51st IEEE Conference on Decision and Control
December 10-13, 2012. Maui, Hawaii, USA

Control of Nonlinear Systems with Symmetries using Chaos

Philipp Reist and Raffaello D'Andrea

Abstract— **We present a method that exploits chaos for the control of systems composed of subsystems with identical nonlinear dynamics and a shared, common control input. Due to symmetry, these systems are uncontrollable in a deterministic sense. However, the systems may be controllable in a stochastic sense when they are driven by process noise. We present a control strategy that exploits the sensitivity of chaotic motion to process noise. The chaotic subsystem trajectories evolve independently on the strange attractor, which enlarges the reachable set of the system. Specifically, we consider the control of a juggling machine bouncing multiple balls on a single, actuated paddle. The goal is to control the balls to a combination of stable periodic orbits. First, a paddle motion is applied that induces chaotic ball trajectories. Then, when the ball states reach the basins of attraction of the desired periodic orbits, the paddle motion is switched to the motion that stabilizes the orbits. Both simulation and preliminary experimental results are presented.**

I. Introduction

We consider the control of dynamical systems composed of N subsystems with identical nonlinear dynamics

$$\dot{x}_i(t) := \frac{d}{dt}x_i(t) = f\left(t, x_i(t), U(t)\right), \ i = 1, \ldots, N \quad (1)$$

where x_i is the subsystem state, and U is the common control input. The system state X is the collection of the subsystem states $X(t) = \{x_1(t), \ldots, x_N(t)\}$. These systems are uncontrollable due to symmetry: If the initial states of the subsystems are identical, $x_i(0) = x_j(0)$, for all i, j, they remain identical, no matter what control is applied: $x_i(t) = x_j(t)$, for all i, j, t. Practical examples of such systems include simultaneous manipulation of multiple parts on vibrating surfaces [1], [2], or teams of micro robots controlled by a magnetic field [3].

The key to control is the observation that any physical realization of the system (1) is more accurately captured by

$$\dot{x}_i(t) = f\left(t, x_i(t), U(t), d_i(t)\right), \ i = 1, \ldots, N \quad (2)$$

where d_i is some process noise signal. Given identical initial conditions of the subsystems, the state trajectories are no longer identical due to the individual noise signals. Roughly speaking, if the subsystems are controllable with respect to the noise signals $d_i(t)$, then the overall system is stochastically controllable: The noise is able to drive the system to a neighborhood of any state X with a probability larger than zero. For precise definitions of stochastic controllability, we refer to the work in [4], [5]. In practice, the control problem may be to achieve efficient transitions of the system state X between the basins of attraction of different equilibria of the system. The efficiency may be measured, for example, by the probability of achieving a transition, or by the expected transition time.

This work was supported by ETH Research Grant ETH-31 11-1. The authors are with the Institute for Dynamic Systems and Control at ETH Zurich, Sonneggstrasse 3, CH-8092 Zurich, Switzerland. Supporting material, including videos, may be found at www.blindjuggler.org.

Fig. 1. The Blind Juggler bouncing four balls on the Cloverleaf paddle. The user interface features two buttons to control the juggling machine.

We present a control strategy for a specific realization of a system described by (2): a juggling machine that bounces multiple balls on a single, actuated paddle, see Fig. 1. The paddle is called the Cloverleaf for its shape and features four concave, parabolic areas that keep the balls from falling off the paddle without feedback. The machine can continuously juggle balls at heights of up to 2 m with a specific periodic juggling motion of the paddle. This paddle motion provides local stability to a vertical, periodic ball trajectory that impacts once per paddle stroke (the ball trajectory you would generate juggling a table tennis ball on a paddle). A detailed stability analysis and more details about the machine may be found in [6], and videos are available in [7].

Bouncing balls exhibit rich dynamical behavior, which has been well studied, see, for example, [8], [9]. One such behavior is that the ball trajectories are chaotic for a specific higher frequency paddle motion. Furthermore, we find multiple distinct, locally stable ball trajectories for the juggling motion of the paddle. These stable trajectories may be combined into *juggling patterns* on the Cloverleaf paddle.

Consider the following control problem: Find a paddle motion that achieves a transition of the balls to a desired juggling pattern, given an initial ball configuration. In the deterministic sense, the balls are uncontrollable: Given identical initial conditions, the ball states remain identical, no matter what paddle motion is applied. In the experimental setup, however, there are disturbances acting on the balls. For example, small breezes of air or random fluctuations of the bounciness of the balls. In addition, we may exploit the sensitivity of chaotic ball trajectories to small disturbances in order to further enlarge the reachable set of the system.

