A computational approach to teaching conservative chaos

Todd Timberlake[a)]
Department of Physics, Astronomy, and Geology, Berry College, Mount Berry, Georgia 30149-5004

(Received 25 July 2003; accepted 30 April 2004)

A computational approach to teaching conservative chaos that is suitable for undergraduates in an upper-level classical mechanics course is discussed. Some ways that computation can be used to facilitate the teaching of several important topics in conservative chaos are described. © *2004 American Association of Physics Teachers.*
[DOI: 10.1119/1.1764559]

I. INTRODUCTION

There are many reasons why one might wish to teach undergraduates about the physics of chaos in conservative (Hamiltonian) systems. Chaos theory is one of the three big ideas to emerge from 20th century physics (along with quantum mechanics and relativity). Students, who may otherwise find classical mechanics a seemingly dead subject, find chaos theory exciting because of its ''newness'' and its visual appeal. The study of chaos in conservative systems is particularly important because it sheds light on the problem of quantum-classical correspondence.[1] There also are many elegant theorems (such as Liouville's theorem and the Poincaré–Birkhoff theorem) that are only applicable to conservative systems.

Despite these motivations, chaos theory is rarely taught to undergraduates. Unlike the study of quantum field theory or general relativity, the study of conservative chaos requires nothing more than calculus, differential equations, and some linear algebra. The problem with teaching chaos is that we can only study the detailed dynamics of a chaotic system computationally. In the past, this requirement meant that the physics of conservative chaos could only be taught to students who had programming skills. Now, however, there are a number of general-purpose computing programs, such as Mathematica and Maple, that can enable students with little or no computing background to quickly pick up the computational skills they need to study conservative chaos.[2]

Once the need for computation ceases to be a roadblock, it actually becomes an advantage. After seeing how their instructor carries out various computations to analyze a model system, students can do their own computational investigations of other systems. In a very short time they can be studying the dynamics of models that no one has ever studied before.[3] In addition, students learn computational skills that can help them in their other courses or research. Computation has become an important tool in nearly all subfields of physics,[4] and knowing how to use general-purpose mathematical software is a useful skill for a physicist.[5]

In this article I present several ways that computation can be used to facilitate the teaching of chaos in conservative systems. These examples are designed to be used in the second semester of a two-semester upper-level sequence in classical mechanics. Among the topics covered in the first semester of the sequence are Hamiltonian mechanics and chaos in iterated functions like the logistic map.[6] The material described in this article is presented at the end of the second semester course, after the students have had the opportunity to become more familiar with Hamiltonian mechanics. The sequence of topics that I have used follows that of Chapter 11 in Ref. 7. Other good sources of information on conservative chaos are available.[1,6–8]

II. TEACHING CONSERVATIVE CHAOS THROUGH COMPUTATION

The computational approach to teaching conservative chaos that I have employed consists of four basic steps. In the first step the instructor prepares several figures that illustrate important concepts in conservative chaos. The instructor then presents a series of lectures in which these concepts and their mathematical underpinnings are discussed and the figures are presented to illustrate key points. Next, the students are given access to the code used to create the figures and asked to solve a few computational problems. Finally, the students are asked to carry out a computational investigation of a nonlinear Hamiltonian system that they have not seen before.

In this section I will present several examples of figures that can be used as described. Although the significance of each figure will be discussed, readers unfamiliar with conservative chaos may need to consult Refs. 1, 6, 7, or 8 for a more detailed discussion of the topics and terminology. The Mathematica code used to create these figures, along with a description of the calculations, can be found at Ref. 9.

The figures are part of an investigation of the standard map.[10] Two-dimensional area-preserving maps provide a convenient way to introduce students to chaos in conservative systems. These maps can display most of the main features of conservative chaos. The dynamics of these maps can be easily visualized with two-dimensional plots. In addition, some of the calculations relevant to conservative chaos, such as finding the tangent map for a fixed point can be carried out analytically.

