# Chaotic pendulum: The complete attractor

Robert DeSerio[a)]
*Department of Physics, University of Florida, Gainesville, Florida 32611*

(Received 16 April 2002; accepted 11 October 2002)

A commercial chaotic pendulum is modified to study nonlinear dynamics, including the determination of Poincaré sections, fractal dimensions, and Lyapunov exponents. The apparatus is driven by a simple oscillating mechanism powered by a 200 pulse per revolution stepper motor running at constant angular velocity. A computer interface generates the uniform pulse train needed to run the stepper motor and, with each pulse, reads a rotary encoder attached to the pendulum axle. Ten million readings from overnight runs of 50 000 drive cycles were smoothed and differentiated to obtain the pendulum angle $\theta$ and the angular velocity $\omega$ at each pulse of the drive. A plot of the 50 000 $(\theta,\omega)$ phase points corresponding to one phase of the drive system produces a single Poincaré section. Thus, 200 Poincaré sections are experimentally available, one at each step of the drive. Viewed separately, any one of them strikingly illustrates the fractal geometry of the underlying chaotic attractor. Viewed sequentially in a repeating loop, they demonstrate the stretching and folding of phase point density typical of chaotic dynamics. Results for four pendulum damping conditions are presented and compared. 
[DOI: 10.1119/1.1526465]

## I. INTRODUCTION

The tremendous interest in nonlinear dynamics has brought with it a need for suitable introductory experiments. Although the chaotic pendulum is an obvious candidate, only a few implementations have been proposed for use in teaching laboratories.[1–3] A torsional pendulum[1] and a mechanical Duffing oscillator[2] have been described, but experimental Poincaré sections were not collected. Three commercial chaotic pendulums were reviewed in 1998[3] and components of one of those systems were used here. Our apparatus is particularly well suited for an introductory treatment because it provides quality data, and its simple design means that the equations of motion can be readily derived and experimentally verified for both regular and chaotic behavior.

Not including a desktop personal computer and an assortment of clamps and rods, the total system cost can be kept below $1000. The rotary encoder and pendulum components are available from Pasco for less than $300;[4] the counter/timer board, connector block, and cable are available from National Instruments for less than $500;[5] and the stepper motor system can be put together for less than $200.[6] A 500 MHz Pentium III PC was available, as was the LabVIEW software development system used for all data acquisition and analysis.[7] A data acquisition program is available on the University of Florida Physics Department web site[8] and does not require a LabVIEW license. The acquisition and analysis programs (virtual instruments) are also available,[8] but are unsupported and require a LabVIEW software license. Other nonlinear analysis and graphing packages such as TISEAN[9] and gnuplot[10] could also be used.

The stepper motor and data acquisition system are critical replacements for the corresponding components currently supplied by Pasco. A Poincaré section requires that the pendulum angle and angular velocity be determined each time the drive system passes a single point in its oscillatory motion, and an optical or magnetic pickup is often placed on the drive mechanism for this purpose. For example, the Daedalon pendulum reviewed in Ref. 3 passes such a pickup signal to its data acquisition interface and is the only one of the three reviewed systems capable of creating Poincaré sections. Moreover, each hardware pickup normally permits only a single Poincaré section to be collected at one time. In effect, our implementation creates a software pickup at every angular position of the motor, thus making possible the creation of Poincaré sections for almost all drive phases and providing accurate three-dimensional phase space coordinates of points along a trajectory.

The mechanical components and the data acquisition hardware and software are described in Sec. II. Data fitting examples and calculations of fractal dimensions and Lyapunov exponents are described in Sec. III.

## II. APPARATUS

A schematic of the apparatus is shown in Fig. 1. The pendulum, rotary encoder, and stepper motor are mounted on a 6 foot, 1/2 in. diameter steel rod. Two identical springs are attached on either side of the pendulum by a string wrapped twice around the pulley to prevent slipping. The end of one spring is driven up and down by a stepper motor running at constant angular velocity via a string passing through a guide hole in a cross rod. The drive amplitude $A$ is adjusted by changing the length of the shaft attached to the motor. The end of the other spring is fixed with a short string to a tuning peg also mounted on the cross rod.

The pendulum axle is part of the rotary encoder which transmits logic pulses (1440/rev) on separate phonojacks for clockwise and counterclockwise rotations. One of the data acquisition counters performs up/down counting of these pulses without any need for preprocessing. The pulse count, directly proportional to the pendulum rotation angle, can be saved to a counter register at any time by strobing the counter's gate.

The stepper motor requires 200 pulses per revolution and runs at a frequency near one revolution per second. The near 200 Hz square wave sent to the controller is divided down from a 20 MHz clock by another data acquisition counter. The frequency resolution (around 10 ppm) and stability (less than 50 ppm drift in a 24 hour period) are more than ad-

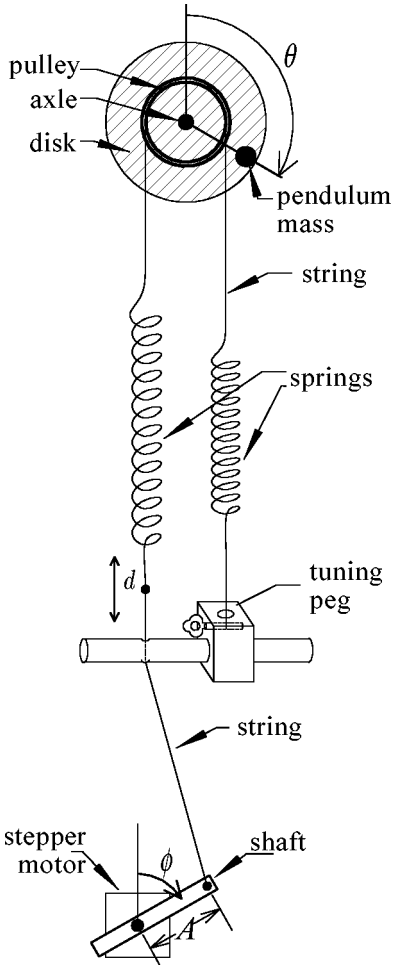


Fig. 1. The chaotic pendulum and drive system. Not shown is the Pasco rotary motion sensor containing the axle and pulley.

equate for this experiment. This square wave also strobes the gate of the angular counter with the result that 200 readings are collected per drive period. The raw readings are transferred to main memory by direct memory access hardware. Thus, all data acquisition is hardware timed and the angular readings stay synchronized to the drive phase even over long runs. After programming the counters for the correct operating modes and drive frequency, the stepper motor pulses are initiated and the direct memory access buffer fills using almost no computer resources. The buffer is written to disk every few seconds and with typical overnight runs of 50 000 drive cycles, two bytes per reading, and 200 readings per period, file sizes are around 20 MB.

