

The Forced Damped Pendulum: Chaos, Complication and Control.

John H. Hubbard

We show that a “simple” differential equation modeling a garden-variety damped forced pendulum can exhibit extraordinarily complicated and unstable behavior. While instability and control might at first glance appear contradictory, we can use the pendulum’s instability to control it. Such results are vital in robotics: the forced pendulum is a basic subsystem of any robot.

Most of the mathematical methods used in this paper were initially developed in celestial mechanics, largely by Poincaré. The literature of the field tends to be quite advanced indeed (see [1, 11]); one object of this paper is to show that computer programs, properly used, can make these advanced topics transparent. All the computer-generated pictures in this paper were produced by the programs *Planar Systems* and *Planar Iterations* [6], both written by Ben Hinkle (now at *Maple*).

1. Some parallels in celestial mechanics

When I was a graduate student, I was amazed by the results of Alekseev concerning a system formed by three bodies obeying Newton’s law of gravitation; see [1] and [11]. As shown in Figure 1, two massive bodies of equal mass move in a plane P on ellipses symmetric around a common focus F , and the third body, the *satellite*, of mass zero, moves on the line L perpendicular to P through F . Once this satellite is launched, its motions are determined uniquely by the gravitational pull of the two massive bodies.

The system has a natural unit of time, the “year”—the time it takes the massive bodies to complete a revolution. Choose a time zero, so that it makes sense to speak of the 0th, 1st, . . . , n th year. Also let x denote the position on the line L , with $x = 0$ corresponding to F .




FIGURE 1. Alekseev’s three-body system.

Alekseev proved that there then exists a number N , which depends on the eccentricity of the orbits of the large bodies, such that given *any* sequence n_1, n_2, \dots of integers at least N , there exists a set of initial conditions that results in the satellite returning to cross the plane P exactly in the n_1 th year, the $(n_1 + n_2)$ th year, etc. In other words, given a specified sequence of years with gaps at least N , it is possible to choose an instant t_0 and a speed $v = x'(t_0)$ so that if the satellite is kicked off at that moment with that speed, it crosses the plane during the desired years: first during the n_1 th year, then n_2 years later, and so on. You can set up the satellite to return in *any* sequence of years you like, so long as the returns are spaced at least N apart.

In particular, there exist unbounded orbits in which the satellite travels arbitrarily far away but always returns, for example the orbit corresponding to the sequence of gaps between crossings $N, N + 1, N + 2, N + 3, \dots$ as well as infinitely many different periodic orbits (for instance $N, N + 12, N + 17, N, N + 12, N + 17, \dots$).

Actually, Alekseev claimed the result only when the eccentricity is “sufficiently small.” He needed to know that his system satisfied some requirements (basically, that a “horseshoe” should be present), and he could verify this only by a perturbation calculation near an explicitly integrable system. Horseshoes are discussed in Section 8.

The pendulum model we explore here exhibits a similar sort of behavior: we can make our pendulum go through any specified sequence of gyrations by correctly choosing the initial conditions. More precisely, by appropriately choosing the position and the velocity of the pendulum at time 0, we can specify whether during each time period (the time period of the forcing term, in our case, 2π) the pendulum goes through the bottom position once clockwise, once counterclockwise, or not at all. For example, we could specify that in each of the first six periods it could go through the bottom position once clockwise, in each of the next three periods it could go through the bottom position once counterclockwise, and in the tenth period oscillate around an upright position All imaginable sequences are possible: once the correct set of initial conditions is chosen, the differential equation governing the system automatically enforces the desired behavior.

2. Differential equations and pendulums

There is only one law in mechanics: $F = ma$ (force equals mass times acceleration). Thus the motion of a pendulum of length l , with a bob of mass m in a constant gravitational field of force g , with friction proportional to the velocity, and forcing $f(t)$ (Figure 2) is modeled by the differential equation

$$\underbrace{f(t) - \gamma l x' - mg \sin(x)}_{\text{force}} = \underbrace{m}_{\text{mass}} \times \underbrace{lx''}_{\text{acceleration}} .$$

The friction term $\gamma l x'$ is a fairly good approximation to reality when the friction is due to air, and the speed of the bob is much less than the speed of sound. The term $mg \sin(x)$ is the force exerted by gravity; the weight of the body is mg , but only the component in the direction of motion contributes to the equation. The forcing $f(t)$ can be created by a current proportional to $f(t)$ through the axis of the pendulum, if the bob is a bar magnet

perpendicular to the axis. In realistic situations (e.g., robot arms), this is the way forcing is really produced.




FIGURE 2. A pendulum being driven by alternating current

We explore the behavior of a pendulum whose motions are described by the particular differential equation

$$\cos(t) - 0.1x' - \sin(x) = x'',$$

in which both mass m and length l equal 1.