The standard map serves as an approximation for several physical systems, most notably the kicked rotor. The standard map can be written as[6]

$$r_{n+1}=r_n-\frac{K}{2\pi}\sin(2\pi\theta_n), \quad \mathrm{mod}\ 1, \tag{1a}$$

$$\theta_{n+1}=\theta_n+r_{n+1}, \quad \mathrm{mod}\ 1. \tag{1b}$$

The variables r and θ are dimensionless phase space variables. Both coordinates are periodic, so for example $\theta=1$ is equivalent to $\theta=0$. It is convenient to think of r as an action variable and θ as the corresponding angle. In some situations it may be useful to think of r and θ as dimensionless polar coordinates. K is a dimensionless parameter that controls the nonlinearity of the system. For $K=0$ the system is integrable.

A. Surfaces of section

The dynamics of a Hamiltonian system takes place in a $2N$-dimensional phase space where N is the number of degrees of freedom. The dynamics is therefore difficult to visualize even for two degrees of freedom. The best way to obtain a visual picture of the dynamics is to construct a surface of section. To construct such a surface, we calculate the trajectories starting from several initial conditions in various parts of the phase space. Instead of trying to view these trajectories in the full $2N$-dimensional phase space, we only examine the points at which these trajectories intersect a two-dimensional plane. As a single trajectory moves through the phase space, it repeatedly intersects this plane, producing a collection of points that provide a visual picture of the trajectory's behavior. Because trajectories cannot intersect each other in Hamiltonian systems, each point on a surface of section is associated with only one trajectory in the full phase space. Therefore a surface of section can be thought of as a two-dimensional map that maps a point at which a trajectory intersects the plane to the next point at which the trajectory intersects the plane.

A two-dimensional, area-preserving map like the standard map can be thought of as generating the surface of section of some physical system. The surface of section is produced by choosing several initial points in phase space and then plotting the sequence of points that result from repeatedly applying the map to each initial point. The set of points that arise from a single initial point can be referred to as a trajectory of the map. If the initial points are distributed throughout phase space, the surface of section will provide an overview of the dynamics of the system. The initial points may be chosen so that they lie on a regular grid in the phase space, or they may be chosen at random. If there are noticeable gaps in the surface of section, then additional initial points can be chosen to fill in the gaps.

For $K=0$ the standard map is integrable. The surface of section for $K=0$ consists of trajectories crossing horizontally from $\theta=0$ to $\theta=1$ at constant values of r. Because θ is a periodic coordinate, each of these trajectories really moves along a circle (or 1-torus). For higher dimensional integrable systems, the trajectories also are confined to (higher dimensional) tori. These tori, on which the integrable trajectories are constrained to move, are known as *invariant tori*.

Figure 1 shows a surface of section for the standard map with $K=0.8$, for which the map is no longer integrable. The surface of section reveals many important features of the map's dynamics. For example, the group of nested elliptical curves surrounding the origin is a structure known as a nonlinear resonance. For two-dimensional (or higher) systems nonlinear resonances occur when the frequency of motion in one dimension is rationally related to the frequency of motion in a different dimension. In the case of a one-dimensional map, nonlinear resonances occur when the frequency of motion is rationally related to the map frequency. For the standard map the point at the origin is mapped back to the origin after each application of the map and is therefore called a fixed point. Nearby points wind around this periodic trajectory as the map is repeatedly applied, forming the elliptical curves that comprise the nonlinear resonance (or resonance island). There also are a number of continuous curves that run horizontally from $\theta=0$ to $\theta=1$. These structures are the remaining invariant tori (or KAM tori after Kolmogorov, Arnol'd, and Moser[11]) that have become distorted by the increase in K, but have not been destroyed. In addition to resonance islands and KAM tori, there also is a disorganized scatter of points in the vicinity of (0.5,0). The motion in this region of the phase space is chaotic.

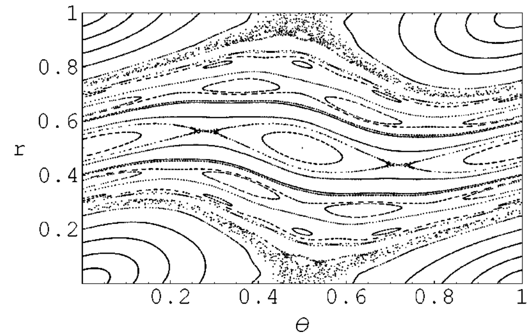

Fig. 1. Surface of section for the standard map with $K=0.8$. Note that the coordinates r and θ are periodic. A large resonance (elliptical island) is visible at the origin. Several higher-order resonances or island chains also are visible. Continuous curves running across the phase space are KAM tori. Chaotic motion is visible in the region of the separatrix around the large resonance island.

