected cycle. These tunable limit sets are effectively regulated by some mechanism that is not troubled by the dynamical bifurcations on the limit set itself.

We have proposed a fairly general theoretical mechanism for building one degree of freedom regulated tunable oscillators. This theoretical point of view is motivated by our laboratory experience with mirror jugglers [1, 2] and our simplified analysis of Raibert's hopper [10]. In brief, the periodic behaviors of a one degree of freedom machine are greatly constrained by the topology of the circle,

¹Throughout this paper the term *oscillator* denotes a nonlinear dynamical system with a globally attracting compact forward limit set. For example, point attractors are trivial oscillators. The term *tunable* denotes oscillators that persist (despite likely bifurcations in the limit set topology) over large volumes in the parameter space that defines the dynamics.

S^1 . We have pointed out that certain consequences of this fact described by dynamicists in recent years [8, 4] are of great use in regulating such simple oscillators.




Figure 1: A tunable oscillator. For each value of the tuning parameter there is a globally attracting limit set. The topological structure of the limit set does not change in the course of smooth variations in the tuning parameter despite the transition through dynamical bifurcations. For example, this caricatures the transition from catching to batting to “limping” described in [2].

2.2 Entrainment of Multiple Oscillators

When two (or more) simple oscillators are coupled together, we would like to imagine that the “cross product of Figure 1 with itself” yields a system depicted in Figure 2. Here, the system has a limit set taking the form of a cross product of the two isolated limit sets (a torus is the cross product of two circles) and a specified phase relation that governs the comparative timing of the two periods.

When the transverse dynamics in each subsystem are “fast” relative to the motion on the cycle (that is, all motions on the plane selected by the tuning parameter converge quickly to the circles in Figure 1) and the coupling between the two subsystems is “weak” then the theory of isolated invariant sets may be used to deduce the situation of Figure 2 [7]. In general, the coupling between the subsystems discussed in this paper is impulsive and will not be likely to satisfy the “weak” properties characteristic of the small perturbations under which such uniform normal hyperbolicity ideas apply.

Much needed is a systems theory for entrainment relevant to robotics settings.




Figure 2: The relative phase of two coupled oscillators. In the case depicted here, each has been tuned to give a period one limit cycle in isolation — their relative phase is entrained at 3:1. This may be seen as the “slope” of the limit curve winding around the torus — the cross product of their isolated limit cycles. While we can build instances of such systems we do not presently have a candidate theory for constructing general coupling mechanisms that yield a specified phase.

3 Notation, Models, and Problem Formulation

This section presents a statement of the one-juggle control problem. The statement is preceded by a sketch of the physical model and followed by some musings about the information structure of robotics problems in general.

3.1 Two Physical Models

The state of a one degree of freedom mass, is denoted $x = (\chi, \dot{\chi}) \in \mathbb{R}^2$. Newton’s equation of free flight,

$$\ddot{\chi} = -\gamma,$$

gives, upon two integrations, a map on the plane,

$$\begin{aligned} F^t(x) &= A_t x - a_t \\ A_t &:= \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} \\ a_t &:= \gamma t \begin{bmatrix} t/2 \\ 1 \end{bmatrix}, \end{aligned} \quad (1)$$

which we shall use throughout the sequel.

Suppose the point mass x deforms and loses some energy during an instantaneous impact in a

manner captured by the *coefficient of restitution*,

$$0 \leq \alpha \leq 1.$$

If x collides with a much more massive body, y then the massive body’s state is unchanged by the impact whereas, the new state of smaller body is given by

$$\begin{aligned} C_y(x)\alpha &= R_\alpha x + (I - R_\alpha)y \\ R_\alpha &:= \begin{bmatrix} 1 & 0 \\ 0 & -\alpha \end{bmatrix}. \end{aligned} \quad (2)$$

3.2 Two Control Problems

The setting of interest is the intermittent interaction between a massive one degree of freedom “robot,”

$$r = (\rho, \dot{\rho})$$

and a lighter body,

$$b = (\beta, \dot{\beta})$$

falling in the earth’s gravitational field and constrained to move along the same axis as the robot. We are interested in studying the evolution of the body’s states in consequence of actions taken by the robot. Assume, for the sake of simplicity, that the robot is a lossless torque actuated mechanical system,

$$\dot{r} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} r + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \tau, \quad (3)$$

that is, a double integrator.