The raw counter readings, spaced 1/200th of the drive period apart in time, are scaled and numerically smoothed and differentiated using Savitsky–Golay filtering.[11] The filters are equivalent to a least-squares polynomial fit for each data point with the fitting region symmetrically surrounding the point. They provide the angle $\theta$, the angular velocity $\omega = d\theta/dt$, and the angular acceleration $\alpha = d^2\theta/dt^2$. Thirty-three point, quartic polynomial filtering was used for all figures and analysis. Simulations with noisy 200 point sine waves showed good filter performance with these parameters. The random uncertainties in the filtered $\theta$, $\omega$, and $\alpha$ values are predicted to be $\sigma_\theta = 0.4$ mrad, $\sigma_\omega = 9$ mrad/s, and $\sigma_\alpha = 0.5$ rad/s$^2$. These are measurement uncertainties only, based solely on an encoder digitizing error of $\pm 1/2$ counts and the covariance matrix for the filters. They do not include effects of dynamical noise[12]—real fluctuations in the variables that subsequently propagate through the dynamics of the system.

As shown in Fig. 1, the pendulum angle $\theta$ is defined from the vertical with $\theta = 0$ for the straight up or inverted position and clockwise angles taken to be positive. The stepper motor phase $\phi$ is also defined from the vertical, and thus the displacement of the driven end of the spring is given by $d = A \cos \phi$.[13]

The springs on either side of the pendulum are relatively soft and stay stretched at all times. With the pendulum mass removed and $d = 0$, the tuning peg is adjusted to get the equilibrium angle to the inverted $\theta = 0$ position. Then, with the pendulum mass $m$ attached and at a distance $L$ from the axis, the potential energy can be expressed as

$$V = \tfrac{1}{2}k(r\theta)^2 + \tfrac{1}{2}k(r\theta - d)^2 + mgL\cos\theta, \tag{1}$$

where $k$ is the force constant of each spring and $r$ is the pulley radius.

A two-term dissipation model provides a reasonable approximation for the energy loss in our apparatus. Eddy-current damping, which can be adjusted by moving a magnet near the rotating aluminum disk, is modeled by a viscosity-like torque of the form $-b\,\omega$. Also present is some degree of dry, or Coulomb, friction arising largely in the axle bearings. This axle friction will be assumed to be constant in magnitude, but opposite to the angular velocity, and will be modeled by a torque of the form $-b'\,\mathrm{sgn}\,\omega$, where $\mathrm{sgn}\,\omega = \omega/|\omega|$. Static friction or other effects as the pendulum momentarily stops each time $\omega$ passes through zero will not be modeled.

With the two damping torques and the conservative torque $\tau_c = -dV/d\theta$, the equations of motion can be set in the autonomous form

$$\frac{d\theta}{dt} = \omega, \tag{2a}$$

$$\frac{d\omega}{dt} = -\Gamma\omega - \Gamma'\,\mathrm{sgn}\,\omega - \kappa\theta + \mu\sin\theta + \epsilon\cos\phi, \tag{2b}$$

$$\frac{d\phi}{dt} = \Omega, \tag{2c}$$

where

$$\Gamma = b/I, \tag{3a}$$

$$\Gamma' = b'/I, \tag{3b}$$

$$\kappa = 2kr^2/I, \tag{3c}$$

$$\mu = mgL/I, \tag{3d}$$

$$\epsilon = Akr/I. \tag{3e}$$

$I$ is the moment of inertia of the pendulum, and $\Omega$ is the angular frequency of the drive.

## III. ANALYSIS

### A. Chaotic attractors

The lower part of Fig. 2 shows the potential $V(\theta)$ in Eq. (1) for $d = 0$. The upper part of the figure is constructed by plotting $\omega = \pm\sqrt{2[E - V(\theta)]/I}$ versus $\theta$ and shows phase

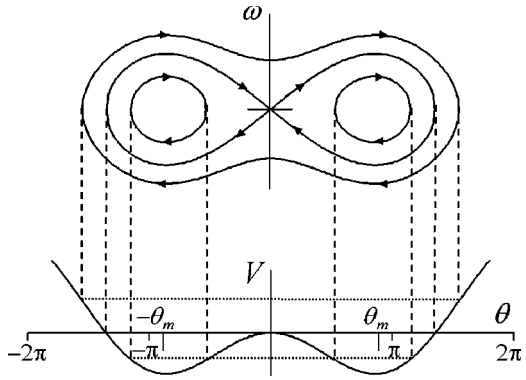


Fig. 2. Lower, near scale drawing of the potential $V(\theta)$ for the system of Fig. 1. Upper, phase space trajectories at three energies.

space trajectories for undriven, undamped vibrations at three representative values of the energy $E$. The turning points of the vibrations are illustrated by the dashed vertical lines and the mechanical energy by the horizontal dotted lines. The direction of motion along a trajectory is indicated by the arrows.

The starting point in the $\theta$–$\omega$ phase plane determines the energy and subsequent motion on the appropriate trajectory. The two smallest loops centered around the elliptic fixed points $(\pm\theta_m,0)$ represent near-harmonic oscillations about the minima in $V(\theta)$, while the outermost trajectory represents ''over-the-top'' oscillations about $\theta=0$. Between these oscillatory trajectories is the separatrix. Started from any point on this $\infty$-shaped trajectory, the pendulum ultimately comes to a stop at the hyperbolic fixed point (0,0).

Data for the driven pendulum undergoing chaotic vibrations are displayed in Fig. 3. The upper part of Fig. 3 demonstrates the nonrepeating motion as a plot of $\theta$ versus $t$. The lower part of Fig. 3 demonstrates the changing mechanical energy as the trajectory spirals into and away from the low energy elliptic points. On the left, a short trajectory segment of seven drive periods can be followed from beginning to end. On the right, an untraceable trajectory of 200 periods demonstrates that although the $(\theta,\omega)$ values are not constrained by energy conservation, their average motion in phase space is similar to those that are (see for example, top of Fig. 2).

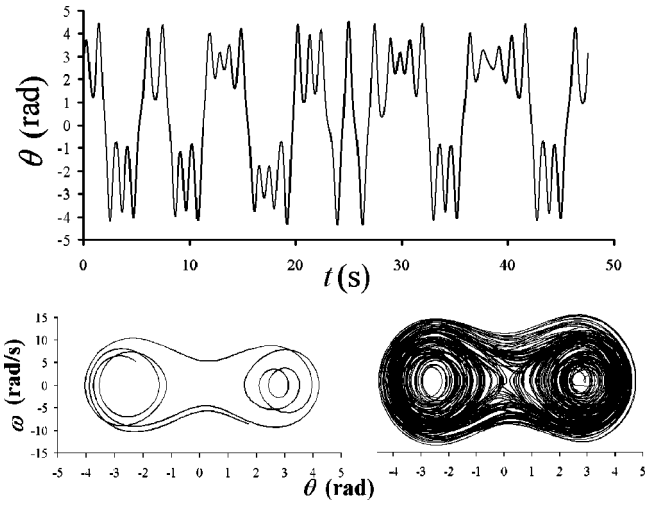


Fig. 3. Upper, experimental $\theta$ versus $t$ data for chaotic motion. Lower, phase space trajectories for a few drive periods (left) and many drive periods (right).