The control strategy is: Induce chaotic ball motions and estimate the ball states as they independently evolve along the strange attractor. When the ball states reach the basin of attraction of the desired juggling pattern, switch to the juggling motion of the paddle. The balls transition to the locally stable juggling pattern and the task is achieved. Preliminary simulation and experimental results provide an empirical proof of concept for the proposed strategy, but also highlight its limitations. For example, relying on a stochastic process to get the state to the basin of attraction may result in long transition times. Including feedback in this transition process may improve the expected transition time and is a rich control problem. We discuss limitations and possible future work in Sections III and IV.

The paper is structured as follows: We first discuss related work in Section I-A, and then introduce the bouncing ball system, stable periodic orbits, and chaos in Section II. We present the proposed control strategy and preliminary simulation and experimental results in Section III and conclude with a discussion in Section IV.

A. Related Work

The control strategy we propose is related to the work in [10], where a single ball is controlled such that it bounces at a desired apex height with a similar approach using chaos. This approach is related to the field of chaos control, introduced in [11], where the OGY method is presented: An unstable periodic orbit of the strange attractor of a system is stabilized using feedback. When the state leaves a predetermined stabilizable region around the orbit for some reason, feedback is switched off, and the system evolves open-loop and chaotically. When the system state re-enters the stabilizable area, feedback is switched back on.

The presented control strategy has a small set of control actions: To switch or not to switch. The concept of achieving tasks by exploiting the natural dynamics of a system with intermittent, limited control is related to the idea of *self-organization* introduced in [12], where different open-loop stable behaviors of particles that bounce between a fixed boundary and an actuated plate are studied.

Lastly, there exists a large body of work on the control of juggling robots. A few examples include the work in [13]–[15]. In [13], a cascading control design for a robot arm juggling a ball in the presence of obstacles is discussed. In [14], minimal feedback from microphones for juggling a ball in an actuated wedge is studied. More recently, a model predictive control strategy for a paddle juggling robot was presented in [15].

II. Dynamics of the Bouncing Ball

We focus on the vertical ball dynamics and assume that the horizontal dynamics of the ball are decoupled and stable. We showed in [6] that these assumptions hold for first-order ball dynamics. The dynamics are modeled by a straightforward hybrid system. Assuming instantaneous impacts, there is only one discrete state that represents the flight phase. We omit this single discrete state in the following. The ball state is $x = (z, \dot{z})$, where z is the ball height and $\dot{z}$ is the ball velocity. During free fall, the ball dynamics are

$$\ddot{z}(t) := \frac{d^2}{dt^2} z(t) = -g \tag{3}$$

where we ignore aerodynamic drag and $g = 9.81\,\mathrm{m\,s^{-2}}$ is the gravitational acceleration. The state of the actuated paddle is $P = (p, \dot{p})$, with p the paddle height and $\dot{p}$ the paddle velocity. The ball is in contact with the paddle if $z = p$ (the ball radius can be ignored, as it is a constant offset in z). The guard condition captures impact events and is $(z(t) - p(t) \leq 0) \wedge (\dot{z}(t) - \dot{p}(t) \leq 0)$, where $\wedge$ denotes a logical *and*. An impact occurs if the ball and paddle position coincide and the ball is not moving away from the paddle. The reset map describes the discrete impact dynamics

$$z^+ = z^- \tag{4}$$

$$\dot{z}^+ = -e_z \dot{z}^- + (1 + e_z)\dot{p}. \tag{5}$$

where we apply Newton's impact law. The superscripts $-, +$ denote pre- and post-impact and $e_z \in [0, 1]$ is the coefficient of restitution of the ball, which models energy loss at impact. In (5), we further assume that the impact does not change the momentum of the paddle, i.e. $\dot{p}^+ = \dot{p}^-$, which is reasonable since the ball mass is much smaller than the paddle mass.

We denote an execution of the hybrid system with

$$x(t) = \phi\left(t, x^0, P(t)\right) \tag{6}$$

where the initial condition is defined as

$$\phi\left(0, x^0, P(0)\right) := x^0 = \left(z^0, \dot{z}^0\right). \tag{7}$$

The set of admissible initial conditions x^0 is

$$\mathcal{I} := \left\{ x^0 \middle| \left(z^0 \geq p(0)\right) \wedge \left(g z^0 + \frac{1}{2}(\dot{z}^0)^2 \leq g H_{\max}\right) \right\}. \tag{8}$$

The ball must start above or on the paddle, and the initial apex height of the ball cannot be above the maximal juggling height $H_{\max} = 2.0\,\mathrm{m}$, above which the ball falls off the paddle [6].