My starting point was the observation by Borelli and Coleman [3] that numerical solutions of this equation are very sensitive to the integration method, step-length, etc., near the initial condition $(x(0), x'(0)) = (0, 2)$. That is, we start with a pendulum hanging down, and hit it with a mallet to give it velocity near 2. This paper is my attempt to understand this instability. The behavior I describe holds not just for the parameters $m, \gamma, l, g, f(t)$ given; they could be varied in a certain range, which I don't know in any detail, but which is large enough so that it would not be difficult to build a real system that behaves like the one described here.

3. A first attempt to understand the motions of the pendulum

The most obvious thing to ask a computer is: what do the motions of the pendulum look like? The following picture shows the motion resulting from 15 different sets of initial conditions. Each graph starts with the position $x(0) = 0$; the initial velocities are evenly spaced between 1.85 and 2.1. The graphs are plotted for $-1 < t < 200$ and $-25 < x < 25$. A word of caution: the overall features of Figure 3 are correct, but the details—exactly which equilibrium each initial condition leads to—might well be wrong. The exponential growth of errors is discussed in Section 11.




FIGURE 3. Fifteen solutions to the differential equation $\cos(t) - 0.1x' - \sin(x) = x''$.

A careful look at the picture suggests that there exists a stable periodic motion $S(t)$ of the pendulum, which you see in the picture many times; of course, $S(t) + 2k\pi$ is another description of the same motion for any integer k ; the letter S stands for “stable.” You will see five different levels of this stable periodic motion: one on the horizontal axis, three above, and one below. The first stable motion above the horizontal axis represents motions that go “over the top” once counterclockwise before settling down, like a child’s swing going over the bar. The next layer up represents motions that go over the top twice counterclockwise before settling down, while the layer below the horizontal axis represents motions that go over the top once clockwise before settling down.

Some motions rapidly settle down to this oscillation, others go through a complicated path before doing so, and yet others do not approach the periodic motion in this amount of time. These appear to be rare, and one might guess that given more time, almost all solutions do settle down. (One that does not is shown in [13, p. 228]; the existence of uncountably many others is proved in Theorem 3.)

An obvious question is: what stable oscillation—what attracting periodic solution—can a motion approach? This seems impossible to understand without another program.

4. The scanning picture

We now look at the whole family of initial conditions: position represented by the horizontal axis, velocity by the vertical axis. We ask the computer to color initial conditions according to the stable oscillation the corresponding solution approaches (if any). This set of initial conditions is called the *basin* of the corresponding *sink*; it is an open subset of \mathbb{R}^2 .

This is best done as follows. First, find the initial values $S_0(0), S'_0(0)$ for one of the attracting periodic solutions, say the one with $-2\pi < S_0(0) < 0$. We call the motion immediately above it S_1 , and the one above that S_2 ; we have $S_k(t) = S_0(t) + 2k\pi$. Next, find a number $r > 0$ such that if

$$|x(0) - S_0(0)|^2 + |x'(0) - S'_0(0)|^2 < r^2,$$

then the motion $x(t)$ is definitely attracted to S_0 . That is, any set of initial values inside the circle of radius r and centered at $(S_0(0), S'_0(0))$, gets arbitrarily close to the solution S_0 (in fact, does so exponentially fast). We rely on computer calculations to determine this, but it would not be hard to provide a rigorous mathematical justification. We are not particularly interested in the points inside that circle; we are just establishing how we know that a motion is attracted to a particular attracting solution: it is attracted to it if it ever enters the circle of radius r around the solution. In our case, we have

$$(S_0(0), S'_0(0)) \approx (-2.0463, .3927) \quad \text{and we can take } r = 0.1.$$

Now we solve the differential equation starting at every point of some grid (in our case, a 600×400 grid—240,000 points!), and sample the solution at times $2\pi, 4\pi, \dots$: this is a substantial computation, taking about two hours even on a fairly fast Mac (200 MHz).

If for some such motion $w(t)$ and some integer $n > 0$ we have

$$|w(2n\pi) - S_k(0)|^2 + |w'(2n\pi) - S'_k(0)|^2 < r^2,$$

we know that this motion is attracted to S_k . Color the point $(w(0), w'(0))$ in the k th color and solve the differential equation for the next point. If after some number of samplings (in our case 30: we integrated solutions for time $60\pi \approx 185$) the solution never falls within r of an attracting solution, leave the initial point white. We obtain Figures 4 and 5.




FIGURE 4. The different colors (hard to appreciate in black and white) represent different basins: which initial conditions are attracted to which sinks. Points colored white may be initial conditions that are never attracted to a sink, but more likely they are attracted to sinks that are off the picture. They could also be attracted to sinks in the picture, but not during the time allowed.




FIGURE 5. In black and white, the four basins of Figure 4 are hard to distinguish. This figure represents just one basin.

5. Lakes of Wada

The colored sets B_k (called, for obvious reasons, the *basins* of the corresponding attracting motions S_k) are immensely complicated.

We show that they form infinitely many Lakes of Wada. Wada was a Japanese mathematician who at the beginning of the 20th century constructed an example of three disjoint, connected open subsets of the unit disc $D \subset \mathbb{R}^2$ such that every point in the boundary of one is in the boundary of the other two [15]. This amazed the mathematical community at the time: if you try to draw three (connected, open) lakes in an island, you would probably soon convince yourself that all three can touch at only two points. Actually, it appears that Brouwer discovered this phenomenon earlier [4].