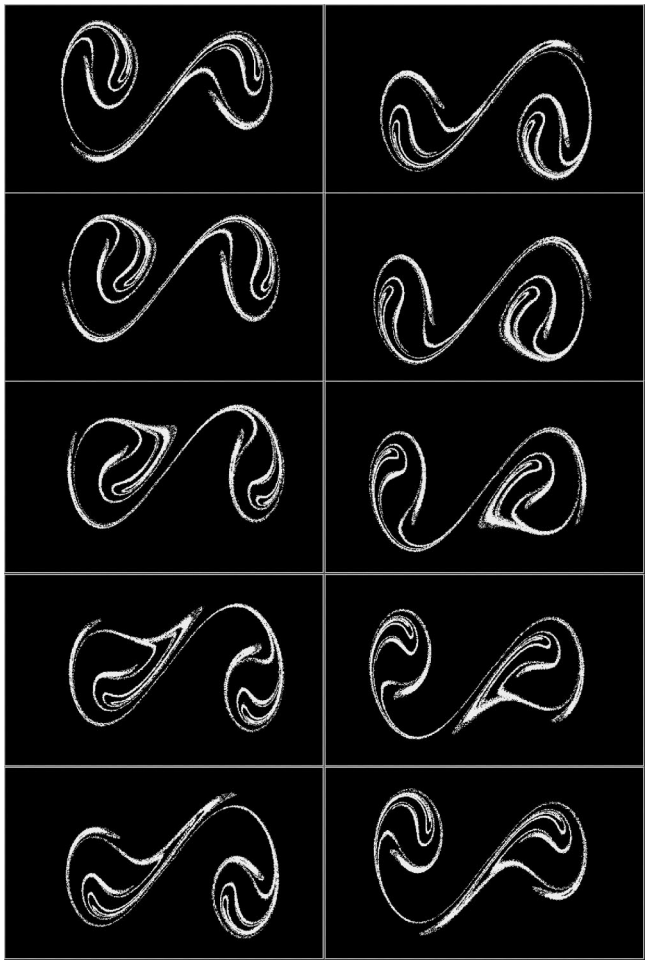

Fig. 4. Ten Poincaré sections at equally spaced phase intervals of the drive cycle—from the top down in the first column then the second. High damping conditions with the magnet set 3 mm from the disk.

For driven motion, the state of the pendulum is completely specified by a triplet of values $(\theta,\omega,\phi)$. Any such triplet can be represented as a point in a three-dimensional phase space and may be considered initial conditions that evolve deterministically according to the equations of motion. By providing unique values for the derivatives of each variable, Eq. (2) specifies a continuous flow which predicts that even a chaotic trajectory will lie on a non-crossing curve in this three-dimensional phase space. That chaotic trajectories will ultimately evolve on a fractal in this phase space is not obvious.

In phase space, chaotic trajectories for dissipative systems are confined to complex regions known as attractors. They are called strange attractors because geometrically they are fractals. The experimental Poincaré sections of Figs. 4 and 5 are also fractals, two-dimensional slices through the full three-dimensional attractor. In Figs. 4 and 5 these Poincaré sections are shown for high and low damping conditions, respectively. Each Poincaré section is a plot of 50 000 $(\theta,\omega)$ phase points, one per drive cycle as the drive phase passes through a particular value of $\phi$. The center of each plot is (0,0) with the $\theta$-axis horizontal and the $\omega$-axis vertical. The scale of the figures is approximately $\pm 6$ rad and $\pm 14$ rad/s. Starting transients, which can produce phase points off the

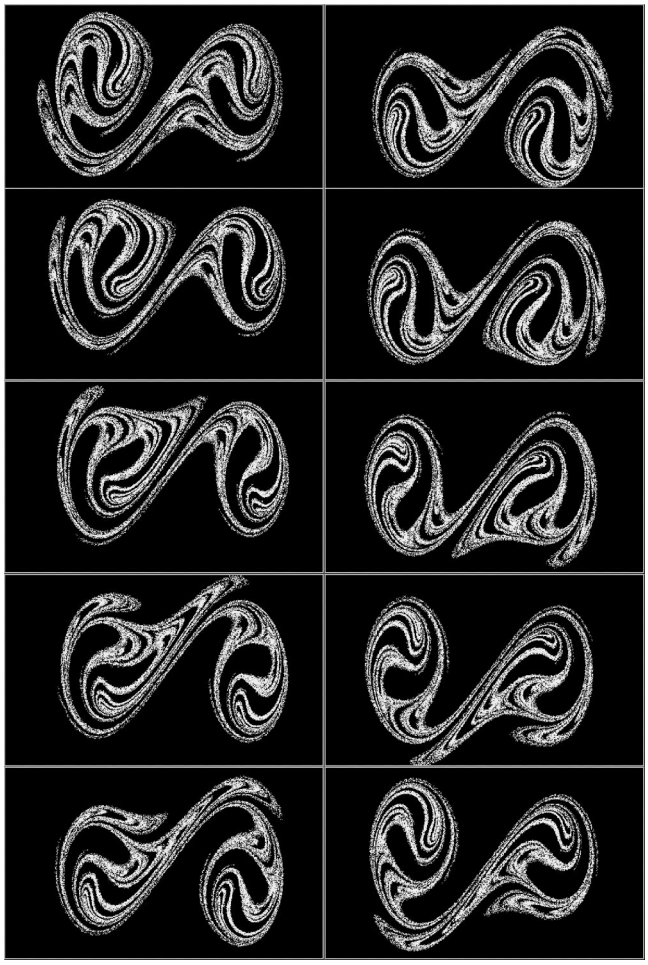

Fig. 5. Ten Poincaré sections at equally spaced phase intervals of the drive cycle, from the top down in the first column then the second. Low damping conditions with the magnet set 18 mm from the disk.

attractor, were eliminated from the plots and analysis by skipping the first 100 drive cycles. The drive frequency for all data sets was 0.82 Hz.

Each figure shows 10 Poincaré sections equally spaced in the drive phase starting at the top left with $\phi=0$. The drive phase increases by 1/10th of a rotation going down the columns and continues from bottom left to top right. The drive is then 180° out of phase between pairs on the same row, and the symmetry between these pairs indicates the spring equilibrium angle was well matched to the inverted pendulum position. If we connect the bottom right section back to the top left, we can then follow a particular feature as it changes position and shape from one section to the next and from one period to the next. Some features can be followed through several drive cycles, stretching and folding until finally thinning out of existence.