By investigating how the surface of section changes as K is varied, we can see that the resonance islands grow as K increases. Eventually these resonances overlap and the KAM tori between them are broken. The overlap of nonlinear resonances leads to chaotic motion because the points in the region of overlap don't know which periodic trajectory to wind around. This chaotic motion always begins in the vicinity of the separatrix, which is the boundary between a nonlinear resonance and the KAM tori that lie just outside. For example, the region of chaotic motion (sometimes called the stochastic layer) in Fig. 1 lies along the boundary between the large resonance island and the KAM tori. Trajectories that start out anywhere in the chaotic region can wander, or diffuse, throughout that region. For sufficiently large K, all KAM tori will be destroyed and the entire phase space will be chaotic.

We also can use surfaces of section to explore the fractal nature of a mixed phase space (one that contains both chaotic and nonchaotic regions) by zooming in on the edge of one of the resonance islands to see that each resonance is typically surrounded by smaller daughter resonances, which are in turn surrounded by even smaller resonances, and so on. This fractal structure illustrates that nonlinear resonances occur on all scales in a chaotic Hamiltonian system.

Problem 1. Create surface of section plots for the standard map with $K=0$, $K=1$, and $K=7$. Describe the dynamics of the map for each value of K.

Problem 2. Determine the approximate value of K for which all resonances other than the one centered at the origin have been destroyed. Determine the approximate value of K for which the resonance centered at the origin is destroyed (at this value of K the phase space will be almost entirely chaotic).

B. Individual trajectories

Plots of individual trajectories also can help us visualize certain aspects of conservative chaos. By choosing the

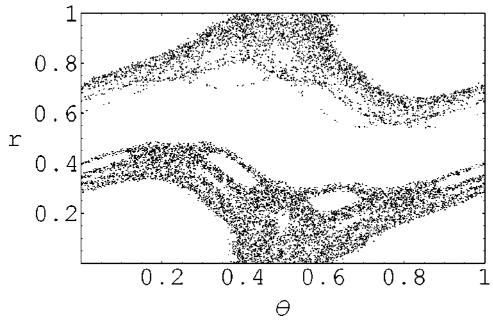

Fig. 2. A portion of a chaotic trajectory for the standard map with $K = 1.13$. The initial point is (0.5, 0.0001). The first 10 000 iterations of the map are shown. Continued iteration would show further diffusion in r, because at this value of K all KAM tori are broken.

proper initial conditions for a given value of K, one can easily create plots of periodic orbits, KAM tori, trajectories trapped in resonance regions, or chaotic trajectories. Simply examine the surface of section, choose an initial point that lies in the region to be examined, and plot the points that result from repeated iterations of the map. Plots of chaotic trajectories can be particularly useful for identifying barriers to chaotic diffusion in phase space. If, for a given value of K, KAM tori still exist in phase space, these structures will block the diffusion of a chaotic trajectory in the r direction. For sufficiently high values of K, all KAM tori will be broken and chaotic trajectories can diffuse freely in r. The last KAM torus to be destroyed will be the one whose winding number (essentially the amount that θ changes with each iteration of the map) is the reciprocal of the golden ratio. Nonlinear resonances form where trajectories have rational winding numbers. The golden ratio and its reciprocal are among the irrational numbers that are most difficult to approximate with rational numbers (because their continued fraction expansions converge slowly), and thus a KAM torus with this winding number will be far from any nonlinear resonance and the last to be destroyed by resonance overlap.[1,8]

Figure 2 shows the first 10 000 iterations of a chaotic trajectory for $K=1.13$ starting at the point (0.5,0.0001). This trajectory appears to be bounded in r, implying that KAM tori still exist for this value of K. However, this picture is somewhat misleading because even broken KAM tori (sometimes called cantori because of their Cantor set-like structure) can provide a partial barrier to chaotic diffusion. The presence of cantori prevents this trajectory from diffusing through all values of r during the first 10 000 iterations of the map. If the first 100 000 iterations were shown, we would see that the trajectory diffuses through the apparent boundaries, indicating that the apparent boundaries are really only partial barriers to diffusion. There are no unbroken KAM tori for values of $K > K^* = 0.971\,6354$.[1]

Problem 3. Choose the initial point (0.6,0.1) and estimate the value of K^* by finding the smallest value of K for which the diffusion of the trajectory in r is no longer blocked during the first 10 000 iterations of the map.

