In: *IEEE International Conference on Robotics and Automation*, **3**, pp.913-918, Georgia, Atlanta.

Open Loop Stable Control Strategies for Robot Juggling

Stefan Schaal and Christopher G. Atkeson

The Artificial Intelligence Laboratory and the Brain and Cognitive Sciences Department,
Massachusetts Institute of Technology, NE43-771, 545 Technology Square, Cambridge, MA 02139

Abstract: *In a series of case studies out of the field of dynamic manipulation (Mason, 1992), different principles for open loop stable control are introduced and analyzed. This investigation may provide some insight into how open loop control can serve as a useful foundation for closed loop control and, particularly, what to focus on in learning control.*

1 Introduction

This paper explores open loop stable control strategies for a variety of juggling tasks. By control strategy we mean the way a movement system structures itself to approach a task. An *open loop stable* control strategy does not use active reaction to respond to perturbations. It uses the geometry of the mechanical device, the kinematics and dynamics of motion, and the properties of materials to stabilize the task execution. It is distinguished from closed loop control strategies by the absence of sensory input to the computing of actuator commands for error compensation. Some open loop controlled devices use no actuators at all (McGeer, 1990).

As has been shown by McGeer's (1990) passive dynamic walking machines, dynamic systems are likely to offer regions in state and control space which are inherently more advantageous to the execution of a task then others. Analogously, such favorable, although not open loop, control strategies may be found in sports where, for instance, the Fosbury Flop in high jumping lifted the entire discipline to new heights. If strategy largely determines the performance of a movement task and not the details of how the strategy is implemented, exploitation of open loop dynamics may be a promising way to find good task strategies.

Open loop control is interesting for several reasons. It is an important technique for automation and can offer cheaper and faster control for some operations. Vibratory feeders and sorters, and the remote center compliance device for peg in hole insertions are good examples. Open loop stable strategies may be used as a core around which closed loop control is organized. This may make closed loop control more robust by reducing the demands on the feedback controller. Finally, understanding open loop stable strategies may aid in understanding the key features to be learned by a robot practicing a task.

In section 2 several juggling tasks will be discussed to illuminate issues of open loop control and their implications. We believe that searching for open loop stability, or at least something which comes close to that, may help the system to substantially reduce control effort, improve performance and speed up learning. Discussions of the case studies are largely deferred to section 3. Each case study has been explored by building an actual machine.

2 Case Studies of Stable Open Loop Control

2.1 Paddle Juggling

In paddle juggling, a ball (or multiple balls) is kept in the air by hitting the ball vertically with a horizontal paddle (a behavior often exhibited by tennis players waiting for a court). Under visual guidance, this is a closed loop task which has been examined by (Aboaf 1988, Bühler 1990, Rizzi 1992a&b, Ballard 1989, and Toshiba 1989). Without information about the ball state, only open loop control is possible. This task received considerable attention in recent years, for the vibrating paddle version (high paddle oscillation frequency with small amplitude) can be shown to exhibit period bifurcations, strange attractors, and chaos-like motion (Lichtenberg & Lieberman, 1982; Guckenheimer & Holmes, 1983; Moon, 1987; Tufillaro et al., 1992). In the following, open loop stable control strategies for paddle juggling will be explored. The emphasis lies on achieving a constant bouncing height and period; control of the horizontal dimensions will be neglected for the moment. The paddle mass is assumed to be much larger than the ball mass.

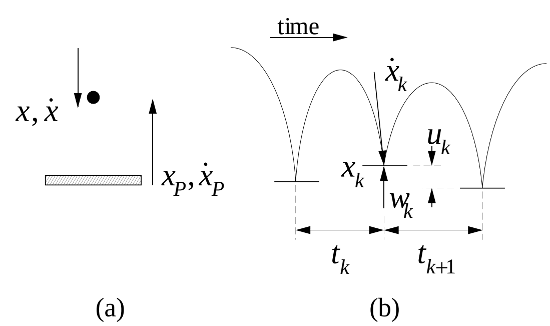

FIGURE 1 (a) sketch of paddle juggling and notation for continuous case; (b) notation for discretized case