The body falls from some initial position and velocity, b , according to Newton’s law (1) and reacts to a collision with the robot at some state r according to the coefficient of restitution law (2). Now let b denote the state of the body just prior to an impact. Suppose the robot impacts with velocity u_2 and allows the body to fall freely for the next u_1 interval of time. Then the state of the body just prior to the next impact is given by

$$f(b, u) := F^{u_1} \circ C_{u_2}(b) \quad (4)$$

Logically, any effect of the robot on the body may be described with regard to this model. Thus any desired robot task may be encoded as a control problem:

Environmental Control Problem: Given a desired body trajectory, $\{b_k^*\}_{k=1}^\infty$, find a robot *impact schedule* $\{u_k\}_{k=1}^\infty$ such that the resulting body trajectory approach asymptotically the desired.

As stated, this logical problem is entirely trivial, since, for fixed time, $u_1 \equiv \text{constant}$, (4) is a completely controllable linear time invariant system. The interest arises when we require a solution in the form of a torque strategy for the robot (3)

Robot Control Problem: Find a continuous a robot control law, $\tau(t)$, such that $r(t)$ in (3) realizes a desired impact schedule, $\{u_k\}_{k=1}^{\infty}$.

This is not a conventional control problem and there is no body of tradition to fall back on in devising a solution.

3.3 Information Flow: Passive Strategies, Sensorless Manipulation, and Feedback

The distinction between the two control systems, (4) and (3), is easily understood and instantly confronted by anyone who has tried get a robot to do a job in the physical world. On the one hand there is the work we would like to robot to do; on the other hand there is the command language within which we are allowed to communicate our desires. Hopefully, the utility of drawing attention to this distinction will be made clear in the next few portions of the paper wherein it affords a comparison between two rather different approaches to the one-juggle problem. I believe that more careful attention to this distinction might advance the growing number of experimental and theoretical discussions of information flow in robotics.

The distinction between data and information has been well noted of late. Irrelevant sensing is a disease [12] or an addiction [11] in this emerging understanding. Despite the recent bad press, it seems safe to assert that no one seriously proposes to eliminate perceptual processing. Instead, we begin to see serious efforts to rationalize the design of sensors [6] or explore the logical limits of perception [5]. These attempt explicitly to illuminate the role of sensing against the dynamics of incremental progress toward an environmental goal state. Perhaps more attention to the distinction between task level and agent level would be of help.

3.3.1 Environments Don't Distinguish Between Passive and Active Strategies

A "passive" strategy in the context of an environmental control problem makes no sense because one is precisely trying to express a job to be done in terms of what some abstract agency ought to

accomplish in the world. However, it is entirely possible to imagine fitting a mechanism into the environment in such a fashion that the mutual interactions accomplish the job with no user inputs to the mechanism. Presumably, there is a trade-off in programmability vs. operational cost. A mechanism that can be fit into an environment to achieve a specified task with no input seems less likely to be capable of achieving any other task in the same or another environment than one that requires input. In this light, the abstracted environmental control problem represents a normative measure against which the relative advantages and disadvantages of mechanisms that accomplish the same task may be compared.