Table I. Mechanical parameters.

| Quantity | Value |
|---|---|
| Pendulum mass, $m$ | 15 g |
| Pendulum mass offset, $L$ | 4.8 cm |
| Disk mass | 128 g |
| Disk diameter | 9.5 cm |
| Pulley radius, $r$ | 1.4 cm |
| Spring constant, $k$ | 2600 dyne/cm |
| Moment of inertia without $m$, $I_0$ | 1400 g cm$^2$ |
| Moment of inertia with $m$, $I$ | 1800 g cm$^2$ |

Simulations of Eq. (2) using a fourth-order Runge–Kutta algorithm with 40 time steps per drive period produced attractors similar to those of the experiment. The Poincaré sections also look very similar to those reported by Peters[1] for $\Gamma'=0$. As can be seen in Fig. 6, the simulations were remarkably similar to but noticeably sharper than their experimental counterparts, whose features become washed out at resolution levels dependent on the measurement errors and dynamical noise present in the apparatus.

The obvious difference between the distribution of points in Figs. 4 and 5 is a result of the different damping conditions. Without axle friction ($\Gamma'=0$), the fractional rate of change of occupied volumes in phase space is predicted to be $-\Gamma$ everywhere. With stronger damping, the larger contraction rate produces more highly peaked ridges in the phase space density. For $\Gamma=0$, phase space volumes are conserved and a simulation with both $\Gamma=\Gamma'=0$ produced Poincaré sections with sharp boundaries and a complex flow pattern; however, the density of phase points was nearly uniform throughout.

### B. System parameters

The mechanical parameters are listed in Table I. The spring constant $k$ was determined in separate measurements of the oscillation period as a function of the mass hanging directly from the spring.

The parameters of Eq. (2) can be determined for each data set from a least squares fit.[14] The angular accelerations $\alpha_i$ obtained via the Savitsky–Golay filters are fit to the function $\alpha(\theta_i,\omega_i,\phi_i)$ provided by the right side of Eq. (2b). The $\theta_i$ and $\omega_i$ in the fitting function are also obtained via the filters, while $\phi_i$ is determined from the data point index $i$ (modulo

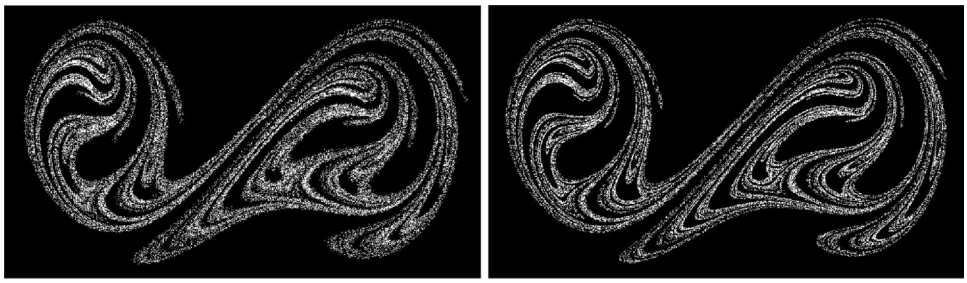

Fig. 6. Comparison between experiment (left) and simulation (right) for the $D=18$ mm damping distance. The simulation parameters were taken from Table II.

Table II. Fitted system parameters [Eq. (3)], capacity dimension $d_0$, Lyapunov exponents $\lambda_i$, and Lyapunov dimension $d_L$ for four magnet damping distances $D$. The estimated standard deviation of the fitted system parameters is represented by the notation s.d.

| $D$ (mm) | $\Gamma$ $s^{-1}$ | $\Gamma'$ $s^{-2}$ | $\kappa$ $s^{-2}$ | $\mu$ $s^{-2}$ | $\epsilon$ $s^{-2}$ | $d_0$ | $\lambda_1$ $s^{-1}$ | $\lambda_2$ $s^{-1}$ | $d_L$ |
|---|---|---|---|---|---|---|---|---|---|
| 18 | 0.05 | 1.22 | 5.81 | 37.52 | 9.97 | 1.73 | 0.62 | −0.81 | 1.77 |
| 7 | 0.12 | 1.10 | 5.95 | 38.09 | 10.07 | 1.68 | 0.57 | −0.83 | 1.69 |
| 5 | 0.31 | 0.90 | 5.83 | 38.08 | 12.05 | 1.63 | 0.50 | −0.87 | 1.57 |
| 3 | 0.53 | 1.18 | 5.81 | 37.61 | 11.98 | 1.49 | 0.42 | −1.06 | 1.40 |
| s.d. | 0.02 | 0.07 | 0.01 | 0.03 | 0.02 | | | | |

200) multiplied by the motor step size $2\pi/200$. The fit is a linear regression of $\alpha$ to the terms $\omega$, $\operatorname{sgn}\omega$, $\theta$, $\sin\theta$, and $\cos\phi$. The terms $\cos\theta$, $\sin\phi$, and a constant are added to the regression to account for possible offset errors in $\theta$, $\phi$, and the spring equilibrium angle. The fitted coefficients for these eight terms determine $\Gamma$, $\Gamma'$, $\kappa$, $\mu$, and $\epsilon$, as well as the three offset angles.

Results for four overnight runs at different damping magnet distances $D$ are given in Table II. Each row is an average from three separate fits to trajectories of 5000 drive cycles at 10 points per cycle. Points with $\omega$ near zero—more specifically, about 1/3 of the points having $|\omega|<3$ rad/s—were excluded from the fits. As demonstrated shortly, this restriction is needed to prevent systematic errors in the two dissipation coefficients. The data for the three fits were from the beginning, middle, and end of each run. Parameter variations within each set were used to estimate the standard deviations given in the last row of Table II. Variations in the fitted spring equilibrium angle and in the offset angles in $\theta$ and $\phi$ were less than 4 mrad and indicate no significant problems due to string-pulley slippage, pulses lost from the rotary encoder, or pulses lost to the stepper motor during these 17 hour runs.

The drive amplitude was 5 cm for the two lower damping settings and had to be increased to 6 cm to obtain chaotic motion at the two higher damping settings. In good agreement with the fits, these amplitudes and the mechanical parameters in Table I give the following parameter values: $\kappa=5.7\ s^{-2}$, $\mu=39\ s^{-2}$, and $\epsilon=11\ s^{-2}$ ($13\ s^{-2}$) for the first (last) two rows.

The increasing $\Gamma$ values in Table II are expected as the magnet is brought closer to the aluminum disk, and the variations in the other parameters from run to run are not unreasonable. The four data sets were taken over a week or two and some control variables may not have been kept constant. For example, a spring may have been changed (for another of the same kind) and the pendulum rotation axis may have been off the horizontal with a degree or so of variability. The axle friction coefficient varied considerably over time, at one point decreasing from 1.7 to 0.9 $s^{-2}$ after the bearings were oiled.