C. Visualizing Liouville flow

Plotting the evolution of a group of points in phase space can help us visualize Liouville's theorem. Liouville's theorem states that the area of any region of phase space is preserved under the application of a Hamiltonian map (or under time evolution in a continuous Hamiltonian system). Mathematically this area preservation occurs because the Jacobian of a Hamiltonian map, defined by

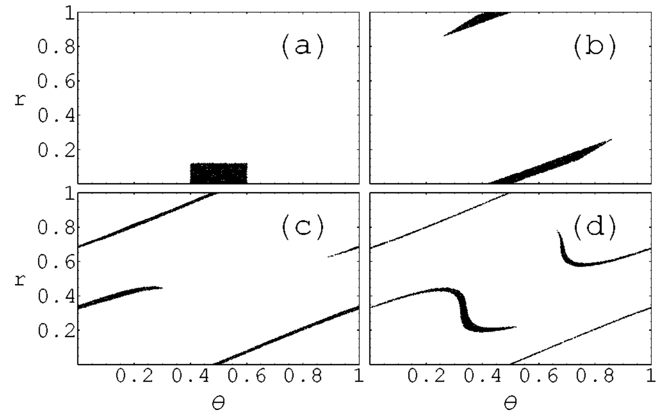

Fig. 3. An illustration of phase space flow and Liouville's theorem. Plot (a) shows 5000 initial points distributed randomly in a rectangular region of phase space. The remaining plots show the distribution of trajectories after (b) one, (c) two, and (d) three iterations of the map with $K=1.5$.

$$J = \begin{pmatrix} \dfrac{\partial \theta_{n+1}}{\partial \theta_n} & \dfrac{\partial \theta_{n+1}}{\partial r_n} \\ \dfrac{\partial r_{n+1}}{\partial \theta_n} & \dfrac{\partial r_{n+1}}{\partial r_n} \end{pmatrix}, \tag{2}$$

has determinant equal to one. The Jacobian for a continuous system can be defined in a similar way.[7]

A proper understanding of Liouville's theorem involves visualizing classical dynamics as a flow of points in phase space. Points flow from one location in phase space to another under application of the map (or under time evolution). This visual picture can be aided by diagrams that illustrate the flow of a large number of initially nearby points. Figure 3 illustrates this phase space flow and the preservation of area required by Liouville's theorem. Figure 3(a) shows the initial conditions: five thousand points randomly selected inside a rectangular region of the phase space. Figures 3(b)–3(d) show the result of applying the map (for $K=1.5$) repeatedly on this set of points. These diagrams illustrate that although the region occupied by the points is stretched and bent, its area remains constant. Thus, the flow of points in a Hamiltonian system is like the flow of an incompressible fluid. An animation composed of several of these plots shows how the group of points evolves under iterations of the map and provides an excellent way to visualize the phase space flow.[12]

Problem 4. Determine the Jacobian for the standard map and show that its determinant is equal to one for all values of K.

Problem 5. Examine the flow of points that lie inside a nonlinear resonance. Contrast this behavior with the flow of points in a chaotic region of the phase space.

D. Behavior of trajectories near a fixed point

Fixed points (which are mapped back to themselves under one iteration of the map) and periodic points (which are mapped back to themselves after several iterations of the map) play an important role in determining the overall dynamics of a Hamiltonian map. According to the Poincaré–Birkhoff theorem,[13] when K is increased from zero to some small but finite value, trajectories with rational winding numbers will form into a sequence of alternating stable (elliptic) and unstable (hyperbolic) periodic points. Periodic points are stable if nearby trajectories remain close and unstable if nearby trajectories move away. The stable periodic points will become the centers of nonlinear resonance islands. The unstable periodic points provide the seeds for chaotic motion (see the discussion of homoclinic tangles in the following). Thus, the stability of a given periodic point is crucial for determining the character of the dynamics in the region of phase space that surrounds the point.