As can easily be verified, the discrete equations of motion (using the notation of Figure 1) yield:

$$\begin{aligned}
\dot{x}_{k+1} &= -\sqrt{\left((1+\alpha)w_k - \alpha\,\dot{x}_k\right)^2 - 2gu_k}\,,\\
x_{k+1} &= x_k + u_k\,, \qquad\qquad (1)\\
t_{k+1} &= \frac{1}{g}\left(\left((1+\alpha)w_k - \alpha\,\dot{x}_k\right) + \dot{x}_{k+1}\right).
\end{aligned}$$

$(\dot{x}_k, x_k)$ denote the velocity and the position, respectively, of the ball just before the moment of impact. The velocity of the paddle at this time is (w_k). After the impact, the paddle shifts its position by the distance (u_k) where the next impact (k+1) will take place. Energy loss during the impact is modeled with a coefficient of restitution (α).

Is there an open loop control strategy which would achieve a stable, simple (one impact per cycle) juggling pattern? It turns out that a sinusoidal driving motion, $x_P = A\sin(\omega t + \theta_0)$, as chosen in nearly all open loop studies, suffices to obtain stability. Appropriately relating (w_k, u_k) to the sine motion results in the following condition for stable fixed points of period (τ), where $\tau = 2\pi / \omega$:

$$\theta_0 = \arccos\left[\frac{\pi g}{Aw^2}\left(\frac{1-\alpha}{1+\alpha}\right)\right] \tag{2}$$

This condition was formulated for the phase (θ_0), assuming the motion of the paddle is constant and the ball has to find the impact phase where a stable periodic motion exists; impact velocity is solely determined by the ballistic flight of the ball of duration (τ).

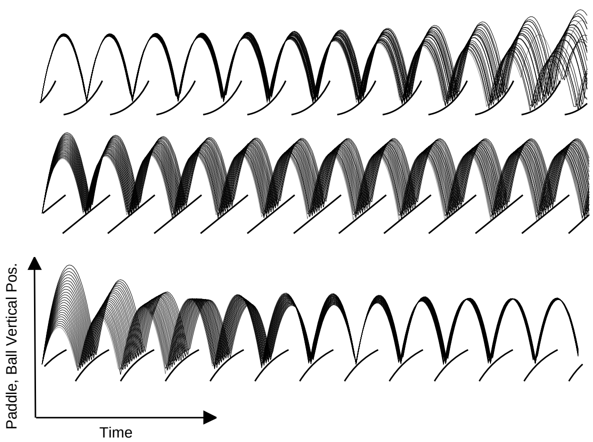

FIGURE 2 Simulation of effect of hit trajectory on open loop stability: (top) positively accelerating hit trajectory; (middle) constant velocity hit trajectory; (bottom) negatively accelerating hit trajectory

The essence of this open loop stability lies in what could be called a focusing hit trajectory (Schaal et al. 1992). As depicted in the simulations of Figure 2, only hit trajectories which are negatively accelerating at impact while the position is still increasing accomplish this stability. The middle row of Figure 2 shows the effect of a constant velocity hit trajectory on a set of ball trajectories with a range of initial velocities. Both the paddle and ball vertical positions are plotted against time. Due to the neutral stability of the constant velocity hit trajectory the trajectories diverge at a rate that is linear in time. The top row of Figure 2 shows the exponential divergence due to a positively accelerating hit trajectory for a set of ball trajectories with a tenth of the range of initial velocities used in the constant velocity case. The bottom row of Figure 2 shows the focusing effect of a negatively accelerating hit trajectory for a wide range of initial ball velocities. Interestingly, in all the literature on closed loop ball juggling this control strategy has not been applied.

Two runs of a simple one-joint robot using a pantograph linkage to maintain a horizontal paddle orientation (Figure 3a) demonstrate the feasibility of this open loop control method (Figure 3b and 3c). A special trampoline-like racket and a ping-pong ball as juggling object even allow open loop stability in the horizontal plane: this racket exerts a restoring force toward the racket center if the ball lands off-center. Data was recorded with a vision system running at 60 Hz; parameters of the paddle movement were $\tau = 0.4\,\text{sec}, A = 0.05\text{m}, \alpha = 0.51$.