A. Periodic Orbits

The periodic juggling motion of the paddle is designed to stabilize a ball bouncing at an apex height of H_1, such that the ball impacts once per paddle stroke. The corresponding ball motion is called the $\mathcal{P}_1^1$-orbit, where we adopt the nomenclature from [9]. The orbit $\mathcal{P}_l^k$ denotes a periodic ball motion where the ball impacts with the paddle l times per

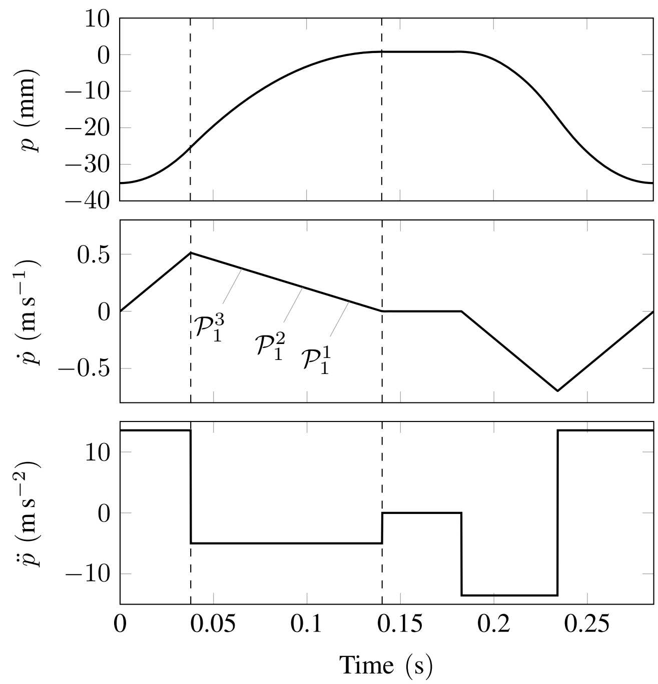

Fig. 2. Juggling paddle motion $P_J(t)$. In order for a $\mathcal{P}_1^k$-orbit to be stable, its nominal impact must occur during the interval where the paddle decelerates with $\ddot{p} = -g/2$, i.e. the interval between the two dashed, vertical lines. The nominal impact times and corresponding paddle velocities of the existing stable $\mathcal{P}_1^k$-orbits are marked in the paddle velocity plot. The required nominal paddle velocity for a $\mathcal{P}_1^4$-orbit to exist is $0.58\,\mathrm{m\,s^{-1}}$, which the paddle never attains.

k paddle strokes. The period time of the orbit is therefore kT_J, where the paddle period is $T_J = 2\sqrt{2H_1/g}$. We show the paddle motion for $H_1 = 0.1\,\mathrm{m}$ in Fig. 2 and denote this motion in the following with $P_J(t)$.

B. Existence and Stability of $\mathcal{P}_1^k$-Orbits

We analyze what locally stable $\mathcal{P}_1^k$-orbits exist for $P_J(t)$ shown in Fig. 2. It may be that there exist other locally stable $\mathcal{P}_l^k$-orbits with $l > 1$, however, for the purpose of introducing the proposed control method, we restrict the set of orbits to the $\mathcal{P}_1^k$-orbits. The ball impacts once per k paddle strokes on a $\mathcal{P}_1^k$-orbit. The apex height, measured from the impact height, of a $\mathcal{P}_1^k$-orbit is

$$H_k = k^2 H_1 \tag{9}$$

where H_1 is the apex height of the $\mathcal{P}_1^1$-orbit. Furthermore,

$$\dot{\bar{p}}_k = \sqrt{2H_k g}\,\frac{1 - e_z}{1 + e_z} \tag{10}$$

is the nominal paddle velocity at impact required for sustained juggling of a $\mathcal{P}_1^k$-orbit given (5). Therefore, the criteria for the existence of a $\mathcal{P}_1^k$-orbit are: 1) the paddle velocity $\dot{p}$ is at least $\dot{\bar{p}}_k$ at any time during $P_J(t)$; 2) the apex height H_k must be less than $2\,\mathrm{m}$, because any larger apex height would cause the ball to bounce off the paddle. In the following, we use the empirically determined coefficient of restitution as a function of the relative impact velocity of the ball: $e_z(\dot{p} - \dot{z}^-) = 0.90 - 0.015(\dot{p} - \dot{z}^-)$. We find that for $P_J(t)$, there exist only $\mathcal{P}_1^1$, $\mathcal{P}_1^2$, and $\mathcal{P}_1^3$ orbits. For $k > 3$, the paddle never reaches the required nominal velocities $\dot{\bar{p}}_k$, see Fig. 2.