Let me sketch the construction as outlined in [15], illustrating the dangers of philanthropy; this is illustrated by Figure 6.

Suppose D is an island cursed with three philanthropists, one of whom wants to bring water to every inhabitant, one tea, and one coffee. At the beginning each has a pond of his own beverage.

First, the purveyor of water digs a system of canals emanating from his pond, and bringing water within 100 meters of every inhabitant, never actually touching the surrounding sea or the other ponds, and forming no loops.

Next, the purveyor of coffee builds a system of canals emanating from his pond, bringing coffee to within 10 meters of every inhabitant, again forming no loops. Since the water canals make no loops, they don't cut off any inhabitants from the coffee pond, so this is possible.




FIGURE 6. Digging the lakes of Wada

Now the purveyor of tea builds his system of canals, bringing tea to within 1 meter of every inhabitant. Next the water purveyor goes back to work, extending his canals (necessarily building narrower ones) to bring water within 10 cm of each inhabitant. And so forth. At the end of this process, the poor inhabitants no longer have any dry land to stand on, but they have water, tea, and coffee as close as they want. What remains of the dry land is in the boundary of all three basins.

Real philanthropists don't seem to behave this way, fortunately. Highway designers, on the other hand . . .

Theorem 1 shows that our pendulum is creating lakes of Wada.

Theorem 1. *The basins B_k have the Wada property: every point in the boundary of one basin is also in the boundary of all the others.*

This is not quite as strong as the preceding statement about philanthropists, where every bit of dry land was in the boundary of all the basins. For the pendulum, all we can prove is that *if* a point is in the boundary of one basin, it's in the boundary of the others. Presumably there is no other dry land, but we don't know how to prove it. True lakes of Wada have been proven to exist in another setting of dynamical systems [7].

The first step in understanding why Theorem 1 is true is to get a grasp on the boundaries of the basins. Most of the material in the next section was developed by Kennedy, Nusse and Yorke; see [9] and [12]. They saw that the basin of a sink often has *saddle points* on its boundary, and that the *stable separatrices* of these saddle points make up the *accessible boundary* of the basin. We will first define these words.

6. Iteration, sinks, saddles, separatrices

Rather than thinking of the differential equation in \mathbb{R}^3 , I find it much easier to think of the *period mapping* (or *Poincaré mapping*) in the plane

$$P : \mathbb{R}^2 \rightarrow \mathbb{R}^2 \quad \text{given by} \quad P : \begin{bmatrix} x(0) \\ x'(0) \end{bmatrix} \mapsto \begin{bmatrix} x(2\pi) \\ x'(2\pi) \end{bmatrix}.$$

This enables me to ignore what motions do between the samples.

There is no real loss if we are interested in long-term behavior: *iterating* m times the mapping P is equivalent to solving the differential equation for time $2m\pi$, sampling the solutions every 2π . But the dynamical objects are now subsets of the plane rather than of space: most people visualize objects in the plane much better than in space. In our case, the planar objects are quite complicated enough.

Seen this way, each point $s_k = (S_k(0), S'_k(0))$ is an attracting fixed point of P , also called a *sink*: $P(s_k) = s_k$ and if a point p is close to s_k (within r of it, for instance), its orbit under P approaches s_k . The basin B_k is exactly the set of points p such that the sequence $p, P(p), P^2(p), \dots$ approaches s_k .

Sinks can also be periodic of period $m > 1$. Such sinks are points p such that $P^m(p) = p$, and such that if a point p_1 is sufficiently close to p , the sequence, $p_1, P^m(p_1), P^{2m}(p_1), \dots$ tends to p . That is, the solution of the differential equation with $(x(0), x'(0)) = p$ is an attracting periodic solution of period $2m\pi$. Our mapping P appears not to have any such points (for these values of the parameters), although proving that it has none may well be an unsolvable problem. But there are infinitely many periodic saddles, as is proved by Theorem 3. And there are infinitely many more whose existence is not guaranteed by that theorem.

Like a sink, a *saddle point* for P corresponds to a periodic solution of the original differential equation, but while sinks are associated with stable equilibria, saddles are associated with unstable equilibria. A periodic solution $(x(t), x'(t))$ of the differential equation gives a saddle $(x(0), x'(0))$ of the period mapping P if there is a surface made up of solutions of the differential equation that tend to the attracting periodic solution as time tends to $+\infty$, and another surface of solutions that tend to the attracting periodic solution as $t \rightarrow -\infty$, i.e., as one travels backwards in time.

An example of a saddle point is the upwards (unstable) equilibrium for an unforced damped pendulum. Almost all solutions are captured by a stable equilibrium. But exceptional solutions exist that take an infinite amount of time to approach the vertical, and other solutions take an infinite amount to fall away from the vertical: these solutions make up two surfaces that intersect along the constant solution corresponding to the unstable

equilibrium. The surface of solutions that tend to the vertical in forward time is the *stable separatrix*, while the surface of solutions tending to the vertical in backwards time is the *unstable separatrix*. The intersection of these surfaces with a Poincaré plane (i.e., the plane $t = 0$) forms two curves, also referred to as separatrices. Think of the separatrices as watersheds: for our unforced pendulum, they separate the initial conditions that go over the top one more time from those that don't make it.