Here we will focus on analyzing the stability of fixed points, but the analysis can be extended to deal with periodic points. The best way to analyze the stability of a fixed point is to construct the stability matrix $\mathbf{P}$ for the point. The stability matrix is a linear map that describes the local dynamics in the vicinity of the fixed point. If the eigenvalues of the stability matrix are complex with unit modulus, the fixed point is stable and nearby trajectories will remain nearby. But if one eigenvalue has absolute value greater than one, the fixed point is unstable and nearby trajectories will move away from the fixed point along a certain axis. For a two-dimensional map the stability matrix is the Jacobian, Eq. (2), evaluated at the fixed point. The eigenvalues of a 2×2 matrix $\mathbf{P}$ are given by $\alpha=[\mathrm{Tr}\,\mathbf{P}\pm\sqrt{(\mathrm{Tr}\,\mathbf{P})^2-4(\mathrm{Det}\,\mathbf{P})}]/2$. Because the map is area preserving, the determinant of the stability matrix is unity and the eigenvalues must come in pairs of the form α and $1/\alpha$. If $|\mathrm{Tr}\,\mathbf{P}|<2$, then the eigenvalues of the stability matrix are complex conjugates with unit modulus and the fixed point is stable. If $|\mathrm{Tr}\,\mathbf{P}|>2$, then the eigenvalues of the stability matrix are real numbers of the form α and $1/\alpha$ with $|\alpha|>1$ and the fixed point is unstable.

The eigenvectors of the stability matrix for an unstable fixed point play an important role in determining the behavior of trajectories in the vicinity of the point. The eigenvector associated with the eigenvalue α indicates the direction along which nearby points move away from the fixed point (the unstable direction). The eigenvector associated with $1/\alpha$ gives the direction along which nearby points move toward the fixed point (the stable direction). This effect can be demonstrated in a way that is similar to the illustration of Liouville's theorem. Begin with a large number of points randomly distributed in the vicinity of the fixed point. As the map is repeatedly applied to these points, the distribution will become compressed along the stable direction and stretched out along the unstable direction, making the unstable direction easy to see. If we apply the inverse map

$$\theta_n=\theta_{n+1}-r_{n+1}, \quad \mathrm{mod}\ 1, \tag{3a}$$

$$r_n=r_{n+1}+\frac{K}{2\pi}\sin(2\pi\theta_n), \quad \mathrm{mod}\ 1, \tag{3b}$$

rather than the forward map, then the distribution will become stretched out along the stable direction, making the stable direction easy to see.

Figure 4 illustrates this effect. The plots were generated using 5000 initial points distributed randomly in a rectangular region around the unstable fixed point (0.5,0) for $K=1.5$. The points that lie along the line of positive slope show the result of applying the forward map four times. The points that lie along the line of negative slope are the result of applying the inverse map four times. The flow of points can be clearly demonstrated by animating these plots.[12]

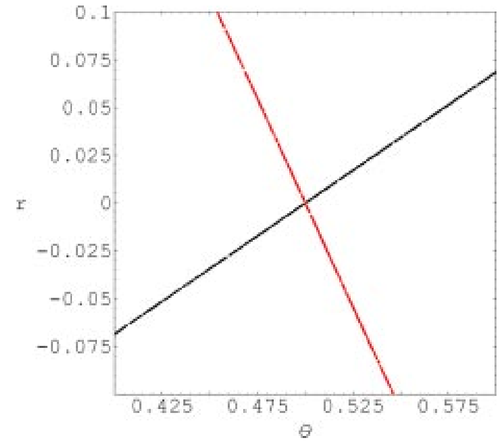

Fig. 4. Stable and unstable directions for the fixed point at (0.5,0) for $K=1.5$. The forward and inverse maps were applied to 5000 points, initially distributed randomly in a rectangular region surrounding the fixed point. After four iterations of the forward map, the points lie along the unstable direction (the line with positive slope). After four iterations of the inverse map, the points lie along the stable direction (the line with negative slope). The modulo 1 restriction in Eq. (1) has been removed to make the results easier to see, because r is a periodic coordinate and a value of r slightly less than zero is equivalent to a value of r slightly less than one.