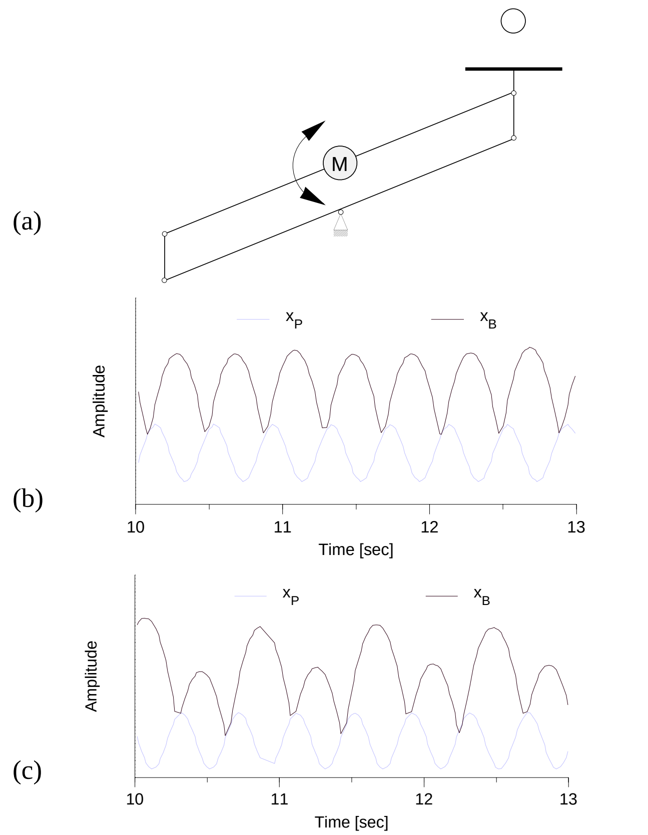

FIGURE 3 (a) setup of juggling robot; (b) period-one juggling motion; (c) period-two juggling motion

It is possible to estimate the size of the basin of attraction of steady paddle juggling. The system has a trapping region for all initial velocities of the ball yielding:

$$\dot{x}_{k=0} \leq \frac{2\sqrt{gA} + (1+\alpha)Aw}{\alpha - 1} \tag{3}$$

The derivation of this bound is similar to Tufillaro et al. (1992) and had to be omitted due to space limitations. Knowing the trapping region, it is sufficient to investigate the basin of attraction only in this region, which is illustrated numerically in Figure 4 for the paddle motion parameters given above. The gray areas denote initial conditions belonging to the basin of attraction for periodic juggling, the white areas physically impossible initial conditions, and the black areas initial conditions which did not lead to periodic juggling (Tufillaro et al., 1992).

The numerical calculation for Figure 4 assesses the size of the basin of attraction as 0.257 of the trapping region, corresponding to a probability P=0.257 that the ball ends up in periodic juggling if it was initially dropped from a large enough height. In the real robot, random effects during bouncing as well as unmodeled parameters (like air resistance for the ping-pong ball) made the basin of attraction significantly larger. Under the given parameter settings, the paddle peak acceleration is greater than gravity so that the ball can never come to rest on the paddle indefinitely. This results in many launches of the ball with some randomization due to mechanical variability in the apparatus, and it turned out that we were unable to avoid getting a periodic juggling pattern (most often period-one, but sometimes period-two or higher).

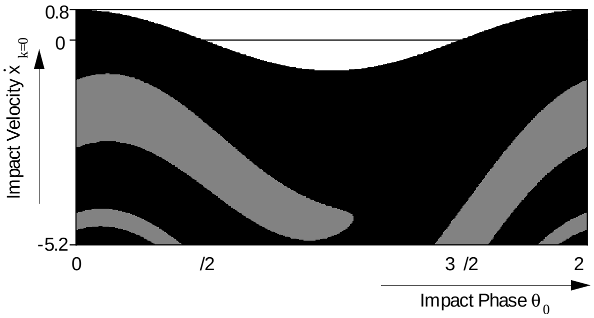

FIGURE 4 Basin of attraction of paddle juggling

The easiest way to paddle juggle more than one ball in phase (so the hits are roughly at the same time) with one actuator is to use an open loop stable hit trajectory. Since the balls fall approximately at the same time it is difficult to manipulate the paddle velocity to actively control each hit. We have used our open loop strategy to paddle juggle three balls simultaneously by attaching three separate paddles to the pantograph mechanism.