Mappings $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ (which might be the period mapping of a time-periodic differential equation in \mathbb{R}^2 , as in our case) usually also have *sources*: fixed or periodic points that repel all nearby orbits. The period mapping P for our pendulum has no sources because P contracts areas by $e^{-2\pi/10} \approx 0.53$, due to the damping [8, vol. 2, chap. 8]. No mapping can simultaneously contract areas and map some region to a strictly larger region, as would have to happen near a source. Of course, P^{-1} has sources wherever P has sinks.

7. Saddles in the boundary of B_k

The computer finds four saddles $p_{k,1}, \dots, p_{k,4}$ in the boundary of each basin. These saddles form two cycles of period 2 (i.e., the solutions of the differential equation with initial values at these saddles have period 4π). The boundary of the basin appears to be made of their stable separatrices, as drawn in Figure 7. We will call these separatrices $\sigma^+(p_{k,i})$: these are the watersheds that separate the solutions falling into the basin from those that don't.




FIGURE 7. The stable separatrices of the saddles of period 2 in the boundary of a basin provide an outline drawing of the basin. Thus this picture is more or less the same as Figure 5, but the stable manifolds would need to be continued for a very long time to get as much resolution as figure 5 provides.

In fact, the preceding statement is not true: the boundaries of the basins are not just the separatrices; they are much more complicated than that. The complication stems from

the fact that all points of the boundary are limits of *sequences*, but not all such points are limits of *paths*. Consider Wada's construction: some points of the boundary of the water are on the edge of some water stream, but most are not. For one thing, points on the edge of a coffee stream are not on the edge of a water stream, even though they are in the boundary of the water: there are water streams arbitrarily close, but tea streams even closer, etc. Such points are *inaccessible* by water: you can reach out to them over other streams, with an arbitrarily small motion, but you cannot reach them in a boat. Most points of the common boundary (the *separator*) are not accessible from the water, coffee, or tea.

Our basins are similar to those of the Wada example. Each includes a central "pond" with four canals leading off from it, which dwindle to become infinitely narrow streams, intermingled with streams belonging to other basins.

In our case, the *inward pointing* unstable separatrix at each of the four saddles is attracted to the sink, as shown in Figure 8, and provides a path from the sink to the stable separatrix of the saddle. Thus the stable separatrix is part of the accessible boundary.

Theorem 2. *The accessible boundary of B_k is exactly the union of the stable separatrices $\sigma^+(p_{k,i})$, $i = 1, \dots, 4$.*

The proof consists of looking at Figure 8.




FIGURE 8. A basin cell; the points $P^{-i}(c_i)$ illustrates the proof of Theorem 2.

The colored neighborhood C_k of the sink s_k (called a *basin cell* in [12]) is bounded by arcs of four stable separatrices $\sigma^+(p_{k,i})$ and arcs of the four unstable separatrices $\sigma^-(p_{k,i})$, which except for endpoints are contained in the interior of the basin. Thus any accessible boundary point q of B_k not in $\bigcup_i \sigma^+(p_{k,i})$ is necessarily outside C_k , and a path

$$\gamma : [0, 1] \mapsto \overline{B}_k, \gamma([0, 1]) \subset B_k$$

joining q to s_k intersects one of these four arcs, in points c_0 . Similarly, the path $P^m(\gamma)$ intersects one of these arcs in a point c_m . The points $z_m = P^{-m}(c_m)$ must be on γ , and must converge to q since for any $\epsilon > 0$, the set $\gamma([0, 1 - \epsilon])$ is a compact subset of B_k . Thus $P^m(\gamma([0, 1 - \epsilon]))$ is inside C_k (or any neighborhood of s_k for m sufficiently large).

But the c_m lie in four compact arcs of $\bigcup_i \sigma^-(p_{k,i})$, hence $P^{-m}(c_m)$ is very close to one of the saddles for m large. So q is one of the saddles $p_{k,i}$, and hence is on its stable separatrix.

This ends the proof of Theorem 2 (or at least a fairly convincing argument; it is not a rigorous proof, as we will discuss in Sections 11 and 13); now to justify Theorem 1.

First, it is enough to show that each accessible point of ∂B_0 (the boundary of B_0) can be approached by every other basin. Indeed, every point of ∂B_0 can be approached by accessible points, so if we can show that each accessible point of ∂B_0 is in the boundary of every other basin, then every point of ∂B_0 is in the boundary of every other basin.