Another way to visualize the behavior of trajectories in the vicinity of an unstable fixed point is to create a flow diagram. A flow diagram shows a vector field that indicates the direction each point will move when the map is applied. Figure 5 shows a flow diagram around the fixed point (0.5,0) for $K=1.5$. Note that the arrows point out along the unstable direction and in along the stable direction of Fig. 4.

Problem 6. Evaluate the stability matrix (Jacobian) for the fixed point of the standard map at (0.5,0). Determine the values of K for which this fixed point is stable or unstable. Find the eigenvalues and eigenvectors of the stability matrix for $K=1.5$. Show that the eigenvector with eigenvalue

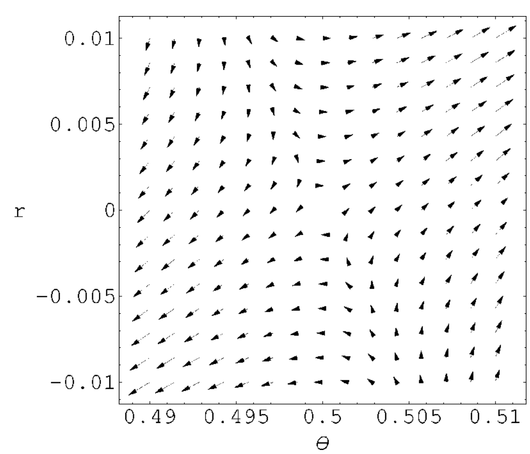

Fig. 5. Flow diagram in the vicinity of the fixed point at (0.5,0) for $K=1.5$. The arrows indicate that nearby points flow away from the fixed point along the unstable direction from Fig. 4 and flow toward the fixed point along the stable direction from Fig. 4. The modulo 1 restriction in Eq. (1) has been removed as in Fig. 4.

greater than one points along the unstable direction of Fig. 4 and that the eigenvector with eigenvalue less than one points along the stable direction.

Problem 7. Evaluate the stability matrix for the fixed point of the standard map at (0,0). Determine the values of K for which this fixed point is stable or unstable. Investigate what happens in the surface of section when this fixed point becomes unstable.

E. Homoclinic tangles

The idea of identifying stable and unstable directions associated with an unstable fixed point can be extended. Imagine constructing the set of all points in the phase space that, if acted on repeatedly with the inverse map, would eventually move in toward the unstable fixed point along the direction of the unstable eigenvector of the stability matrix. This structure is called the unstable manifold of the fixed point and the manifold as a whole is invariant under application of the map. Close to the fixed point itself this structure is a straight line in the direction of the unstable eigenvector. Farther from the fixed point this line will begin to curve and turn, but it will never cross itself. Similarly, we can construct the stable manifold of the fixed point (the set of points that are eventually mapped in along the stable eigenvector by the forward map). The stable manifold will be a straight line along the stable eigenvector in the vicinity of the fixed point, but will curve when it is farther away. In an integrable system, the unstable manifold of a particular fixed point will connect smoothly with the stable manifold of the same, or of a different, unstable fixed point. These smoothly connecting manifolds form a separatrix. In a chaotic system these manifolds do not connect smoothly. In fact, they don't connect at all. Instead they cross each other an infinite number of times (while never crossing themselves). If the stable and unstable manifolds emanate from the same unstable fixed point, then this structure is called a homoclinic tangle. If they are from different fixed points, it is called a heteroclinic tangle. The intersection points are called homoclinic (or heteroclinic) points. Each homoclinic (or heteroclinic) point is mapped to another such point.

We can easily imagine that motion in the vicinity of a tangle will be quite complex. The stable and unstable manifolds may occupy only a small region of the phase space. Yet in this small region these manifolds must cross each other an infinite number of times without crossing themselves. To accomplish this feat, the manifolds must make tighter and tighter turns as they twist back and forth in the phase space. This complex web of intersections between the two manifolds gives rise to chaotic motion in the region occupied by the tangle.