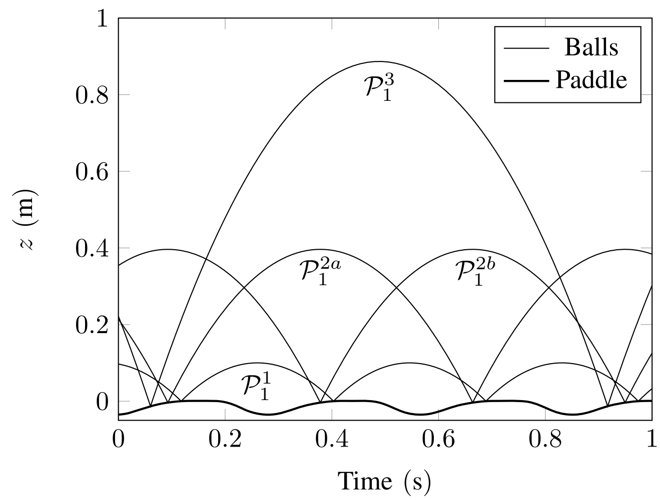

Fig. 3. All existing, stable periodic $\mathcal{P}_1^k$-orbits for the juggling paddle motion $P_J(t)$ shown in Fig. 2. The nominal height of the $\mathcal{P}_1^1$-orbit is $H_1 = 0.1\,\mathrm{m}$ and the nominal impact height of this orbit is at $0\,\mathrm{m}$. The phase-shifted versions of the $\mathcal{P}_1^3$-orbit are omitted for clarity.

We may now analyze the local stability of the $\mathcal{P}_1^k$-orbits that exist. In [6], we showed open-loop, local stability of the $\mathcal{P}_1^1$-orbit when the paddle is decelerating during impact, see Fig. 2. Furthermore, we showed that this stability is independent of the apex height of the ball. Therefore, the local stability of a $\mathcal{P}_1^k$-orbit directly follows from the local stability of the $\mathcal{P}_1^1$-orbit, given that the nominal impact of the $\mathcal{P}_1^k$-orbit occurs during the decelerating phase of $P_J(t)$. This is indeed the case for the $\mathcal{P}_1^1$, $\mathcal{P}_1^2$, and $\mathcal{P}_1^3$ orbits, and we marked the corresponding nominal impact times in Fig. 2.

For $k > 1$, balls skip at least one paddle period between impacts, and we can further distinguish the orbits by their phase. For example, there are two distinct ball motions for the $\mathcal{P}_1^2$-orbit, where the impacts with the paddle are at alternate strokes. This distinction is denoted with letters: $\mathcal{P}_1^{2a}$ and $\mathcal{P}_1^{2b}$, see Fig. 3. Analogously, the $\mathcal{P}_1^3$-orbit results in three distinct motions. Therefore we find a total of 6 distinct locally stable $\mathcal{P}_1^k$-orbits for $P_J(t)$.

C. Chaos

We induce chaotic ball motions with the periodic paddle acceleration profile

$$\ddot{p}(t) = \begin{cases} a_C, & 0 \le s < T_C/2 \\ -a_C, & T_C/2 \le s < T_C \end{cases} \tag{11}$$

where a_C is an acceleration constant and

$$s = (t + T_C/4) \bmod T_C \tag{12}$$

is shifted such that the lowest paddle heights occur at $t_i = iT_C$, for any integer i. The initial condition is $P_C(0) = P_J(0)$, and the resulting periodic paddle motion is denoted with $P_C(t)$. The period time T_C is obtained from

$$A_C = \frac{1}{2} a_C \left(\frac{T_C}{4}\right)^2 \tag{13}$$

where A_C is the paddle position amplitude. In Fig. 4, we show a sampled strange attractor of the system that we

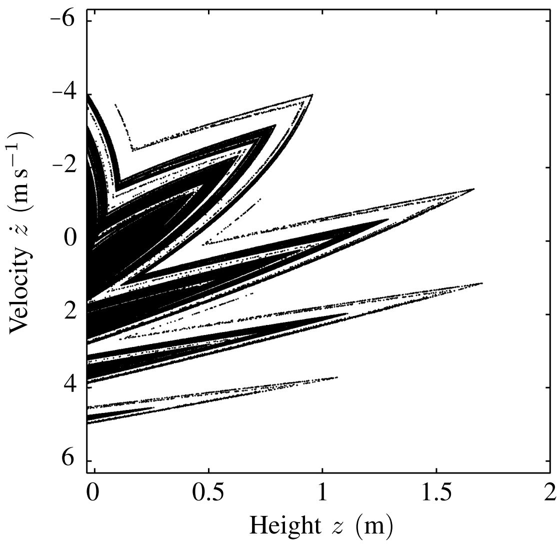

Fig. 4. Strange attractor of the chaotic ball motions for $a_C = 14\,\mathrm{m\,s^{-1}}$ and $A_C = 0.03\,\mathrm{m}$. The ball states are sampled at t_i, when the paddle is at the lowest height in $P_C(t)$.