So far the paddle has been driven sinusoidally which is a rather arbitrary choice, essentially for mathematical convenience. It would be far more interesting if the system developed its own way of dealing with the juggling task. Schaal & Sternad (1992) examined learning of paddle juggling with genetic algorithms (Holland, 1992). They demonstrated that a learning system can find focusing trajectories just from reinforcing small variance juggling patterns and minimization of jerk $(\ddot{x}_P)$. When feedback of the ball state was also provided, the focusing trajectory was still found, although it emerged out of a coupled oscillator-like representation in which the ball drives the paddle by some linear coupling terms (Bühler, 1990). Particularly in the latter case, the paddle trajectory was clearly not a sine wave anymore.

2.2 Ball-in-a-Wedge Juggling

The next juggling task constitutes a slight extension of the paddle juggler. Instead of juggling the ball vertically with a racket, it is juggled horizontally between two walls, arranged in a V-form (Figure 5). Both walls can pivot together about the tip of the wedge. The angle between the walls is adjustable but fixed during experimental runs, and the two walls move in unison.

Driving the wedge with a sinusoidal motion and trying to rely on the principle of self-focusing trajectories does not suffice to find an open loop control law. From the analysis of the conservative, non-oscillating ball-in-a-wedge, Lethihet & Miller (1986) found a dependence of the system's stability on the wall angle (θ). The result of their Liapunov stability examination is depicted in Figure 6.

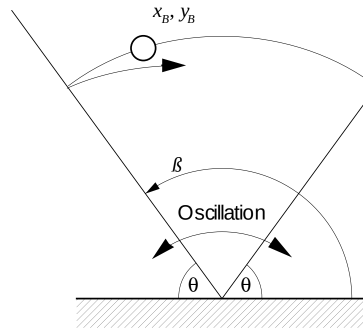

FIGURE 5 Ball-in-a-wedge juggling system

The conservative system is Hamiltonian and, by virtue of its first integral, i.e., the Hamiltonian, its stability can be graphed in this way. Due to energy conservation, maximum stability does not go beyond neutral stability (Liouville's theorem, Symon, 1971). This corresponds to a Liapunov exponent of zero (note: numerical instabilities in calculating the data of Figure 6 also give some incorrect values slightly below zero). Liapunov exponents greater than zero indicate instability (or chaotic motion).

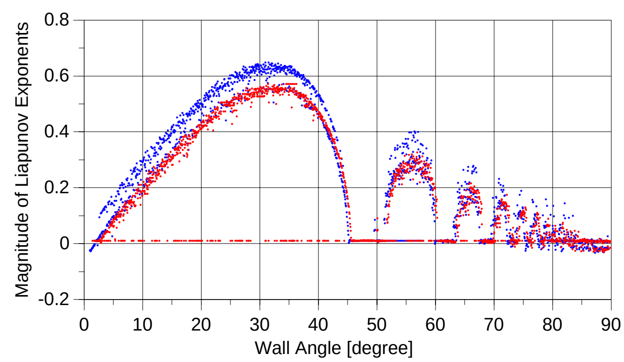

FIGURE 6 Liapunov exponents of ball-in-wedge system as a function of the wall angle (θ)

A linear stability analysis of the oscillating, thus non-autonomous, dissipative system yields similar results. Figure 7 depicts the magnitude of the four eigenvalues as a function of the wall angle (θ) for a coefficient of restitution $\alpha = 0.8$. It should be noted that the stability analysis was based on a discrete version of the system's equations (cf. Lethihet & Miller, 1986). Hence, stable regions are characterized by all eigenvalues in the interval [-1,1]. The basin of attraction could be estimated as in the previous section.

Figure 8 shows a run of a real ball-in-a-wedge machine. From the impact phase, one recognizes again the principle of a focusing trajectory which, as well as the wall angle, is crucial to achieve stable juggling.