3.3.2 There are no Sensorless Robots or Environments — Only Sensory Disconnects Between Them

It is hard to imagine why anyone would spend time thinking about a control system whose state was entirely inaccessible to observation. This would be a system that doesn't do anything. Thus, it seems pointless to speak of sensorless strategies for solving the environmental control problem or the robot control problem when considering either in isolation. To the extent that the state matters in the problem, it may be reconstructed from the accumulated history of its effects. In contrast, it is fascinating to consider the extent to which the state of the environment need be communicated to the robot. A sensorless robot strategy, then, is one wherein the robot receives no information about the environment's state. The Shannon juggler is such a sensorless strategy: it employs a robot that may require self measurements and, as shown below, induces an environmental feedback law. Distinguishing robot from environment state measurements affords the possibility of comparing strategies that accomplish similar tasks with different communications costs between the two.

4 Solution I: The Mirror Juggler

In his thesis work with me, Bühler proposed an approach to the environmental control problem as follows. Consider a robot control strategy for (3) that achieves the trajectory

$$r = m(b) \quad (5)$$

where m is a second order function, that is,

$$m(b) = \begin{bmatrix} \mu(b) \\ D\mu(b) \cdot [\dot{\beta}, -\gamma] \end{bmatrix},$$

so that the second component of the image is the derivative of the first along any body flight trajectory.

4.1 The Closed Loop Environmental Dynamics

Collisions occur on the curve

$$\mathcal{S} = \{b : \sigma(b) = 0\}; \quad \sigma(\beta, \dot{\beta}) := \beta - \mu(\beta, \dot{\beta})$$

which we now assume is the graph of some function of velocity, s ,

$$\mathcal{S} = \{(\beta, \dot{\beta}) : \beta = s(\dot{\beta})\}.$$

This implies that the robot's velocity at impact is a function of the body's,

$$v(\dot{\beta}) := D\mu(s(\dot{\beta}), \dot{\beta}) \cdot \begin{bmatrix} \dot{\beta} \\ -\gamma \end{bmatrix}.$$

Another way of stating that collisions occur on \mathcal{S} is to note that the impact schedule induced by (5) has the property of rendering \mathcal{S} an invariant set with respect to the logical system (4). Thus, the time of flight between impacts, t_d , is such that for every robot velocity at impact, $\dot{\beta}$,

$$\sigma \circ f((s(\dot{\beta}), \dot{\beta}), (t_d, v(\dot{\beta}))) = 0.$$

Assume again that this curve is the graph of a function of velocity

$$t_d = \tau_d(\dot{\beta}).$$

It is now clear that the robot's actions induce an impact schedule of the form

$$u = g_m(b) = \begin{bmatrix} \tau_d(\dot{\beta}) \\ v(\dot{\beta}) \end{bmatrix}, \quad (6)$$

leading to an effective closed loop system

$$f_m(b) = f(b, g_m(b)). \quad (7)$$

4.2 Analysis Lags Behind Empirical Results

Bühler's choice of mirror law,

$$\mu(\beta) := \left(\kappa_0 + \kappa_1 \left[\eta(\beta, \dot{\beta}) - \eta^* \right] \right) \beta \quad (8)$$

$$\eta = \frac{1}{2} \dot{\beta}^2 + \gamma \beta,$$

accomplishes the desired task in the laboratory [2]. Moreover, assuming that the robot delivers the mirror trajectory (5) exactly, the resulting environmental dynamics (7) takes the form of a negative Schwartzian unimodal map [1] for which there can be found a very nice regulation theory [10]. This theory yields the tunable oscillator depicted in Figure 1. Namely, by tuning the gains $\kappa_0, \kappa_1, \eta^*$, one passes from a zero period "catch" to a period one "juggle" to a period two "limp," and on up the well explored period doubling route to chaos [2].

In point of fact, Bühler's laboratory implementation instantiated a version of the mirror policy that coupled the "gedanken" one degree of freedom oscillator described here with a transverse period zero oscillator to obtain a planar juggle, as depicted in Figure 3 [3]. This has been general-




Figure 3: Bühler's planar one-juggle couples a vertical period one oscillator to a horizontal period zero (attracting equilibrium state) oscillator [2].

ized by Rizzi to the spatial one juggle symbolized in Figure 4 [13]. We are well on the way to an an-




Figure 4: Rizzi's spatial one-juggle couples a vertical period one oscillator to two horizontal period zero (attracting equilibrium state) oscillators [13].

alytical stability proof for the situations depicted in Figure 3 and Figure 4, but there does not seem yet to have emerged any theoretical perspective that might guide the analysis, much less future synthesis.