Second, it is enough to know that the four *outward pointing branches* of the unstable separatrices for the four accessible saddles in ∂B_0 enter every basin. Indeed, if the four unstable separatrices $\sigma^-(p_{0,i})$, for $i = 1, 2, 3, 4$, enter B_n , then the inverse images of B_n accumulate to $p_{0,i}$, hence to the entire stable separatrix $\sigma^+(p_{0,i})$. This shows a little more: if all four $\sigma^-(p_{0,i})$ enter B_n , then no curve can enter B_0 without crossing a stream of B_n , i.e., entering B_n .

Third, rather than show that the outward-pointing part of each $\sigma^-(p_{0,i})$ enters all the basins B_n , for n any integer, it is enough to show that it enters the two neighboring basins B_1 and B_{-1} . We can prove this by induction. Figure 9 shows that the four separatrices $\sigma^-(p_{0,i})$, $i = 1, 2, 3, 4$ enter the basins B_{-1} and B_1 .

Now suppose they enter B_k for some $k > 1$. But they cannot enter B_k without entering B_{k+1} , because the $\sigma_{k,i}^-$ enter B_{k+1} , so that their inverse images give streams of B_{k+1} , which they must ford to enter B_k .




FIGURE 9. All four of the unstable separatrices from the points $p_{0,i}$ enter both B_1 and B_{-1} .

8. Solutions not attracted to the sinks

In this section we use techniques mainly due to Smale [14] to show that the differential equation for our pendulum has trajectories that carry out any specified sequence of gyrations. During one time interval $I_k = [2k\pi, 2(k+1)\pi)$ a solution $(x(t), x'(t))$ may satisfy $x(t) = 0 \pmod{2\pi}$ exactly

[-1]	once with $x' < 0$,
[0]	never,
[1]	once with $x' > 0$,
[NA]	none of the above.

These events correspond to the pendulum crossing the downward position exactly once clockwise, not crossing it, crossing it once counterclockwise, or doing something else. In particular, the attracting solutions belong to the “none of the above” category, because they cross the downward position twice during each period. So, eventually, do all solutions that are attracted to them. Thus Theorem 3 describes solutions entirely contained in the separator, which are never attracted to one of the sinks.

Theorem 3. *Given any bi-infinite sequence of events $\dots E_{-1}, E_0, E_1, \dots$ with $E_k \in \{-1, 0, 1\}$ (but not [NA]), there exists a solution of our differential equation that during each time interval $[2k\pi, 2(k+1)\pi)$ will “do” E_k .*

Thus given any sequence of gyrations one might choose, there is a solution that does exactly that. In particular, any sequence of E_i of period m and that sum to 0 over one such period corresponds to a periodic cycle of period m for P . Theorem 3 is very similar to Alekseev’s theorem, and is proved the same way: by exhibiting a Smale horseshoe. In Alekseev’s case this requires a delicate perturbation argument; we show how the computer can make such a result transparent.

We have found a sequence of fixed sinks s_k that correspond to the downward equilibrium of the unforced pendulum. There is also a sequence of fixed saddles corresponding to a periodic solution of the original differential equation of period 2π near the unstable upward equilibrium. If you draw a sequence of quadrilaterals Q_k roughly aligned with the stable and unstable separatrices of these fixed saddles, as in Figure 10, you expect the image of such a quadrilateral to be compressed in the stable direction and stretched in the unstable direction, becoming long and filiform.




FIGURE 10. The quadrilaterals Q_{-1} , Q_0 , Q_1 , together with the forward and backwards images of Q_0 .

Below we will describe the set of points

$$Q_k(E_0, E_1, \dots, E_N) = \{p | P^n(p) \in Q_{k+E_0+\dots+E_{n-1}} \text{ for } 0 \leq n \leq N\}.$$

Let A_0, B_0, C_0, D_0 denote the corners of Q_0 , as shown in Figure 10. The set $P(Q_0)$ is the curvilinear quadrilateral Q'_0 , shaded in Figure 10, with vertices A'_0, B'_0, C'_0, D'_0 . The key property of the image is that it crosses the quadrilaterals Q_1 and Q_{-1} , as well as itself, in each case going from top to bottom (or bottom to top), with the top A_0B_0 and bottom C_0D_0 mapping outside these quadrilaterals.

This implies that each of $Q_0([-1])$, $Q_0([0])$, $Q_0([1])$ forms a full-width subrectangle of Q_0 . Figure 11 shows the forward and backwards images of Q_0 , Q_{-1} and Q_1 , and a blow-up of showing how these intersect Q_0 . Indeed the backwards images (light shading) form full-width subrectangles. Of course, Q_1 and Q_{-1} also contain such subrectangles $Q_1(E_0)$, etc. The inverse image $P^{-1}(Q_{E_0}(E_1))$ is then again a (thinner) full-width subrectangle $Q_0(E_0, E_1)$.

Continuing this way, we see that for any finite sequence (E_0, E_1, \dots, E_N) , the corresponding set $Q_0(E_0, E_1, \dots, E_N)$ is a full-width subrectangle of Q_0 . Finally, the assignment of an infinite forward trajectory restricts the initial position to an infinite intersection of nested full-width subrectangles of Q_0 ; such an intersection is a connected subset of Q_0 connecting one side of Q_0 to the other. In fact, it is a smooth curve, but this requires writing some inequalities.