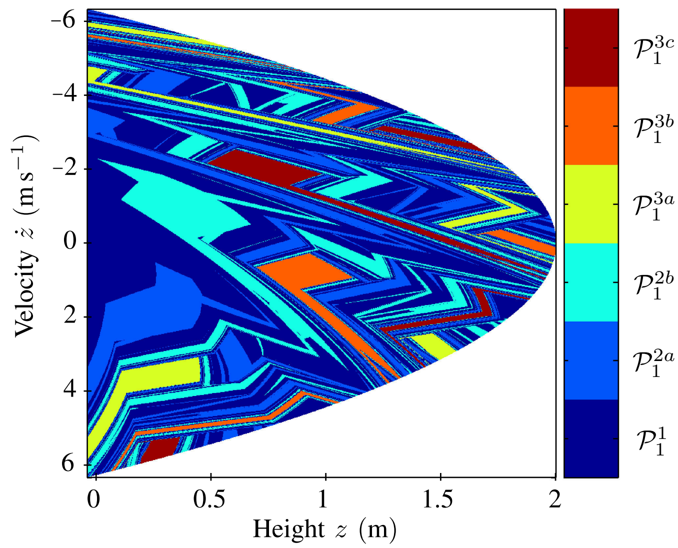

Fig. 5. Sampled basins of attraction. The color predicts the resulting orbit given the respective ball state at the beginning of $P_J(t)$ shown in Fig. 2.

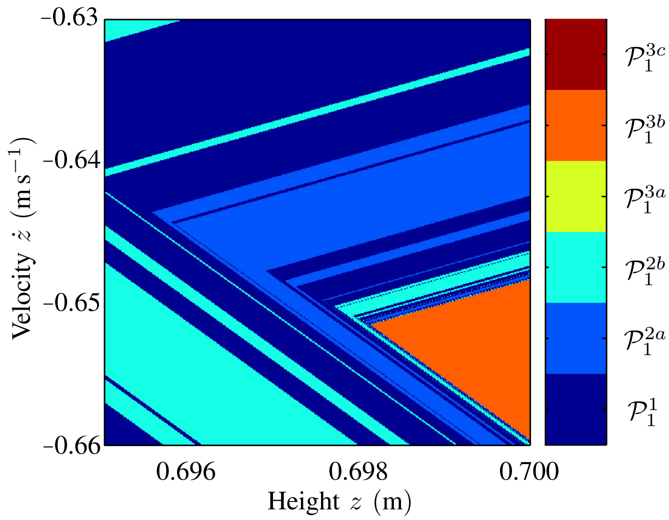

Fig. 6. High-resolution detail of a section of the basin of attraction plot in Fig. 5, showing the fractal-like structure of the basin edges.

obtained by simulation with $a_C = 14\,\mathrm{m/s^2}$ and $A_C = 0.03\,\mathrm{m}$. The simulation is without process noise. The specific choice of parameters is discussed in Section III-C. The ball states are sampled at t_i, when the paddle is at its lowest height in $P_C(t)$. The key property of the chaotic ball trajectories is the sensitive dependence on initial conditions and process noise. This property causes the motions of the balls to quickly diverge and become independent, which is key to the control strategy we introduce below.

III. Controlling Juggling Patterns

We call a combination of stable periodic orbits a juggling pattern. The control problem is how to achieve a given juggling pattern from any initial configuration of the balls. A straightforward control strategy that achieves this task is: Induce chaotic ball motions with the paddle motion $P_C(t)$ and switch to $P_J(t)$ if the balls are predicted to converge to the desired juggling pattern after switching.

A. Basins of Attraction

We limit the candidate switching times to t_i, where the paddle reaches the lowest height in $P_C(t)$. This allows straightforward, smooth transitions from $P_C(t)$ to $P_J(t)$. In order to predict the orbit that a ball transitions to, we approximate the basins of attraction of the orbits. The basin of attraction of an orbit $\mathcal{P}_l^k$ is defined as the set of initial conditions that cause an execution to converge to the orbit when $P_J(t)$ is applied:

$$\mathcal{B}\left(\mathcal{P}_l^k\right) := \left\{ x^0 \;\middle|\; \lim_{t\to+\infty} \phi\left(t, x^0, P_J(t)\right) \to \mathcal{P}_l^k \right\}. \qquad (14)$$

The basins are determined by simulation: We grid the set of initial conditions $\mathcal{I}$ (8) and numerically evaluate the executions $\phi\left(t, x^0, P_J(t)\right)$, checking for convergence to a periodic orbit. The simulations are without process noise. The resulting basins of attraction are shown in Fig. 5. Even though we did not limit the convergence check to $\mathcal{P}_1^k$-orbits, we have not found any other orbits.