We have devised empirical coupling rules for building two ball jugglers out of the cross product of two isolated one juggles that work quite well in the laboratory as symbolized in Figure 5. These have been very successfully implemented both on the plane [2] as well as in space [13]. But we do not presently know how to prove that the coupling



Figure 5: We have built planar and spatial two-jugglers using a coupling rule and phase regulation control term in generalizations of the mirror law [13].

rules and phase regulation terms are correct.

In short, our lack of an entrainment theory for coupled oscillators continues to impede analytical understanding relative to empirical capability.

5 Solution II: The Shannon Juggler

Atkeson and Schaal have achieved experimentally the same task through a very different strategy — one inspired by earlier work of Shannon. I now present the stability analysis (a somewhat stylized version) of this scheme.

5.1 A Lossless Relaxation Oscillator




Figure 6: A gravity driven relaxation oscillator.

Consider a robot control strategy that achieves the relaxation oscillator depicted in Figure 6. This trajectory characterizes the behavior of a perfectly elastic ball, $r(t) = (\rho, \dot{\rho})(t)$, bouncing in the earth’s gravitational field subject to lossless collisions with the ground at impact and lossless free flight in between. Fixing coordinates, say the ball collides with the ground at time $t = 0$, in the state $-r_0 = (0, -\dot{\rho}_0)$. Lossless collisions correspond to

the coefficient of restitution unity, so that the evolution of the ball’s state from zero height at the “post bounce” velocity, $+\dot{\rho}_0$, is now given by

$$r(t) = F^{\tau_p(t)}(r_0) \quad (9)$$

where τ_p denotes time modulo the period of flight, $T := 2\dot{\rho}_0/\gamma$.² The parameter, $\dot{\rho}_0$ is assumed positive throughout the sequel.

5.1.1 Coupling a Lossy Relaxation Oscillator

Suppose now that a much lighter ball is constrained to fall along the same vertical axis as the previous one. Denote its position and velocity by $\bar{b} = (\bar{\beta}, \dot{\bar{\beta}})$, measured at the moment the heavy ball just touches the ground, $r(t) = -r_0$.³ Again, we will suppose no energy losses in flight.

There is a region, $\mathcal{B}_1 \subset \mathbb{R}^2$, in the light ball’s state space, such that if $\bar{b} \in \mathcal{B}_1$ then there will be a collision with the heavy ball,⁴

$$\begin{aligned} 0 &= [F^{\tau_c}(\bar{b})]_1 - [F^{\tau_c}(r_0)]_1 \\ &= [A_{\tau_c}(\bar{b} - r_0)]_1 = \bar{\beta} + \tau_c(\dot{\bar{\beta}} - \dot{\rho}_0) \end{aligned}$$

before the heavy ball next collides with ground. Solving this equation for τ_c gives the function,

$$\tau_c(\bar{\beta}, \dot{\bar{\beta}}) := \bar{\beta}/(\dot{\rho}_0 - \dot{\bar{\beta}}), \quad (10)$$

that computes the time to collision. Imposing the inequalities

$$0 \leq \tau_c(\bar{b}) \leq T,$$

makes explicit the extent of the region \mathcal{B}_1 ,

$$\begin{aligned} \bar{\beta} &> 0 \\ \left[\frac{1}{2}\gamma, \dot{\rho}_0\right] \left[\begin{array}{c} \bar{\beta} \\ \dot{\bar{\beta}} \end{array}\right] &< \dot{\rho}_0^2. \end{aligned} \quad (11)$$

The state of such a light ball at the time, T , of the heavy ball’s next collision with ground (at state $-r_0$), may now be modeled by flying the two

²That is, τ_p denotes the piecewise linear periodic function coincident with the identity on the interval

$$0 \leq t < T.$$

³Note the use of the barred variables denotes the introduction of a frame of reference at variance with the one introduced in Section 3. These coordinates greatly simplify the derivation of the dynamics. We will shift back to the standard coordinates later in the paper.