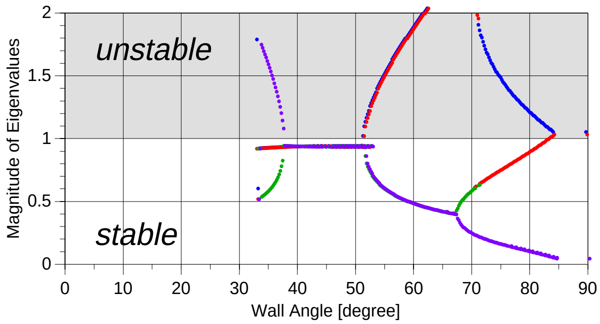

FIGURE 7 Eigenvalue analysis of ball-in-a-wedge as a function of the wall angle (θ) (note: physically implausible regions of (θ) are clipped)

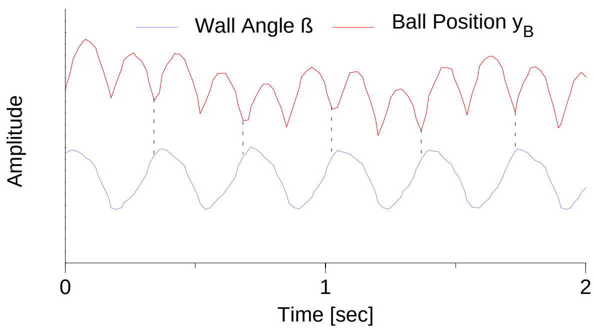

FIGURE 8 Run of the real ball-in-a-wedge system with $\theta = 49°$ using a tennis ball (recorded with a 60Hz vision system) Note: the dashed lines indicate the impact phase of ball and wall on one side of the wedge (for notation see Figure 5)

2.3 Five Ball Juggling

Although the transfer from conservative systems to dissipative nonautonomous systems, as done in the previous section, is not possible in general, it raises the issue of the role of dissipation in dynamic systems. Pushing dissipation to an extreme, the ball-in-a-wedge may be transformed into another juggling task. Setting the coefficient of restitution to 0 and supporting the ball at impact so that it cannot roll away results in regular ball juggling. Ball juggling involves catching, carrying, and releasing a ball on the proper trajectory. After the ball's ballistic flight the catching hand has to be placed appropriately in the ball's trajectory, which may require velocity matching with the incoming ball to reduce the impact. From the moment of contact between ball and hand the control phase begins. During the carrying phase the ball has to be guided to the desired release conditions from where it continues its cycle autonomously. Since the throw of the balls is unlikely to be precisely repeatable, a juggling system must be able to tolerate a variety of incoming ball trajectories and guide the balls from these different catch conditions to the intended release state.

Changing the ball-in-a-wedge as described above, one ball could be thrown back-and-forth between the walls. In order to juggle more than one ball, the balls must travel on distinct trajectories, and they should travel for a rather long time to facilitate the coordination of the other balls. These requirements can be met by letting the balls bounce once on the floor before catching them again. This "bounce juggling" makes catching easier because the balls are caught roughly at the top of their trajectories where the ball velocity is at its lowest.

Several years ago, Claude Shannon found a rather simple solution to a robot bounce juggler. He built a machine essentially consisting of a motor attached at the center of a rod which has two catchers mounted at each end (Figure 9). By driving the motor sinusoidally and adjusting the distance of the catchers, the motor frequency and amplitude, and the height of the setup above the floor, it is possible to find a configuration in which the balls are juggled in a stable fashion. A drum was used to provide an elastic floor. Juggling three balls requires one full oscillation during the flight of a ball. The modified Shannon juggler which was built for this investigation could even juggle five balls for one to two minutes. Here, two full oscillations had to elapse during the flight phase of one ball which demands a much faster driving frequency.