A similar argument shows that any finite backwards trajectory restricts the final position to a full-height subrectangle of Q_0 , and an infinite backwards trajectory leads to a connected subset joining A_0B_0 to C_0D_0 (again in fact a smooth curve). If $X, Y \subset Q_0$ are connected subsets, with X joining D_0A_0 to B_0C_0 and Y joining A_0B_0 to C_0D_0 , then $X \cap Y \neq \emptyset$. Thus there is a point realizing any prescribed symbolic trajectory.




FIGURE 11. The forward images of Q_{-1} , Q_0 , Q_1 , and their intersections with Q_0 . At right a blow-up of Q_0 .

Finally, I claim that the points of $Q_0([-1])$, $Q_0([0])$, $Q_0([1])$ realize the events $[-1]$, $[0]$, and $[1]$ respectively. Figure 12 shows the images of $Q_0(+1)$ and $Q_0(-1)$ at times

$$0, \frac{2\pi}{8}, \frac{4\pi}{8}, \frac{6\pi}{8}, \frac{8\pi}{8}, \frac{10\pi}{8}, \frac{12\pi}{8}, \frac{14\pi}{8}, \frac{16\pi}{8}.$$




FIGURE 12. How the quadrilaterals move during one period.

The first set certainly seems to cut the line $x = 2\pi$ exactly once with $y > 0$; the second set seems to cut the line $x = 0$ once with $y < 0$.

9. Controlling the pendulum

Imagine that the pendulum is massive, and is being used as a flywheel to control some very delicate operation, like polishing the mirror of a telescope. An array of lasers is constantly monitoring the operation, deciding on the fly whether the pendulum should turn clockwise, counterclockwise, or wait until the mirror has been repositioned.

The previous section showed that there are motions of the pendulum performing any specified sequence of gyrations, in particular the one required a posteriori by the polisher. But on second thought this seems useless: these motions are extremely unstable, and the slightest error in the initial condition destroys them, as well as any perturbation of the differential equation itself. But if the machine is to perform any work, this inevitably perturbs the differential equation, in a way that is essentially unpredictable (you cannot predict how much work one swipe of the polisher will accomplish), and in any case we don't know ahead of time the sequence of swipes and stops the task will require.

On third thought, we see that the instability of the specified motions is exactly what should make them useful! Suppose that our array of sensors controls the current $f(t)$ that is forcing the pendulum, changing it from $\cos(t)$ to something like

$$(1 + a(t)) \cos(t) \quad (\text{amplitude modulation}) \quad \text{or} \\ \cos((1 + a(t))t) \quad (\text{frequency modulation}),$$

where $a(t)$ represents the fine-tuning necessary to achieve the desired sequence of gyrations. The point is that we do not have to figure out what sequence we want ahead of time: the sensors can react to the polishing of the telescope on the fly, computing the adjustment $a(t)$ that is necessary. It is because of the instability that you can keep $a(t)$ small and still realize any sequence of gyrations: you don't need to grind to a halt, compute, and start up again; the corrections can be done smoothly. A useful analogy is skiing: a beginning skier plants his skis well apart, seeking stability, which is fine until he tries to turn and discovers he can't. An expert skier, with skis parallel and touching, is highly unstable, and a slight wiggle of the hips allows him to negotiate a mogul. Of course he doesn't plot his entire path at the top of the mountain; he calculates the slight adjustments $a(t)$ as they are needed.

Theorem 4. *For any sequence of events E_0, E_1, \dots and any sufficiently small disturbance $b(t)$ of the forcing term $\cos t$, there exists a function $a(t)$ of the same order of magnitude as $b(t)$ and an initial condition $x(0), x'(0)$ such that the solution of the differential equation*

$$x'' + 0.1x'(t) + \sin(x) + b(t) = (1 + a(t)) \cos t$$

with those initial conditions realizes the specified sequence of events.

This result is fairly obvious: choose $a(t)$ as the pendulum approaches the upwards position so as to speed it up or slow it down as required. The problem is how to compute the $a(t)$, in terms of available data. Clearly $a(t)$ should depend only on the values of b up

to time $t - 2\pi$; it should not depend on the specified sequence of events very far ahead, as this is unknown. How small can $a(t)$ be made? How far ahead in the required sequence of events does it need to look? How sensitive is it to small errors in the sensors? . . .

10. Control and celestial mechanics

To return to celestial mechanics for a moment, it is interesting to note that when sending a spaceship to visit the outer solar system, NASA uses the instabilities of the differential equations describing gravity in much the same way as we have used the instabilities of the pendulum. It is well beyond present-day engineering to send a spaceship out of the solar system by simply using its fuel to accelerate it. Instead, it is allowed to “fall” into the sun, with an orbit that passes close to Venus. It then loops around Venus; we can imagine that it is the “satellite” in the three body system consisting of itself, Venus, and the sun.