The structure and shape of the basins are not straightforward; in fact, we find fractal-like structure at the edges of some basins, see the enlarged area shown in Fig. 6. In order to avoid interpolation issues, we approximate the basins by unions of convex polygons, which also allows us to efficiently check if a ball is within a basin. We manually performed these approximations and chose conservative approximations with the purpose of avoiding the complex structure at the edges of the basins.

B. Control Strategy

The control strategy for achieving a juggling pattern where ball j of N balls converges to orbit $\mathcal{P}_{l_j}^{k_j}$ given initial condition x_j^0 is:

1) Apply paddle motion $P_C(t)$ and estimate ball states at candidate switching times t_i.

2) If the condition

$$\phi\left(t_i, x_j^0, P_C(t_i)\right) \in \mathcal{B}\left(\mathcal{P}_{l_j}^{k_j}\right), \forall j \in \{1, \ldots, N\} \quad (15)$$

is fulfilled, switch the paddle motion to $P_J(t)$. When all balls are within the respective basins of attraction, switch to $P_J(t)$ that stabilizes the orbits.

In practice, we use a straightforward improvement to the algorithm: Instead of immediately switching from $P_C(t)$ to $P_J(t)$, we introduce a variable waiting time $\tau \in [0, T_C)$, where the paddle is at rest, $P_W(\tau - t_i) = P_J(0)$, before the juggling motion starts. This adds the set of states

$$\phi\left(t_i + \tau, x_j^0, P_W(\tau - t_i)\right), \ \tau \in [0, T_C) \quad (16)$$

to check in (15). If we find a τ for which the switching condition (15) holds, we start $P_J(t)$ after the paddle is at rest for τ. In practice, we discretize the interval $[0, T_C)$ and check the ball states at each time step. A simulation result illustrating the strategy is shown in Fig. 7.

C. Simulation Results

The control strategy is evaluated in simulation for two balls and a juggling pattern that combines a $\mathcal{P}_1^{2a}$ and a $\mathcal{P}_1^{2b}$ orbit, where the balls reach an apex height of $0.4\,\mathrm{m}$ and impact with the paddle at alternating paddle strokes. Furthermore, we do not require that the balls reach the specific orbits, but just the specific combination, i.e. Ball 1 may be in $\mathcal{P}_1^{2a}$ and Ball 2 in $\mathcal{P}_1^{2b}$, or vice versa. This is motivated by the fact that on the experimental platform, a spectator could not tell the absolute phase of the ball motions.

The simulation of the chaotic ball motions includes process noise that is introduced at impact

$$\dot{z}^+ = -e_z \dot{z}^- + (1 + e_z)\dot{p} + d \quad (17)$$

where d is assumed to be white noise, i.e. independent between impacts, and is drawn from an empirically determined distribution. For the following experiments, the initial condition for the balls is at rest at a height of $0.1\,\mathrm{m}$, which is a realistic initial condition for the experimental system after juggling up the balls from resting on the paddle. The paddle motion is set to $P_C(t)$, causing chaotic ball motions, and we switch based on the strategy described earlier. We record the time until the paddle juggling motion $P_J(t)$ is started. In 500 experiments, we measured an average switch time of $t_{\text{avg}} = 6.8\,\mathrm{s}$ and a maximal switch time of $35.9\,\mathrm{s}$. Assuming that the event of switching is independent between candidate switching times $t_i = iT_C$, we may calculate the expected switch time

$$\mathrm{E}[t_i] = \sum_{i=1}^{\infty} iT_C(1-q)^{i-1}q = \frac{T_C}{q} \quad (18)$$

where q is the probability of switching at t_i. Given the measured mean switch time t_{avg} and $T_C = 0.26\,\mathrm{s}$, we find that $q \approx T_C/t_{\text{avg}} = 3.8\%$.

The straightforward improvement with variable wait times τ reduces the average switch time roughly by a factor of three. The average switch time for a strategy where only $\tau = 0$ is allowed is $t^0_{\text{avg}} = 18.9\,\mathrm{s}$, and the maximal switch time we measured was $115\,\mathrm{s}$. Another observation is that the variable wait time decreases the dependency of the average switch time on the choice of attractor parameters. In Table I, we state the measured average switch times in 500 experiments each, given the acceleration a_C of $P_C(t)$. The amplitude was left fixed at $A_C = 0.03\,\mathrm{m}$. Since the average switch times are similar for the range of accelerations, we chose $a_C = 14\,\mathrm{m\,s^{-2}}$ as it results in a manageable thermal load on the motor in the experimental setup if $P_C(t)$ is applied for a longer time. For a different juggling pattern, however, the transition times might vary significantly. Furthermore, the set of paddle motions described by a constant acceleration a_C and amplitude A_C is limited. There may be other paddle motions that generate attractors that result in lower transition times. In future work, we will explore the concept of shaping an attractor, which may include feedback control, see for example *targeting* methods described in the discussion. A related challenge is the control of systems where certain regions of the state space must be avoided in the chaotic transitions (for example to avoid destruction, see the discussion of chaotic control of an experimental double pendulum in [10]).