The goodness of fit can be assessed by the fit residuals which are defined as $\delta\alpha_i=\alpha_i-\alpha(\theta_i,\omega_i,\phi_i)$. They are the deviations of the measured $\alpha_i$ from the fit values $\alpha(\theta_i,\omega_i,\phi_i)$. Root mean square (rms) deviations were in the range 0.6–0.8 rad/$s^2$—comparable to the acceleration error $\sigma_\alpha=0.5$ rad/$s^2$ expected from the Savitsky–Golay filter and less than 3% of the rms angular acceleration, which was approximately 30 rad/$s^2$ for all data sets.

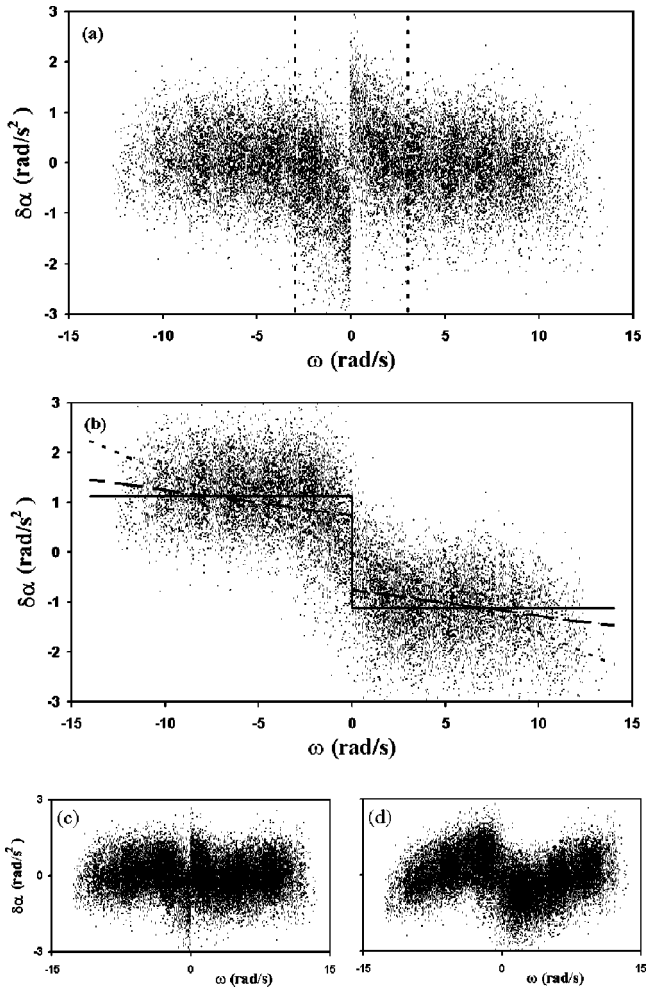


Fig. 7. (a) Fit residuals versus $\omega$. Points between the dotted vertical lines were not included in the fit. (b) Deviations in (a), but without the axle friction contribution subtracted. Solid line, the axle friction contribution from the fit. Dashed line, best fit when all points are included in the fit, resulting in the fit residuals shown in (c). Dotted line, best fit using viscous damping only and resulting in the fit residuals shown in (d).

The reason for excluding data points with $|\omega|<3$ rad/s is demonstrated in Fig. 7. Plot (a) shows the residuals for one of the three fits for the $D=7$ mm data set. The deviations are from a range of $\theta$, $\omega$, and $\phi$ values, but they are plotted versus $\omega$ alone. Thus, the dotted lines at $\omega=\pm3$ rad/s isolate the discontinuous ''kink'' of deviations not considered in the fit. Placing a finger over the excluded data helps show that the deviations that are considered in the fit are evenly distributed about $\delta\alpha=0$ for all $\omega$.

The cause of the kink is made more obvious in Fig. 7(b) which shows the same deviations, but without subtracting the fitted axle friction contribution. Instead, this contribution is shown separately by the solid line. By construction, the residuals of Fig. 7(a) are the differences between the deviations and the solid line fit in Fig. 7(b). Despite the experimental noise of comparable size, the data clearly agree with the fitted axle friction term, although its discontinuous change at $\omega=0$ occurs gradually—over a $\pm3$ rad/s range—in the experimental deviations. This gradual change is an expected smoothing effect of the Savitsky–Golay filters when applied to data with discontinuities. Although the exact smoothing range is difficult to predict, a rough estimate can be obtained

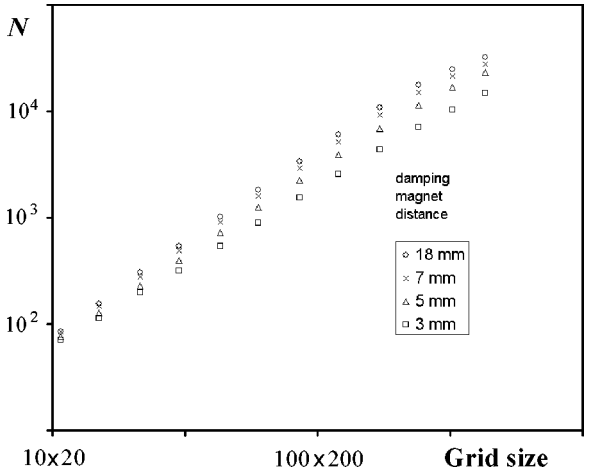


Fig. 8. A log–log plot of the number of grid boxes needed to cover the Poincaré section versus the inverse size of the boxes for the four damping magnet settings. The slope is the capacity dimension $d_0$ given in Table II.

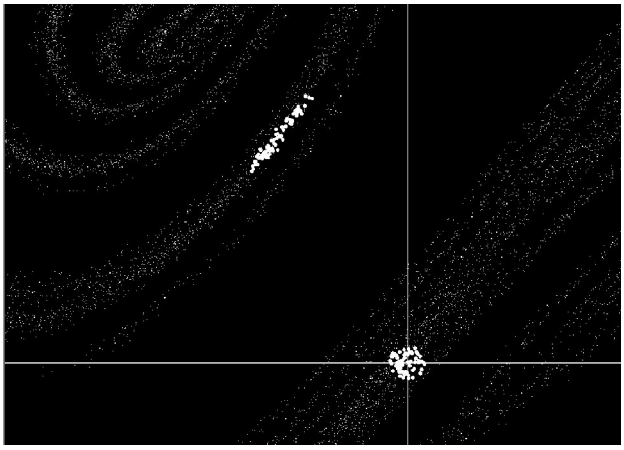

Fig. 9. A small area of a Poincaré section showing how an initial set of $\delta\mathbf{u}$ (at the cursor) move and change shape when propagated one period. The new set is the bright, elongated patch slightly above and left of center.

from the total time interval $\delta t \approx 0.2$ s for collecting the 33 raw encoder readings used for each $\theta_i$, $\omega_i$, and $\alpha_i$. In that interval, a pendulum moving with a given angular acceleration $\alpha$ changes its angular velocity over a range $\delta\omega \approx \alpha\,\delta t$. The rms value $\alpha \approx 30$ rad/s$^2$ gives $\delta\omega \approx 6$ rad/s, which is acceptably close to the 3 rad/s value determined empirically. The smoothing is responsible for the kink and would cause systematic errors in $\Gamma$ and $\Gamma'$ were the fits not restricted to data points with $|\omega| > 3$ rad/s.