We can gain some appreciation for the complexity of a tangle by plotting a portion of the stable and unstable manifolds associated with an unstable fixed point. Figure 6 shows a portion of the stable and unstable manifolds associated with the fixed point of the standard map at (0.5,0) with $K=1.5$. A plot of the unstable manifold can be constructed by taking several initial points that are displaced from the fixed point along the direction of the unstable eigenvector of the stability matrix but at random (though small) distances. Applying the forward map to these points will cause them to spread out along the unstable manifold. Successive iterations of the map will generate a more complete picture of the unstable manifold. A plot of the stable manifold can be constructed in an analogous way, with points distributed along the direction of the stable eigenvector and using the inverse map. The combination of these two plots gives a partial picture of the homoclinic tangle. The complete tangle is an infinitely complicated structure and can never be viewed in its entirety. However, the more iterations of the map one uses to construct the tangle, the more of its complexity can be seen. An animation showing successive iterations of this process can effectively illustrate the increasing complexity of the homoclinic tangle. Imagining the extrapolation of this process to an infinite number of map iterations allows us to visualize the complexity of the homoclinic tangle that leads to chaotic motion.

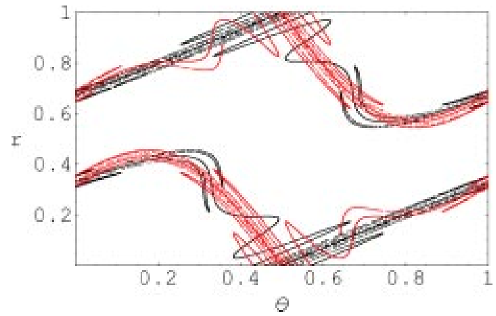

Fig. 6. A portion of the homoclinic tangle surrounding the unstable fixed point of the standard map at (0.5,0) with $K=1.5$. The complete homoclinic tangle is infinitely more complicated than the curve shown here.

Problem 8. Create a plot of the homoclinic tangle for the fixed point (0.5,0) with $K=0.8$. How is this homoclinic tangle different from the one shown in Fig. 6? Describe what you would expect to see if you created a plot of the homoclinic tangle for the same fixed point with $K=3$. Test your hypothesis.

F. Exponential divergence of trajectories and Lyapunov exponents

Because unstable fixed points play such an important role in determining the dynamics of the system, it would be useful to have a way of quantifying their instability. The eigenvalue, α, of the stability matrix with absolute value greater than one can play this role, but another useful measure of the instability of a fixed point is the Lyapunov exponent. The Lyapunov exponent is defined as $\lambda=\ln|\alpha|$. Trajectories that start close to a typical unstable fixed point will move away from the fixed point exponentially. The distance between the two points after n iterations of the map is $d_n \approx d_0 e^{\lambda n}$. Stable fixed points (with eigenvalues of unit modulus) have $\lambda=0$. The exponential divergence of nearby trajectories is illustrated in Fig. 7. The plot shows the natural log of the distance, $\ln d_n$, between two trajectories as a function of n with $K=1.5$. One trajectory is the fixed point at (0.5,0) while the other trajectory begins at $(0.5+10^{-8},0)$ and moves away from the fixed point. The linear increase of $\ln d_n$ for $n<17$ indicates that the distance between the two trajectories is increasing exponentially. For $n>17$ the second trajectory begins to move closer to the fixed point, indicating that the exponential divergence does not continue indefinitely. The exponential divergence may cease because the trajectory has run up against barriers in phase space, or it may simply be that the trajectory has moved far enough from the unstable fixed point so as to be no longer repelled.

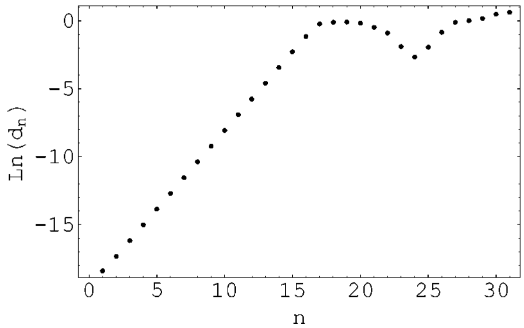

Fig. 7. An illustration of the exponential divergence of neighboring trajectories. The plot shows the natural log of the distance between the fixed point at (0.5,0) and an initially nearby ($\delta\theta = 10^{-8}$ and $\delta r = 0$) trajectory as a function of the number of map iterations for $K = 1.5$. Note that the curve increases linearly until $n = 17$, indicating that the distance between the two trajectories is increasing exponentially. The modulo 1 restriction of Eq. (1) has been removed to allow for the accurate calculation of distances.