TABLE I

AVERAGE SWITCH TIMES GIVEN ATTRACTOR ACCELERATION

a_C ($\mathrm{m\,s^{-2}}$)	13	14	15	16	17	18	19
t_{avg} (s)	6.9	6.8	6.6	7.2	6.5	7.3	6.6
t^0_{avg} (s)	17.7	18.9	20.9	20.1	21.5	27.5	25.7

D. Experimental Results

We evaluated the strategy with the experimental setup shown in Fig. 1. The goal was to achieve the same two-ball juggling pattern as described in the simulation results. The ball states at candidate switching times are estimated from ball impact times, and the paddle velocities and positions at impact. The impact times are measured with two microphones and the paddle position and velocity are obtained from the motor encoder data.

Because the available experimental data is limited, we provide only qualitative results that highlight the limitations of the proposed control strategy. In experiments, the strategy succeeds in generating the desired juggling pattern, but the observed performance is considerably lower than in simulation. First of all, due to current hardware limitations, the average switch time is higher because the microphones do not detect every impact. To ensure that the decision to switch is based on good data, we only switch when the ball state estimate is from a consistent series of measured impact times. Consistency is checked by comparing the last measured impact time to the impact time predicted by the previously estimated ball state. This reduces the number of candidate switching times and explains the longer observed switch times. Furthermore, the strategy currently fails to achieve the desired juggling pattern in roughly 50% of trials. There are two main reasons for this lower performance: 1) The basins obtained in simulation will differ from the basins of the real system due to modeling errors; and 2) The ball state estimates are not perfect. A video of two successful experiments is available in [7].

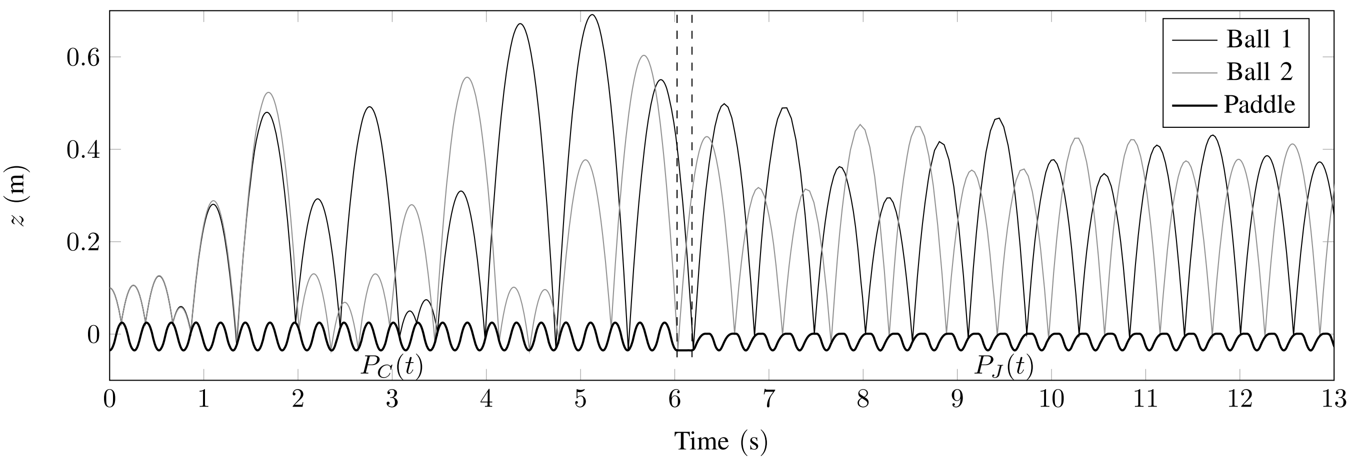

Fig. 7. Simulation result of control strategy. Two balls start at rest at an initial height of $0.1\,\mathrm{m}$. The paddle motion $P_C(t)$ and process noise cause the ball trajectories to quickly diverge. The first vertical, dashed line marks the candidate switch time t_{23} for which a τ was found such that the ball states at $t_{23}+\tau$, i.e. at the second dashed line, are in the basins of a $\mathcal{P}_1^{2a}$ and a $\mathcal{P}_1^{2b}$-orbit, respectively. After the paddle is at rest for τ, the orbit stabilizing paddle motion $P_J(t)$ is started and the balls transition to the desired juggling pattern.