The systematic errors can be demonstrated in unrestricted fits, that is, those including all 50 000 points in the trajectory. In all such fits, $\Gamma$ increases and $\Gamma'$ decreases, with no significant changes to the other parameters. For the data set under consideration, the unrestricted $\Gamma$ and $\Gamma'$ were 0.18 s$^{-1}$ and 0.75 s$^{-2}$ compared to their restricted counterparts 0.13 s$^{-1}$ and 1.13 s$^{-2}$. The sloping $z$ shape of the dashed line in Fig. 7(b) shows these alternate fit contributions—an axle friction term with $\Gamma' = 0.75$ s$^{-2}$ and a 0.05 s$^{-1}$ viscous term. (Keep in mind that a 0.13 s$^{-1}$ viscous term is already subtracted.) The final residuals for the unrestricted fit are given in Fig. 7(c) and show a systematic error—a small downturn in the negative $\omega$ direction and an upturn in the positive $\omega$ direction—indicative of the overestimated $\Gamma$. These alternate residuals are significantly smaller in the restricted region, but this is where the discontinuity in the model and the smoothing effects of the Savitsky–Golay filters are expected to cause systematic deviations. Thus, although their rms value of 0.67 rad/s$^2$ is worse than the 0.61 rad/s$^2$ for the restricted fit, because it is better than the 0.71 rad/s$^2$ for the restricted fit over all points, the fit must be restricted to avoid obtaining the flawed coefficients.

Fits without the axle friction term were also performed. The fits compensated by increasing the viscous coefficient with no significant changes to the other parameters. For this data set, $\Gamma$ increased to 0.28 s$^{-1}$. The sloping dotted line in Fig. 7(b) represents the additional 0.15 s$^{-1}$ viscous term. The residuals for this fit are shown in Fig. 7(d) and have an rms value of 0.78 rad/s$^2$. The distinct sawtooth pattern in this plot is a result of greatly overestimating $\Gamma$ to take into account the effects of axle friction.

### C. Capacity dimension and Lyapunov exponents

The capacity dimension $d_0$ was determined for the $\phi = 0$ Poincaré sections using the standard box-counting algorithm.[15] The $\theta$–$\omega$ plane is divided into a rectangular grid and the number of rectangles $N$ containing at least one phase point is determined. Values of $N$ are redetermined as the linear size $s$ of the rectangles is decreased. The slope of the $\log N$ versus $\log 1/s$ graph is the capacity dimension and would have a slope of 0, 1, or 2 for a set of points covering a finite number of phase space points, lines, or areas.

Figure 8 shows the results for the four data sets and demonstrates a small decrease in slope as the damping increases. Starting with a $10\times20$ grid and just under 100 grid boxes occupied, the predicted straight line (fractal) behavior can be seen to continue for more than an order of magnitude in the linear dimension. Two effects limit the range of fractal behavior as the grid size decreases. First, the finite number of points in the data set imply that the smallest grid boxes will ultimately become empty of phase points because of the paucity of data and not because the underlying Poincaré section has no density in those boxes. The graph will then flatten out as $N$ approaches the number of data points. The last 2 to 4 points in the plots of Fig. 8 showed this effect and were not included in the determination of $d_0$ given in the seventh column of Table II. The second limiting effect arises from experimental noise which washes out fractal features at scales commensurate with the noise. If the data set has enough points, the graph slope changes to that of the dimension of the grid at grid scales much smaller than the noise. Larger data sets would be needed to see this effect.

As shown in Fig. 9, a small group of neighboring phase points in one area of a Poincaré section map to a new neighborhood one period later. Evidence of chaotic motion's extreme sensitivity to initial conditions can be seen by following such a group of phase points in time. Within a few drive periods they become dispersed throughout the Poincaré section. Lyapunov spectra characterize the long term exponential growth (and decay) in the separation of very close phase points. For our four data sets, they are given in columns eight and nine of Table II. The calculation, based on the algorithm of Eckmann and Ruelle,[16] is briefly described next.

A flow map $\mathbf{F}(\mathbf{u})$ relates any phase $\mathbf{u}_i = (\theta_i, \omega_i)$ to its value $\mathbf{u}_{i+1} = \mathbf{F}(\mathbf{u}_i)$ one time step later. Although $\mathbf{F}(\mathbf{u})$ cannot be

expressed in closed form, it can be determined from the equations of motion.[17] It provides an alternate description of the pendulum dynamics and is well suited for experimental determination of the Lyapunov exponents. A time step $\tau = T/10$, that is, 1/10th of the drive period was acceptable for our data sets, and thus the computation begins by constructing and storing in computer memory 10 equally spaced Poincaré sections like those of Figs. 4 and 5.

Starting from two slightly different phases on the same Poincaré section $\mathbf{u}_i$ and $\mathbf{u}_j = \mathbf{u}_i + \delta\mathbf{u}_j$, the evolution of the deviation $\delta\mathbf{u}_j$ to its value $\delta\mathbf{u}_{j+1} = \mathbf{u}_{j+1} - \mathbf{u}_{i+1}$ one time step later—called a tangent mapping—can be modeled by a Taylor expansion about the point $\mathbf{u}_i$.

$$\delta\mathbf{u}_{j+1} = \frac{\partial \mathbf{F}}{\partial \mathbf{u}} \delta\mathbf{u}_j + \cdots, \tag{4}$$

where the $2\times 2$ Jacobian matrix $\partial\mathbf{F}/\partial\mathbf{u}$ is evaluated at $\mathbf{u}_i$. The ellipsis represents terms containing higher order derivatives which become negligible in the limit of infinitesimal $\delta\mathbf{u}_j$.