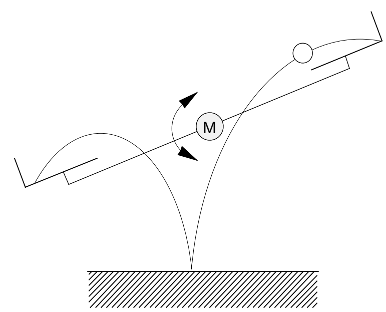

FIGURE 9 Sketch of the Shannon Juggler

When the linkage and the driving frequency are adjusted correctly, the basin of attraction of this juggling is solely determined by the size of the catchers. The catchers are padded with shock absorbing material (bean bags or some special foam) to replace the active catch of human juggling. As long as the balls hit the catchers far enough from the catcher edges, the balls will be guided to the throw location in the catchers by centripetal and gravitational forces. The padding of the catchers implements an inelastic impact. Thus, a repeatable start position for the throw is achieved. Five-ball juggling requires an accurately symmetric driving motion as well as some additional damping of the vibrating drum head. The range of parameter settings for robust five-ball juggling is narrow, while three-ball juggling is much easier to achieve.

2.4 Devil Sticking

By applying the ingredients of the previous sections, more demanding juggling tasks can be investigated. Devil sticking requires manipulating a center stick with two hand sticks by hitting the center stick back and forth between the hand sticks (Figure 10a). A sketch of a devil sticking robot is given in Figure 10b. Three direct drive motors allow positioning of the hand sticks anywhere in space. The hands sticks have passive compliance through the indicated springs which work in parallel with oil dampers. This implements a passive catch. The center stick does not bounce when hitting the hand stick and, therefore, requires an active throwing motion by the robot. To simplify the problem, the center stick was constrained to move on the surface of a sphere by a boom attached perpendicularly to the center stick at its center of mass. For small movements the center stick stays approximately in a plane. The boom also provides a way to measure the current state of the center stick. The juggling robot uses its top two joints to perform planar devil sticking.

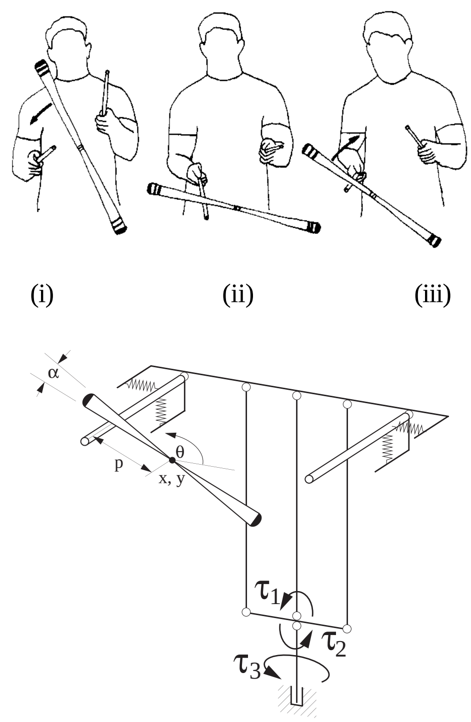

FIGURE 10 (a) Human devil sticking: (i) flight phase, (ii) catch phase, (iii) throw phase; (b) sketch of devil sticking robot; (from van Zyl, 1991)

The task state was defined as the predicted location of the center stick when it hits the hand stick held in a fixed nominal position in absolute space. Standard ballistics equations for the flight of the center stick are used to map flight trajectory measurements $(x(t), y(t), \theta(t))$ into a task state:

$$\mathbf{x} = (p, \theta, \dot{x}, \dot{y}, \dot{\theta}) \tag{4}$$

where (p) is the distance between the hand stick and the center of mass of the center stick, (θ) is the angle of the center stick, and $(\dot{x}, \dot{y}, \dot{\theta})$ is the velocity vector of the center stick. The task command is given by a displacement of the hand stick from the nominal position (x_h, y_h) which specifies where and when the catch occurs, a center stick angular velocity threshold to trigger the start of a throwing motion $(\dot{\theta}_t)$ which specifies how long to hold the center stick, and a throw velocity vector (v_x, v_y) which specifies the characteristics of a fixed distance constant velocity throwing motion.

$$\mathbf{u} = (x_h, y_h, \dot{\theta}_t, v_x, v_y) \tag{5}$$

Only juggling patterns that are symmetric are considered.