⁴Denoted in the following by equating the position components — $[\cdot]_1$ — of both balls.

balls to collision, $C_{F^{r_0}(r_0)} \circ F^{\tau_c}(\bar{b})$, and then flying the light ball onward for time $T - \tau_c$. This yields a map $F^T \circ \bar{C}_{r_0}(\bar{b})$ where

$$\bar{C}_r(\bar{b}) := F^{-\tau_c(\bar{b})} \circ C_{F^{r_0}(\bar{b})(r)} \circ F^{\tau_c(\bar{b})}(\bar{b})$$

Note that

$$\bar{C}_r(\bar{b}) = A_{[-\tau_c]} R A_{\tau_c}(\bar{b} - r_0) + r_0.$$

By definition of τ_c , it follows that both

$$A_{\tau_c}(\bar{b} - r_0) = \begin{bmatrix} 0 \\ \dot{\bar{\beta}} - \dot{\rho}_0 \end{bmatrix},$$

and

$$\alpha \dot{\bar{\beta}} = -\alpha \tau_c (\dot{\bar{\beta}} - \dot{\rho}_0),$$

so that

$$\bar{C}_{r_0}(\bar{b}) = -\alpha(\bar{b} - r_0) + r_0, \quad (12)$$

and

$$F^T \circ \bar{C}_{r_0}(\bar{b}) = -\alpha A_T(\bar{b} - r_0) - r_0.$$

5.1.2 Local Stability Analysis of the Coupled Oscillator System

If the next state of the light ball after the first collision is also in \mathcal{B}_1 then its state after a second collision is given by composing the last function with itself. Assuming that all further iterates, remain in \mathcal{B}_1 as well, it would follow that the state of the light ball at the time of the heavy ball's k^{th} bounce with the ground is given by the recursion

$$\begin{aligned} \bar{b}_{k+1} &= \bar{f}_1(\bar{b}_k) \\ &:= -\alpha A_T \bar{b}_k + (\alpha - 1) A_{-\xi_1} r_0 \end{aligned} \quad (13)$$

where $\xi_1 := \alpha T / (1 - \alpha)$. We now demonstrate the existence of an open set of points in \mathcal{B}_1 whose iterates remain in that set.

The map, \bar{f}_1 has a unique fixed point,

$$\begin{aligned} e_1 &:= -(I + \alpha A_T)^{-1} \cdot (1 - \alpha) A_{-\xi_1} r_0 \\ &= \frac{1 - \alpha}{1 + \alpha} A_{-\xi_1} r_0 \end{aligned}$$

where $\epsilon_1 := 2\alpha T / (1 - \alpha^2)$. Now

$$e_1 = \begin{bmatrix} 4\alpha \dot{\rho}_0^2 / (1 + \alpha)^2 \gamma \\ -(1 - \alpha) \dot{\rho}_0 / (1 + \alpha) \end{bmatrix} \in \mathcal{B}_1,$$

since the first component is positive and

$$0 < \dot{\rho}_0^2 - \left[\frac{1}{2} \gamma, \dot{\rho}_0 \right] e_1 = \frac{2\dot{\rho}_0^2}{1 + \alpha}.$$

Moreover, $-\alpha A_T$ has eigenvalues inside the unit circle for any choice of $\dot{\rho}_0$. Thus, some open neighborhood of e_1 remains in \mathcal{B}_1 under iteration by \bar{f}_1 . This open neighborhood takes e_1 as its limit.