This system is similar to Alekseev’s (somewhat more complicated: a Poincaré section would need to be 4-dimensional rather than 2), and one can prescribe an orbit so that the space ship steals a tiny amount of potential energy from Venus, speeding up enormously in the process, and ends up in a very unstable state where a small push by guidance rockets can put it on the path to Jupiter.

This scenario is then repeated near Jupiter, Saturn, and Uranus, with the spaceship each time gaining momentum, and using small pushes to head itself in the direction of the next destination. Thus the chaos of the solar system is essential to its exploration.

11. What is proved?

To what extent does this paper prove anything? As written, no statement is proved anywhere: for the punchline we just looked at a computer picture. How do we know that these pictures are right? I will not address the possibility that the programs have essential bugs and are computing something other than what I think, or the esoteric possibility that the computer arithmetic is wrong. But even if the computer is computing exactly what I think, that is still only an approximation to solutions of the differential equations; we need to quantify the quality of the approximation. The contribution of round-off error also should be addressed.

Actually, many of the results are not hard to prove rigorously, namely all those where we have to show that after time 2π , solutions are within some fairly large ϵ of the value suggested by the computer drawings.

Good estimates of long-term errors of numerical approximations to solutions of differential equations are notoriously hard to come by, but that is not really a problem here. First, we do not need *good* estimates (solutions only need to be accurate to about 0.1); second, the time considered is *not long* (2π); and most important, the differential equation has a small *Lipschitz constant* ($\sqrt{2.001} < 1.42$). Errors in solutions to differential equations grow at most exponentially, at a rate e^{kt} , where t is time (in our case, 2π) and k is the Lipschitz constant; with $k < 1.42$, errors grow at a fairly small interest rate, and can be controlled for a short time.

Using these numbers, a straightforward computation using the *fundamental inequality* ([8, Chapters 4 and 6]) shows that if the initial velocity satisfies $|x'(0)| < 3$, then Euler's method with step-length $h = 0.000002$ gives results accurate to 0.1 after time 2π . Moreover, the same inequality shows that round-off error contributes a much smaller error yet. This is not a good way to do such numerics; better numerical methods give much better estimates [5]. For instance, formula (14) of [2] can be used to show that the fourth order Runge-Kutta method with step 0.005 has more than the needed precision.

A word of caution, though. The elementary bound above says that errors of all types are multiplied by at most $e^{2\pi 1.42} \approx 7500$ over one time period. It is not too difficult to improve this to $e^{2\pi 1.1} \approx 1000$, and one could improve it further. But one could not improve it very much further.

Consider for example the completely unavoidable error caused by the computer's inability to handle numbers with infinite precision. If it handles numbers to 16 significant digits, you may think you are starting at a saddle point, but your initial error (the distance between the saddle point and where you really are) may be as great as 10^{-16} . The largest eigenvalue λ of the linearization of P at the fixed saddles in the Q_k is about 321 (according to the computer). As long as you are in the region where P is approximately its linearization at this saddle, errors of all types are expanded by a factor of λ over one time period, and hence λ^m over m time periods. So after m iterations the error will have mushroomed to $10^{-16}(321^m)$: for $m = 7$ the initial minute error will have grown to 35. But already for an error of 1, you will have been booted out of the region where the linearization is a reasonable approximation to reality.

Thus no numerical method can guarantee even one digit of accuracy after six time periods, if we are computing with 16 significant digits. In fact, the reality is much worse than that, and I wouldn't trust anything after four time periods without some good reason.

12. *A posteriori bounds*

Good reasons to trust solutions are available: I advocate extrapolation, as described in [8, Chapter 3]. At the moment, this only works for fixed step-length, but for a Poincaré mapping of a differential equation, fixed step-length is probably best anyway. For other possible methods, consult [10].

Denote by $u_h(t)$ the numerical approximation to the solution of some differential equation given by the standard fourth order Runge-Kutta method, with $u_h(0) = a$. Then the theory asserts that for each fixed t the approximation $u_h(t)$ converges to the value of the solution $u(t)$, and that we have an asymptotic development

$$u_h(t) = u(t) + Ch^4 + o(h^4).$$

The exponent 4 is a feature of this approximation procedure; other procedures have different exponents.

If for some h we know $u_h(t)$, $u_{h/2}(t)$, and $u_{h/4}(t)$, and we assume that we have an asymptotic development of the form $u_h(t) = u(t) + Ch^k + o(h^k)$ for some k , we can extrapolate the values of k and of C from the values of the approximate solutions:

$$k = \frac{1}{\log 2} \log \left| \frac{u_h(t) - u_{h/2}(t)}{u_{h/2}(t) - u_{h/4}(t)} \right| \quad \text{and} \quad C = \frac{2^k}{2^k - 1} \frac{u_h(t) - u_{h/2}(t)}{h^k}.$$