The main calculation is to determine the Jacobian matrices at sequential points along a trajectory. The computation first finds all phases $\mathbf{u}_j$ on a single Poincaré section within some fixed ellipse centered on $\mathbf{u}_i$. It then determines the deviations $\delta\mathbf{u}_j$ and the corresponding values $\delta\mathbf{u}_{j+1}$ one time step later. The program then performs two linear regressions—one for $\delta\theta$ and one for $\delta\omega$. If we let $\delta x$ represent either $\delta\theta$ or $\delta\omega$, the regressions can be expressed as

$$\begin{aligned}\delta x_{j+1} = {} & a_x + b_{x\theta}\delta\theta_j + b_{x\omega}\delta\omega_j + c_{x\theta\theta}(\delta\theta_j)^2 \\ & + c_{x\omega\omega}(\delta\omega_j)^2 + c_{x\theta\omega}(\delta\theta_j)(\delta\omega_j).\end{aligned} \tag{5}$$

The constant terms $a_x$ are included to take into account possible offset errors in $\mathbf{u}_i$ and $\mathbf{u}_{i+1}$. The quadratic terms $c_x$ account for second-order terms in the Taylor series expansion and decrease the errors in the linear terms that arise from the finite size of the $\delta\mathbf{u}$ used in the fit.[18] The four coefficients of the linear terms, $b_{\theta\theta}$, $b_{\theta\omega}$, $b_{\omega\theta}$, and $b_{\omega\omega}$, are taken as estimates of the Jacobian matrix elements at the point $\mathbf{u}_i$.

A fixed neighborhood ellipse was used when finding the initial deviations $\delta\mathbf{u}_j$; radii were chosen at 0.15 rad in $\theta$ and 0.30 rad/s in $\omega$. The number of deviations in each fit typically varied between 30 and 80, although a point was occasionally eliminated ($<1\%$) from the analysis when there were less than 10 deviations available. The rms deviations of the fits described by Eq. (5) were typically around 2 mrad for the $\delta\theta$ fit and 20 mrad/s for the $\delta\omega$ fit.

The product of $N$ sequential Jacobian matrices gives the evolution of the smallest $\delta\mathbf{u}$ through $N$ time periods. For large enough $N$, this product is predicted to have eigenvalues $\exp(\lambda_1 N\tau)$ and $\exp(\lambda_2 N\tau)$, where $\lambda_1$ and $\lambda_2$ are the Lyapunov exponents. The eigenvalues of this product are experimentally determined via $Q-R$ decompositions of each Jacobian matrices along the trajectory, say $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_N$. The decomposition expresses a given matrix as a product $\mathbf{QR}$ of an orthonormal matrix $\mathbf{Q}$ and an upper diagonal matrix $\mathbf{R}$. For the Lyapunov algorithm, each Jacobian matrix is first multiplied on the right by the prior $\mathbf{Q}$ before performing the decomposition,

$$\left.\frac{\partial \mathbf{F}}{\partial \mathbf{u}}\right|_{\mathbf{u}_i} \mathbf{Q}_{i-1} = \mathbf{Q}_i \mathbf{R}_i. \tag{6}$$

(The starting $\mathbf{Q}_0$ is taken as the identity matrix.) The product of the Jacobian matrices then becomes

$$\left.\frac{\partial \mathbf{F}}{\partial \mathbf{u}}\right|_{\mathbf{u}_N} \cdots \left.\frac{\partial \mathbf{F}}{\partial \mathbf{u}}\right|_{\mathbf{u}_2} \left.\frac{\partial \mathbf{F}}{\partial \mathbf{u}}\right|_{\mathbf{u}_1} = \mathbf{Q}_N \mathbf{R}_N \cdots \mathbf{R}_2 \mathbf{R}_1. \tag{7}$$

The single remaining $\mathbf{Q}_N$ can be neglected when determining the eigenvalues of the product on the right-hand side. Because all the $\mathbf{R}$ matrices are upper diagonal, the eigenvalues of their product are the product of their diagonal elements. Equating this eigenvalue to the prediction, $\exp(\lambda_k N\tau)$, and taking natural logarithms then gives

$$\lambda_k = \lim_{N\to\infty} \frac{1}{N\tau} \sum_{i=1}^{N} \ln R_i^{kk}. \tag{8}$$

All reported exponents are from at least 1000 drive periods following 100 periods used to let the $Q$ matrices stabilize. For both experimental and simulated data sets, the individual exponents varied by $\pm 0.03\ \mathrm{s}^{-1}$ as the starting point of the trajectory changed, which may be taken as a rough estimate of their uncertainty.

The Lyapunov program was also applied to the simulations. Without axle friction, the sum of the Lyapunov exponents is predicted to be equal to $-\Gamma$.[19] For a simulation created with a viscous damping coefficient of $\Gamma = 0.10\ \mathrm{s}^{-1}$, the Lyapunov exponents were $(0.61, -0.71)\ \mathrm{s}^{-1}$. For $\Gamma = 0.30\ \mathrm{s}^{-1}$ they were $(0.63, -0.93)\ \mathrm{s}^{-1}$. Both results are in excellent agreement with this prediction. However, for the experimental results of Table II, the sum of the Lyapunov exponents is systematically below $-\Gamma$ from the fit by the amounts 0.13, 0.14, 0.06, and $0.11\ \mathrm{s}^{-1}$ for the four data sets. The direction and rough size of this shift indicated it is likely due to the effects of axle friction. (Recall that $\Gamma$ increases by about $0.15\ \mathrm{s}^{-1}$ if axle friction is modeled entirely by additional viscous damping.) The effect was also observed in the simulations. When an axle friction contribution of $\Gamma' = 0.8\ \mathrm{s}^{-2}$ was added to the two simulations above, both exponent sums decreased by $0.09\ \mathrm{s}^{-1}$.

An additional check on the consistency of the Lyapunov exponents and the fractal dimension is provided by the Kaplan–Yorke conjecture.[20] For a two-dimensional mapping, the Lyapunov dimension is defined as

$$d_L = 1 + \frac{\lambda_1}{|\lambda_2|}, \tag{9}$$

and is predicted to be equal to the capacity dimension. The Lyapunov dimension is given in the last column of Table II and is seen to be in reasonable agreement with the capacity dimension given in column seven.

Before concluding this section, it is worth pointing out the value of multiple Poincaré sections for our Lyapunov exponent calculations, which were initially attempted using only a single Poincaré section. Monitoring intermediate computations and graphical displays showed that 10–20% of the trajectory points had $\delta\mathbf{u}_{j+1}$ that were widely dispersed over the Poincaré section. The local fits to determine the Jacobian matrix were noticeably poor in such cases with large deviations between the measured and fitted $\delta\mathbf{u}_{j+1}$. Several attempts were made to deal with the issue algorithmically, for example, by using smaller neighborhoods and/or including cubic and quartic terms in the fit, but the fits and resulting

Lyapunov exponents remained problematic. It was only after using 10 Poincaré sections ($\tau = T/10$) that poor Jacobian fits were eliminated and trustworthy Lyapunov exponents were obtained. Although further study is needed, it is at least apparent that having multiple Poincaré sections simplifies the algorithm.