5.1.3 Relative Phase of Entrainment

The locally attracting 1:1 phase relationship between the two balls at the stable fixed point, e_1 , is depicted in Figure 7 A global analysis would re-

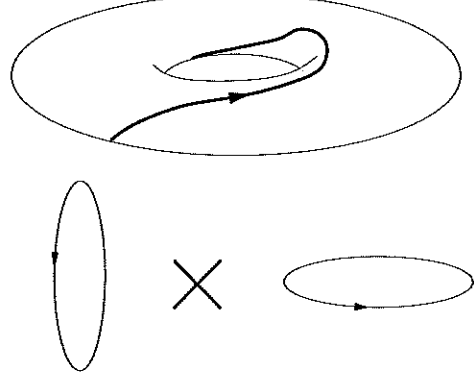


Figure 7: The locally attracting 1:1 entrained trajectory of the Shannon Juggler.

quire explicit treatment of the possible subharmonics of the light ball relative to the heavy ball. These correspond to multiple bounces by the heavy ball in between collisions with a “higher flying” light ball.

For example, to find the 2:1 subharmonic consider the set

$$\mathcal{B}_2 := F^{-T}(\mathcal{B}_1) \cap \{\bar{b} \in \mathbb{R}^2 : \bar{\beta} > 0\},$$

of light ball states that cannot collide with the heavy ball until its second bounce. For such a state, $\bar{b} \in \mathcal{B}_2$ we may model the state at the following bottom of the heavy ball as

$$\bar{f}_2(\bar{b}) = F^T \circ \bar{C} \circ F^T(\bar{b})$$

and since

$$F^T(\bar{b}) = A_{T/2} [A_{T/2} \bar{b} - 2r_0],$$

this simplifies to

$$\bar{f}_2(\bar{b}) = -\alpha A_{2T} + (3\alpha - 1) A_{\xi_2} r_0$$

where $\xi_2 := 4\alpha T / (3\alpha - 1)$.

In general, the $k : 1$ subharmonic trajectories (assume for now that $k \geq 1$) originate in the set

$$\mathcal{B}_k := F^{-(k-1)T}(\mathcal{B}_1) \cap \{\bar{b} \in \mathbb{R}^2 : \bar{\beta} > 0\},$$

of light ball states that cannot collide with the heavy ball until its k^{th} future bounce. For such

a state, $\bar{b} \in \mathcal{B}_k$ we may model the state at the following bottom of the heavy ball as

$$\begin{aligned}\bar{f}_k(\bar{b}) &= F^T \circ \bar{C} \circ F^{(k-1)T}(\bar{b}) \\ &= -\alpha A_{kT} + [(2k-1)\alpha - 1] A_{\xi_k} r_0\end{aligned}$$

where $\xi_k := k^2 \alpha T / [(2k-1)\alpha - 1]$. To determine for how long such subharmonics can persist; it is instructive to examine e_k , the unique fixed point of \bar{f}_k . Specifically, if

$$\tau_c \circ F^{kT}(e_k) \leq T$$

then e_k is a valid fixed point of the overall system. It is evidently asymptotically stable since $-\alpha A_t$ has eigenvalues inside the unit circle no matter what the value of t .

It is also clear that there may be $1:k$ harmonic trajectories — multiple impacts occurring as a result of a light ball with relatively small energy colliding again and again with the heavy ball as the latter moves through a single period.

A global model, \bar{f} , of the trajectory of the light ball, \bar{b} seems to require a careful piecing together of the various subharmonic, \mathcal{B}_k , and harmonic, \mathcal{B}_{-k} regions. It is entertaining to speculate about the various conditions on α that would or would not result in a single globally asymptotically stable periodic orbit. At present, because of the complicated appearance of multiple (sub)harmonic periodic orbits — some of them obviously locally attracting — this question remains a matter of speculation.