IV. Discussion

The potential for spending long time periods in chaotic transitions can be reduced using control. In future work, we may use feedback to apply a paddle motion that directs the balls more quickly to the basins of attraction. This is a rich control problem, especially given the physical limitations of the paddle motion (if the paddle could be accelerated arbitrarily fast, the control problem would be trivial once the ball trajectories are separated). A related approach is *targeting*, introduced in [16]: The system state may be directed towards a target region on the strange attractor using feedback control. Since the evolution of a chaotic trajectory is sensitive to small perturbations, arbitrary states on the attractor may be reached quickly by applying small control inputs. The challenge is to implement such a strategy for multiple balls, and in the presence of measurement noise, since a targeting policy will be sensitive to small measurement errors.

The experimental results highlight the importance of extending the approach to incorporate data collected with the real system. Any model-based approximation of the basins will differ from the basins of the real system. The challenge is to adjust the approximations based on measured data (for example the fact that an experiment failed) or online estimates of the coefficients of restitution of the balls. The basins may also be evaluated online: simulate the system with estimated system parameters and check if more conservative approximations of the basins may be reached. Conservative approximations may be efficiently obtained with tools such as sum of squares programming [17]. Directly using conservative approximations leads to long transition times, and there is a trade off between aggressive and conservative switching policies: An aggressive policy may attempt a transition multiple times in the time it takes the system to reach conservative basins. Lastly, we can address imperfect state estimates with a probabilistic switching policy: Given measured process noise statistics, the distributions of the ball states can be approximated. Switching is then based on the probability of the states being in the basins.

V. Acknowledgments

This work was supported by ETH Research Grant ETH-31 11-1. The authors would like to thank Marc-André Corzillius, Alex Braun, and Matthew Donovan for their great help with the experimental setup.

References

[1] P. J. Swanson, R. R. Burridge, and D. E. Koditschek, "Global Asymptotic Stability of a Passive Juggler : A Parts Feeding Strategy," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 1983.

[2] T. H. Vose, P. Umbanhowar, and K. M. Lynch, "Toward the Set of Frictional Velocity Fields Generable by 6-Degree-of-Freedom Oscillatory Motion of a Rigid Plate," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2010.

[3] D. R. Frutiger, B. E. Kratochvil, K. Vollmers, and B. J. Nelson, "Small, fast, and under control: wireless resonant magnetic micro-agents," *Int. J. Rob. Res.*, vol. 29, no. 5, 2010.

[4] Y. Sunahara, T. Kabeuchi, Y. Asada, S. Aihara, and K. Kishino, "On stochastic controllability for nonlinear systems," *Automatic Control, IEEE Transactions on*, vol. 19, no. 1, 1974.

[5] J. Zabczyk, "Controllability of stochastic linear systems," *Systems & Control Letters*, vol. 1, no. 1, 1981.

[6] P. Reist and R. D'Andrea, "Design and Analysis of a Blind Juggling Robot," *Robotics, IEEE Transactions on*, vol. PP, no. 99, 2012.

[7] "Blind Juggler Website." [Online]. Available: www.blindjuggler.org

[8] P. J. Holmes, "The dynamics of repeated impacts with a sinusoidally vibrating table," *Journal of Sound Vibration*, vol. 84, Sept. 1982.

[9] R. Leine and T. Heimsch, "Global uniform symptotic attractive stability of the non-autonomous bouncing ball system," *Physica D: Nonlinear Phenomena*, May 2011.

[10] T. L. Vincent, "Control using chaos," *IEEE Control Systems Magazine*, vol. 17, no. 6, 1997.

[11] E. Ott, C. Grebogi, and J. A. Yorke, "Controlling Chaos," *Phys. Rev. Lett.*, vol. 64, no. 11, 1990.

[12] J. Baillieul, "Self-organizing Behavior in a Simple Controlled Dynamical System," in *Stoch. Theory and Adaptive Control*. Springer, 1992.

[13] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek, "Sequential Composition of Dynamically Dexterous Robot Behaviors," *Int. J. Rob. Res.*, vol. 18, no. 6, 1998.

[14] R. Ronsse, P. Lefevre, and R. Sepulchre, "Rhythmic Feedback Control of a Blind Planar Juggler," *IEEE Transactions on Robotics*, vol. 23, no. 4, 2007.

[15] P. Kulchenko and E. Todorov, "First-exit model predictive control of fast discontinuous dynamics : Application to ball bouncing," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2011.

[16] T. Shinbrot, C. Grebogi, J. A. Yorke, and E. Ott, "Using Small Perturbations to Control Chaos," *Nature*, vol. 363, 1993.

[17] P. Parrilo, "Semidefinite programming relaxations for semialgebraic problems," *Mathematical Programming*, vol. 96, no. 2, 2003.