## IV. CONCLUSIONS

This paper demonstrates the benefits of using a stepper motor running at constant angular velocity to drive a chaotic pendulum. In particular, we have shown how to use common data acquisition hardware and software to synchronize angular readings of the pendulum with the phase of the drive system. The apparatus produces phase-space data suitable for quantitative studies of nonlinear dynamics. Poincaré sections can be displayed and analyzed, and phase space mixing is strikingly demonstrated by viewing sequential Poincaré sections as a video loop.

Linear regression analysis was used to determine the parameters of the model equations of motion and verify the two-term dissipation model. Poincaré sections from simulations based on the fit parameters show remarkable agreement with those from experimental data. The calculation of fractal dimensions and Lyapunov spectra were performed using algorithms implemented in the LabVIEW programming language. Their dependence on the strength of the viscous damping coefficient was observed and found to be in reasonable agreement with theoretical predictions.

Much more can be done with the apparatus. Studies of dynamical noise in conjunction with noise reduction techniques[21] hold the promise of bringing the fractal nature of the attractor to finer scales. In this regard, it would be interesting to try to reach finer fractal scales by improvements in the apparatus, for example, better vibration isolation, smoother axle bearings, and higher rotary encoder resolution. Adding computer-controlled electromagnetic damping would make possible the study of the control of chaos.[22,23] Analysis using time delay embedding[15,24] could also be performed and comparisons between the two representations should help students gain a better understanding of both. Other potential functions could also be investigated by changing the springs, the pendulum mass, or the spring equilibrium position. One could study how the attractor changes for a shallower double-well potential or for a triple- or quadruple-well potential.

As an introductory experiment in nonlinear dynamics, this version of the chaotic pendulum provides an excellent entry point into the field. Additionally, modifications and enhancements to the hardware and software should make possible a wide array of experiments and studies suitable at the undergraduate and graduate level.

[a)]Electronic mail: deserio@phys.ufl.edu

[1]Randall D. Peters, ''Chaotic pendulum based on torsion and gravity in opposition,'' Am. J. Phys. **63**, 1128–1136 (1995).
[2]J. E. Berger, Jr. and G. Nunes, ''A mechanical duffing oscillator for the undergraduate laboratory,'' Am. J. Phys. **65**, 841–846 (1997).
[3]James A. Blackburn and Gregory L. Baker, ''A comparison of commercial chaotic pendulums,'' Am. J. Phys. **66**, 821–830 (1998).
[4]The Chaotic Pendulum and the Rotary Motion Sensor, Pasco Scientific, 10101 Foothills Blvd., Roseville, CA 95747-7100, and <www.pasco.com>.
[5]The PCI-6601 timer/counter board, CB-68LP connector block, and R6868 ribbon cable, National Instruments, 11500 N. Mopac Expwy., Austin, TX 38759-3504 and <www.ni.com>.
[6]Our stepper motor system uses a MO61-FD301 Slo-Syn stepper motor (Servo Systems, Co., 115 Main Rd., Montville, NJ 07045-0097 and <www.servosystems.com>) and an Allegro 5804 BiMOS II Unipolar Stepper-Motor Translator/Driver, Allegro MicroSystems, Inc., 115 Northeast Cutoff, Worcester, MA 01606 and <www.allegromicro.com>.
[7]The LabVIEW programming language, National Instruments, 11500 N. Mopac Expwy., Austin, TX 38759-3504 and <www.ni.com>.
[8]The University of Florida Advanced Physics Laboratory web site is at <www.phys.ufl.edu/courses/phy4803L>.
[9]The TISEAN nonlinear time series analysis package can be found at <www.mpipks-dresden.mpg.de/~tisean>.
[10]The gnuplot plotting software web site is at <www.gnuplot.info>.
[11]Abraham Savitsky and Marcel J. E. Golay, ''Smoothing and differentiation of data by simplified least squares procedures,'' Anal. Chem. **36**, 1627–1639 (1964).
[12]J. P. M. Heald and J. Stark, ''Estimation of noise levels for models of chaotic dynamical systems,'' Phys. Rev. Lett. **84**, 2366–2369 (2000).
[13]There are small corrections to $d = A\cos\phi$ that depend on the drive amplitude $A$ and the distance from the motor shaft to the guide hole (60 cm for our setup). For the largest drive amplitude used ($A = 6$ cm), there will be a 0.15 cm downward shift in the midpoint of the oscillation (from its value for $A = 0$) and an additional $\cos 2\phi$ term of amplitude 0.15 cm.
[14]G. L. Baker, J. P. Gollub, and J. A. Blackburn, ''Inverting chaos: extracting system parameters from experimental data,'' Chaos **6**, 528–533 (1996).
[15]Henry D. I. Abarbanel, Reggie Brown, John J. Sidorowich, and Lev Sh. Tsimring, ''The analysis of observed chaotic data in physical systems,'' Rev. Mod. Phys. **65**, 1331–1392 (1993).
[16]J.-P. Eckmann, S. Oliffson Kamphorst, D. Ruelle, and S. Ciliberto, ''Liapunov exponents from time series,'' Phys. Rev. A **34**, 4971–4979 (1986).
[17]$\mathbf{F}(\mathbf{u})$ would also depend on the drive phase $\phi$, and for small enough $\tau$ would be given by $\mathbf{F}(\mathbf{u}) = \mathbf{u} + \tau\mathbf{G}(\mathbf{u})$ where the $\theta$ and $\omega$ components of $\mathbf{G}$ are the right-hand sides of Eqs. (2a) and (2b), respectively.
[18]Reggie Brown, Paul Bryant, and Henry D. I. Abarbanel, ''Computing the Lyapunov spectrum of a dynamical system from an observed time series,'' Phys. Rev. A **43**, 2787–2806 (1991).
[19]G. L. Baker and J. P. Gollub, *Chaotic Dynamics: An Introduction*, 2nd ed. (Cambridge U.P., Cambridge, 1996), Chap. 5.
[20]J. L. Kaplan and J. A. Yorke, ''Chaotic behavior in multidimensional difference equations,'' Lect. Notes Math. **730**, 204–227 (1979).
[21]Eric J. Kostelich and Thomas Schreiber, ''Noise reduction in chaotic time-series data: A survey of common methods,'' Phys. Rev. E **48**, 1752–1763 (1993).
[22]C. Grebogi, E. Ott, and J. A. Yorke, ''Controlling chaos,'' Phys. Rev. Lett. **64**, 1196–1199 (1990).
[23]Gregory L. Baker, ''Control of the chaotic driven pendulum,'' Am. J. Phys. **63**, 832–838 (1995).
[24]G. L. Baker and J. P. Gollub, *Chaotic Dynamics: An Introduction*, 2nd ed. (Cambridge U.P., Cambridge, 1996), Chap. 6.