5.2 The Effective Impact Schedule: a Feedback Law

As shown above, no matter how many other stable harmonics arise, the 1:1 behavior has an attractor at e_1 a fixed point of the global dynamics, \bar{f} . By entraining the light ball to a stable oscillation of the same period at fixed phase, the heavy ball has “solved” the environmental control problem posed earlier. We may now examine the “decisions” of the heavy ball, considered as a robot, in terms of this originally posed problem. For this purpose, it is convenient to introduce the change of coordinates

$$h(\bar{b}) := F^{\tau_c(\bar{b})}(\bar{b}) \quad (14)$$

This transformation takes a body state, \bar{b} , at the bottom of the robot’s flight and computes the resulting body state, (just before) the instant of impact — the state b in the notation of Section 3.

The function

$$\tau_d(\bar{b}) := \tau_c \circ \bar{f}(\bar{b}) + t_p - \tau_c(\bar{b})$$

computes the body’s time of flight between an impact resulting from state \bar{b} and the next impact, the one resulting from state $\bar{f}(\bar{b})$. The function

$$v(\bar{b}) := [0, 1]^T F^{\tau_c(\bar{b})}(r_0)$$

computes the robot’s velocity at impact resulting from the body state \bar{b} at the robot’s bottom. Thus, the relaxation oscillator, $r(t)$ induces an effective impact schedule that depends upon the body’s state.

In the coordinates of system (4), this feedback law takes the form

$$u = g_m(b) = \begin{bmatrix} \tau_d \circ h^{-1}(b) \\ v \circ h^{-1}(b) \end{bmatrix}. \quad (15)$$

It induces the effective closed loop dynamics

$$f_s(b) = f(b, g_m(b)) = h \circ \bar{f}_s \circ h^{-1}(b).$$

This reveals the task level feedback implicit in the Shannon juggler.

6 Conclusion

Both the Mirror and Shannon juggling strategies offer a solution to the stable ball batting problem. Both have been implemented quite successfully and generalized to machines and environments of higher degrees of freedom. The mirror strategy results in a limit set with an essentially global domain of attraction. It admits a rather complex regulation capability (Figure 1) but admits no harmonics. The Shannon juggler — at least in the stylized version developed here — presents the opposite situation where one observes a discontinuous piecewise linear map revealing a simple regulation structure (locally linear periodic orbits) and a complex entrainment structure. In particular, the Shannon strategy, since it will give rise to multiple stable harmonics in general, can only offer local convergence to a specifically desired cycle. On the other hand, the Shannon strategy requires no active sensory information be presented to the robot from the environment.

By introducing the task encoding formalism of Section 3 that distinguishes the environmental control problem from the robot control problem, this paper offers an explicit expression (15) of the environmental control actions that the Shannon juggler performs. It should be very fruitful to compare this impact schedule with that induced by the mirror law (6). One imagines the construction of various compromise schemes — for example, a scheme that might enlarge the domain of attraction of the 1:1

Shannon fixed point by introducing a little bit of actively acquired information about the environment. However, in the absence of a complete characterization of the Shannon juggler's dynamics it is not at all clear how to think about this.

It should be noted that both the Shannon and mirror strategies lend themselves to implementation on real robots but that neither offers a specification of what the continuous torque policy should be for a general purpose actuator.⁵ One would hope for a more complete encoding prescription, a mapping, $\phi(r, b)$, for which the closed loop system coupled to the environment

$$\begin{aligned} \ddot{\rho} &= \phi(\rho, \dot{\rho}, \beta, \dot{\beta}) \\ \ddot{\beta} &= -\gamma + \delta(\beta - \rho) \cdot C_r(b) \end{aligned} \quad (16)$$

(here, δ is the delta function) generates the entrained relation between robot and environment we desire. Again, this hope seems to founder on the present unavailability of an entrainment theory for strongly coupled oscillators.