Problem 9. Find the Lyapunov exponent for the fixed point at (0.5,0) with $K = 1.5$. Show that it is approximately equal to the slope of the line on which the points in Fig. 7 fall (for $n < 17$). Then find the Lyapunov exponent for this fixed point with $K = 0.8$. How is the Lyapunov exponent related to the size of the homoclinic tangle (from Problem 8)?

III. CONCLUDING REMARKS

I have presented a number of ways in which computation can facilitate the teaching of conservative chaos to undergraduates. By combining numerical exploration with analytical calculations, students can gain a deeper understanding of chaotic dynamics in conservative systems. Students can be assigned problems, similar to those presented in this article, as the material is covered in class. After they have completed these assignments, students will be able to take on larger, independent projects. For instance, they may be asked to investigate the dynamics of a new map using the tools they have learned in studying the standard map. Other two-dimensional area-preserving maps can be found in the recommended texts.[1,6,8] I have taught this material in the second semester of a two semester upper-level classical mechanics sequence, but one could add to this material to create a course for undergraduate that focuses entirely on nonlinear dynamics. Such a course could include material on one-dimensional maps (iterated functions), continuous systems, and dissipative chaos. Information on these topics can be found in Refs. 1, 6–8.

ACKNOWLEDGMENT

The author would like to thank Dr. Charles Lane for his helpful suggestions in the preparation of this article.

[a]Electronic mail: ttimberlake@berry.edu

[1]Linda E. Reichl, *The Transition to Chaos in Conservative Classical Systems: Quantum Manifestations* (Springer-Verlag, New York, 1992).

[2]Charles H. Patterson, ‘‘Two approaches to teaching computational physics,’’ Comput. Sci. Eng. **4**(6), 64–68 (2002).

[3]Todd Timberlake, Justin Harvey, Matt Lewis, and Wesley Taylor, ‘‘Beyond the standard map: A numerical investigation of indifferent fixed points’’ (unpublished).

[4]Harvey Gould, ‘‘Computational physics and the undergraduate curriculum,’’ Comput. Phys. Commun. **127**, 6–10 (2000).

[5]Robert L. Zimmerman and Fredrick I. Olness, *Mathematica For Physics* (Addison Wesley, San Francisco, 2002); Patrick T. Tam, *A Physicist's Guide to Mathematica* (Academic Press, San Diego, 1997); Marvin L. De Jong, *Mathematica for Calculus-Based Physics* (Addison-Wesley, San Francisco, 1999); Ronald Greene, *Classical Mechanics With Maple* (Springer-Verlag, New York, 1995).

[6]Robert C. Hilborn, *Chaos and Nonlinear Dynamics*, 2nd ed. (Oxford University Press, Oxford, U.K., 2000).

[7]Louis N. Hand and Janet D. Finch, *Analytical Mechanics* (Cambridge University Press, Cambridge, U.K., 1998).

[8]Michael Tabor, *Chaos and Integrability in Nonlinear Dynamics* (Wiley, New York, 1989).

[9]See EPAPS Document No. E-AJPIAS-72-011408 for a Mathematica notebook containing the code used to produce the figures in this article which can be found at ⟨http://fsweb.berry.edu/academic/mans/ttimberlake/cc/⟩. Other materials related to teaching Hamiltonian chaos, including lecture slides and examples of student work, can also be found at this site. A direct link to this document may be found in the online article's HTML reference section. The document may also be reached via the EPAPS homepage (http://www.aip.org/pubservs/epaps.html) or from ftp.aip.org in the directory /epaps. See the EPAPS homepage for more information.

[10]Boris V. Chirikov, ‘‘A universal instability of many dimensional oscillator systems,’’ Phys. Rep. **52**, 263–379 (1979).

[11]Louis N. Hand and Janet D. Finch, *Analytical Mechanics* (Cambridge University Press, Cambridge, U.K., 1998), pp. 433–436.

[12]All of the animations mentioned in the article can be found at the website mentioned in Ref. 9.

[13]Louis N. Hand and Janet D. Finch, *Analytical Mechanics* (Cambridge University Press, Cambridge, U.K., 1998), pp. 439–442.