The importance of geometry in the ball-in-a-wedge task can be applied to devil sticking. A simplified stability analysis indicated that juggling patterns at steep center stick angles, e.g., $\theta \in [45°, 135°]$ are more stable than those at larger center stick angles, for instance, $\theta \in [15°, 165°]$ (Schaal et al., 1992). The former case achieved clearly better results when implemented with an open loop throwing motion: a fixed throw pattern was triggered through sensing the impact between hand stick and center stick. Figure 11 shows a run for the steep juggling pattern. The shallow pattern accomplished only 4-5 hits (not shown here).

Due to the improved open loop stability, an active LQR controller achieved quite a performance improvement with respect to an LQR controller implementing the less favorable juggling strategy (several hundred hits on average against 60-80 hits on average). The LQR controller was built by calculating a pair of setpoints yielding the left-right symmetry requirement, then iteratively finding the controls which achieve a one-hit juggle from the left setpoint to the right setpoint and vice versa, and finally collecting empirical data around the setpoints to build a linear model and subsequently the LQR controller.

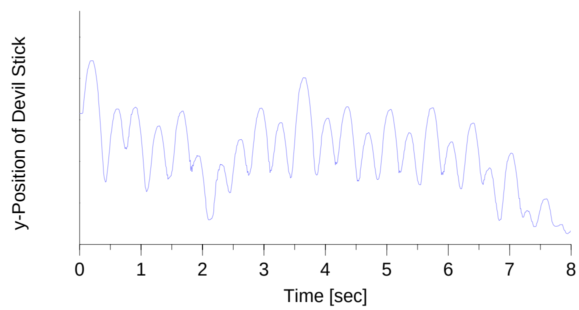

FIGURE 11 Open loop run of devil sticking robot with triggering of fixed throw movement through impact sensing. The graph shows the y coordinate of the center of mass of the devil stick (cf. Figure 10b).

3 Discussion

The examples of section 2 all had a common feature: a large amount of freedom in how to approach the given problem. Instead of forcing a desired trajectory or movement pattern on the system, the design goal was to find solutions exhibiting open loop stability. This reduces computational effort for any control system and uses mechanics to help implement the task (Mason, 1992).

The Shannon juggler literally implements the task mechanically. Aside from relying on clever design and correct tuning of mechanical parameters, its function was mainly accomplished by an energy reset strategy: at the end of each cycle of a ball, the ball's remaining kinetic energy was absorbed, and the ball was re-launched with the desired initial conditions.

Focusing trajectories, as used for the paddle juggler, belong in the class of open loop control design strategies. In the case of the ball-in-a-wedge, a pure trajectory control design strategy is insufficient to accomplish open loop stable control. Additional control parameters are required, i.e., parameters which are kept fixed during the run of a system, but which may be changed between runs. The key control parameter is the geometry of the wedge, namely its angle (θ). Control parameter space usually offers such advantageous regions for a given task, as some arrangement of controls harmonizes more favorably with the system's intrinsic dynamics.

Devil sticking is the most complex jof our uggling examples. It combines two of the preceding strategies: the passive catch of the center stick is an energy reset, the angle of the juggling pattern belongs to geometry variations. Since a focusing trajectory has not been found so far, the system still needs sensory feedback. However, efforts to find open loop-like schemes resulted in clearly improved performance of a closed loop controller.

A crucial property of movement systems is their loss of energy, on which all concepts of open loop stability rely. Liouville's theorem (Symon 1971) shows that conservative systems, i.e., systems which do not lose any energy (as in celestial physics), cannot be asymptotically stable; a change in energy caused by a small perturbation persists. Therefore, dissipation is desirable. At one extreme it can be used as a complete control strategy as demonstrated for the Shannon juggler. At another extreme, a system which cannot dissipate its energy fast enough will need extra control effort to be steered appropriately. Systems with energy dissipation must have some source of energy to avoid running down. McGeer (1990) used an inclined walking surface to drive passive dynamic walking. We have used open loop motion of paddles or hands to drive the system. Finding the right compromise between energy efficiency and stabilization due to energy loss is a topic for future research.