In fact, to be honest, the fragility of the analysis offered here lies in embarrassing contrast to the robustness of entrainment phenomena demonstrated by a growing number of workers in the area of dynamical robot dexterity.⁶ We are forced to integrate in closed form the smooth pieces of relaxation oscillators. This has never proven terribly daunting in the one degree of freedom "gedanken" cases of the kind investigated here or, say, in [10]. But in problems involving higher degrees of freedom our ability to integrate exactly deteriorates quickly. Although the dynamics are simple, the need to compute time of flight exactly, for example as in (10), leads to transcendental equations. Appeal to the implicit function theorem permits merely local analysis [1]. A qualitative theory is needed for coupled oscillators in robotics.

Acknowledgements

I acknowledge with pleasure discussions with Avis Cohen, John Guckenheimer, Phil Holmes, Robert MacKay, as well as the many other participants

⁵This statement does not do adequate justice to the pleasing simplicity of Atkeson and Schaal's juggling implementation [14]. In their work, the one degree of freedom actuator is given a fixed torque and the motion profile is achieved by recourse to a fixed linkage. Presumably, a wide range of oscillators could be emulated by affixing a suitable linkage to uniform sinusoidal motion in this manner. One wishes for an analytical perspective offering the means to induce so neatly the same menu of oscillators directly through the actuator's torque inputs.

⁶To say nothing of the truly amazing behaviors of coupled biological oscillators.

in the Special Semester on Spatial and Temporal Dynamics of the Université de Montréal Centre de Recherches Mathématiques whose comments and tutorial efforts have greatly advanced these musings.

Portions of this work were completed during my stay with Dr. Tatsuo Arai of the Japan MITI Mechanical Engineering Laboratory. I am indebted to Dr. Arai for his support, and to MITI for my appointment as Visiting Fellow.

References

- [1] M. Bühler and D. E. Koditschek. From stable to chaotic juggling. In *Proc. IEEE International Conference on Robotics and Automation*, pages 1976–1981, Cincinnati, OH, May 1990.
- [2] M. Bühler, D. E. Koditschek, and P. J. Kindlmann. Planning and control of a juggling robot. *International Journal of Robotics Research*, (to appear), 1993.
- [3] Martin Bühler. *Robotic Tasks with Intermittent Dynamics*. PhD thesis, Yale University, New Haven, CT, May 1990.
- [4] Pierre Collet and Jean-Pierre Eckmann. *Iterated Maps on the Interval as Dynamical Systems*. Birkhäuser, Boston, MA, 1980.
- [5] Bruce Donald and Jim Jennings. Constructive recognizability for task-directed robot programming. In *IEEE Int. Conf. Rob. Aut.*, 1992.
- [6] M. A. Erdmann and M. T. Mason. An exploration of sensorless manipulation. *IEEE J. Rob. and Aut.*, 4(4):635–642, Aug 1988.
- [7] Neil Fenichel. Persistence and smoothness of invariant manifolds for flows. *Indiana U. Math. J.*, 21(3):193–226, 1971.
- [8] John Guckenheimer. Sensitive dependence to initial conditions for one dimensional maps. *Communications in Mathematical Physics*, (70):133–160, 1979.
- [9] D. E. Koditschek. Task encoding: Toward a scientific paradigm for robot planning and control. *Journal of Robotics and Autonomous Systems*, 9:5–39, 1992.
- [10] D. E. Koditschek and M. Bühler. Analysis of a simplified hopping robot. *International Journal of Robotics Research*, 10(6), Dec 1991

- [11] M. Mason. Kicking the sensor habit. *AI Magazine*, 1993.
- [12] D. Miller. A twelve-step program to more efficient robots. *AI Magazine*, 1993.
- [13] A. A. Rizzi and D. E. Koditschek. Further progress in robot juggling: The spatial two-juggle. In *IEEE Int. Conf. Robt. Aut.*, page (to appear), Atlanta, GA, May 1993.
- [14] Stefan Schaal and Christopher G. Atkeson. Open loop stable control strategies for robot juggling. In *IEEE Int. Conf. Rob. Aut.*, pages 3:913-918, Atlanta, GA, May 1993.