Now suppose we calculate $u_{h/2^m}(t)$ for a range of values of m , focusing on the expression for k above. The theory says that as m increases, the value of k should approach 4, but that doesn't take round-off error into account; typically the value of k approaches 4 as m increases, then veer away from 4 as round-off error takes over. If there is a range of values of m where k is close to 4, the approximation is happening the way the theory predicts, and we can probably trust the corresponding estimate of the error. The following data illustrates this for our differential equation, solved for $0 \leq t \leq 16\pi$, i.e., for 8 periods. We start with the two initial positions (7.15859, 0.14097) and (7.16859, 0.14097). The extrapolations we find are

steps	first solution		second solution	
	order	error	order	error
6				
12		22.45		86.22
24	3.07	2.67	1.05	41.48
48	-1.79	9.31	2.61	6.77
96	3.26	0.96	-0.15	6.84
192	-0.44	1.31	-1.09	14.64
384	-2.01	5.27	3.02	1.80
768	-0.06	5.48	4.96	0.057
1536	5.13	0.16	4.19	0.003

Thus, the first approximation never becomes reliable; the order is never close to 4. In particular, there is no reason to think that the quantity in the "error" column is actually an estimate of the error. But the second appears to be converging nicely, with the order approaching 4, and probably the error estimate of 0.003 is reliable. Thus although any estimate we make *a priori* for a bound for the error is bound to be wildly pessimistic, after the computation we can make a good guess as to how reliable it is.

13. Questions and observations

- (1) Are there any periodic sinks other than the attracting fixed points we found? I have no idea how to attack this problem. For one thing, I don't trust computer drawings on this point: in many instances I eventually found sinks whose basins were too small to be visible on computer drawings unless you knew where to look. For another, the answer might depend in the most delicate way on the parameters: there definitely are other attracting fixed points when the forcing term is $1.22 \cos t$ instead of $\cos t$; for example, there is a sink of period 3, where solutions go from the point with coordinates $x = -1.29785, y = 1.0025$ to the point $x = -1.3349, y = -0.21286$, to the point $x = -3.004469, y = 0.17586$, and then back to the first point In fact, with those parameters there are at least two more sinks of period 3, in addition to all the translates of the three sinks by 2π .

In fact, this problem may be unsolvable. John Milnor's candidate for the simplest unsolvable problem of mathematics is the question: "Does the polynomial $x^2 - 1.5$

have an attracting cycle?” Of course, if it does, one can find it with a finite amount of work. But if it doesn’t, there may be no proof of this fact.

- (2) Is the complement of all the basins B_k of measure 0? This would mean that with probability 1 every initial point is attracted to a sink. I think this is the case, but have no solid grounds for this belief. Even the computer isn’t very definite, and besides, this is one point where numerical error might really be important: the perturbations of the period mapping due to errors of integration and round-off might affect the probability of being attracted to a sink.

Acknowledgements

I thank Stan Wagon, Mukund Thattai, and Robert Terrell for their many helpful comments and suggestions.

Bibliography

1. V.M. Alekseev, On the capture orbits for the three-body problem for negative energy constant, *Uspekhi Mat. Nauk* 24 (1969) 185–186.
2. L. Bieberbach, On the remainder of the Runge-Kutta formula in the theory of ordinary differential equations, *Zeitschrift für angewandte Mathematik und Physik* 2 (1951) 233–248.
3. R. Borelli and C. Coleman, Computers, lies and the fishing season, *College Math. Jour.* 25 (1994) 401–412.
4. L.E.J Brouwer, Zur Analysis Situs, *Math. Annalen* 68 (1910) 422–434.
5. P. Henrici, *Discrete variable methods in ordinary differential equations*, Wiley, New York, 1962.
6. B. Hinkle, J. Hubbard, and B. West, *Planar Systems and Planar Iterations*, to be published by Springer Verlag. These programs constitute the latest version of *MacMath* by Hubbard and West, originally published by Springer Verlag in 1991.
7. J. Hubbard and R. Oberste-Vorth, Hénon mapping in the complex domain II: Projective and inductive limits of polynomials, *Proceedings of the NATO summer conference in Hillerod*, 1993.
8. J. Hubbard and B. West, *Differential Equations, a dynamical systems approach*, Vol. I, Springer Verlag, New York, 1992 and Vol. II, Springer Verlag, New York, 1994.
9. J. Kennedy and J. Yorke, Basins of Wada, *Physica D* 51 (1991) 213–255.
10. R. Knapp and S. Wagon, Check your answers . . . but how? , *Mathematica in Education and Research*, 7-4.
11. J. Moser, Stable and random motions in dynamical systems, *Annals of Mathematics Studies*, Princeton University Press, Princeton, N.J., 1973
12. H. Nusse and J. Yorke, Wada basin boundaries, preprint March 1994, 1–51.
13. D. Schwalbe and S. Wagon, *VisualDSolve: Visualizing Differential Equations with Mathematica*, Springer/TELOS, N.Y. 1997.
14. D. Smale, Diffeomorphisms with many periodic points, *Differential and combinatorial topology*, edited by S. S. Cairns, Princeton University Press, 1965, pp. 63–80.
- 15 K. Yoneyama, Theory of continuous set of points, *Tohoku Math. J* 11–12 (1917) 43.

*Department of Mathematics
Cornell University
White Hall
Ithaca, NY 14853
jhh8@cornell.edu*