Open loop stable control strategies may be a key in facilitating control system design and, therefore, learning control. A system with improved open loop performance is less dependent on the sampling rate and the accuracy of its sensor readings. It becomes easier to control and more robust. Being open loop stable in certain regions, the system may be able to decrease attention during the open loop phases and use its computational power elsewhere. As we demonstrated with the devil sticking robot, good open loop and feedback strategies are complementary. To find open loop stability becomes equalivalent to the detection of relevant parameters of a task. This may be especially helpful in high dimensional problems.

All dynamic tasks introduced in this paper were solved by special purpose systems. If one were asked to solve such tasks with more complex and general mechanical systems such as multi degree of freedom robot arms, then it might be possible to derive the same open loop strategies by choosing appropriate constraints on the general system. If the complex system confined itself to mimicking the simple machine, redundant degrees of freedom might be resolved more easily. A major topic for research is how these general purpose devices can learn to constrain themselves appropriately.

4 Acknowledgments

Support was provided under Air Force Office of Scientific Research grant AFOSR-89-0500 and by Siemens Corporation. Support for the first author was provided by the German Scholarship Foundation and the Alexander von Humboldt Foundation. Support for CGA was provided by a National Science Foundation Presidential Young Investigator Award, an Alfred P. Sloan Research Fellowship, and the W. M. Keck Foundation Associate Professorship in Biomedical Engineering.

5 References

Aboaf, E. W., 1988: "Task-Level Robot Learning", MS Thesis, MIT Artificial Intelligence Laboratory.

Aboaf, E. W., C. G. Atkeson, and D. J. Reinkensmeyer: "Task-Level Robot Learning", *Proceedings, IEEE International Conference on Robotics and Automation*, April 24-29, 1988, Philadelphia, Pennsylvania. See AI Memo 1006, "Task-level Robot Learning: Ball Throwing", MIT Artificial Intelligence Laboratory for a more complete description of this work.

Aboaf, E. W., S. M. Drucker, and C. G. Atkeson: "Task-Level Robot Learning: Juggling a Tennis Ball More Accurately", *Proceedings, IEEE International Conference on Robotics and Automation*, May 14-19, 1989, Scottsdale, Arizona.

Ballard, D., 1989: Personal communication.

Bühler, M., 1990: "Robotic Tasks with Intermittent Dynamics", Ph.D. Thesis, Yale University Electrical Engineering, New Haven, CT.

Guckenheimer, J.; Holmes, P., 1983: *"Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*", Springer.

Lehtihet, H. E., Miller, B. N. , 1986: "Numerical Study of a Billiard in a Gravitational Field", *Physica 21 D*, 93-104.

Lichtenberg, A.J.; Lieberman, M.A., 1982: "*Regular and Stochastic Motion*", New York: Springer.

Mason, M., 1992. Personal communication.

McGeer, T., 1990. "Passive Walking with Knees," *Proceedings, IEEE International Conference on Robotics and Automation*, May 13-18, 1989, Cincinnati, Ohio.

Moon, F.C., 1987: "*Chaotic Vibrations : An Introduction for Applied Scientists and Engineers*", New York: Wiley (1987).

Rizzi, A. A.; Whitcomb, L.L.; Koditschek, D.E., 1992a: "Distributed Real Time Control of a Spatial Robot Juggler", *IEEE Computer*, 25(5).

Rizzi, A. A.; Koditschek, D.E., 1992b: "Progress in Spatial Robot Juggling", *Proceedings of IEEE International Conference on Robotics and Automation*, Nice, France.

Schaal, S.; Atkeson, C.G.; Botros, S., 1992: "What Should Be Learned?", *Proceedings of Seventh Yale Workshop on Adaptive and Learning Systems*, pp.199-204.

Schaal, S.; Sternad, D. ,1992: Strategies in Motor Learning and Motor Coordination, To appear in: "*SFI Lectures on Complex Systems 1992*", Santa Fe, NM: Santa Fe Institute.

Symon, K. R., 1971: *Mechanics*, Addison-Wesley Reading, Mass.

van Zyl, G., 1991: "Controlling a Devil Sticking Robot", Master Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Toshiba, 1989: Personal communication.

Tufillaro, N.B.; Abbott, T.; Reilly, J., 1992: "*An Experimental Approach to Nonlinear Dynamics and Chaos*", Addison-Wesly.