

Chapter 3

Discrete time dynamics

Gentles, perchance you wonder at this show; But wonder on, till truth make all things plain.
 — W. Shakespeare, *A Midsummer Night's Dream*

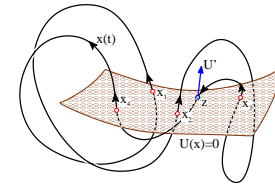
THE TIME PARAMETER in the definition of a dynamical system can be either continuous or discrete. Discrete time dynamical systems arise naturally from flows. In general there are two strategies for replacing a continuous-time flow by iterated mappings; by cutting it by Poincaré sections, or by *strobing* it at a sequence of instants in time. Think of your partner moving to the beat in a disco: a sequence of frozen stills. While ‘strobing’ is what any numerical integrator does, by representing a trajectory by a sequence of time-integration step separated points, strobing is in general not a reduction of a flow, as the sequence of strobed points still resides in the full state space \mathcal{M} , of dimensionality d . An exception are non-autonomous flows that are externally periodically forced. In that case it might be natural to observe the flow by strobing it at time intervals fixed by the external forcing, as in example 7.7 where strobing of a periodically forced Hamiltonian leads to the ‘standard map.’

section 2.1

In the *Poincaré section method* one records the coordinates of a trajectory whenever the trajectory crosses a prescribed trigger. This triggering event can be as simple as vanishing of one of the coordinates, or as complicated as the trajectory cutting through a curved hypersurface. A Poincaré section (or, in the remainder of this chapter, just ‘section’) is *not* a projection onto a lower-dimensional space: Rather, it is a local change of coordinates to a direction along the flow, and the remaining coordinates (spanning the section) transverse to it. No information about the flow is lost by reducing it to its set of Poincaré section points and the return maps connecting them; the full space trajectory can always be reconstructed by integration from the nearest point in the section.

Reduction of a continuous time flow to its Poincaré section is a powerful visualization tool. But, the method of sections is more than visualization; it is also a fundamental tool of dynamics - to fully unravel the geometry of a chaotic flow,

Figure 3.1: A trajectory $x(t)$ that intersects a Poincaré section \mathcal{P} at times t_1, t_2, t_3, t_4 , and closes a cycle $(\hat{x}_1, \hat{x}_2, \hat{x}_3, \hat{x}_4)$, $\hat{x}_k = x(t_k) \in \mathcal{P}$ of topological length 4 with respect to the section. In general, the intersections are not normal to the section. Note also that the crossing z does not count, as it is in the wrong direction.



one *has to* quotient all of its symmetries, and evolution in time is one of these (This delphic piece of hindsight will be illuminated in chapter 10).

3.1 Poincaré sections



A continuous time flow decomposes the state space into Lagrangian ‘spaghetti’ of figure 2.2, a union of non-intersecting 1-dimensional orbits. Any point on an orbit can be used to label the orbit, with the state space thus reduced to a ‘skew-product’ of a $(d-1)$ -dimensional space \mathcal{P} of labeling points $\hat{x}_j \in \mathcal{P}$ and the corresponding 1-dimensional orbit curves \mathcal{M}_j on which the flow acts as a time translation. However, as orbits can be arbitrary complicated and, if unstable, uncontrollable for times beyond the Lyapunov time (1.1), in practice it is necessary to split the orbit into finite trajectory segments, with time intervals corresponding to the shortest recurrence times on a non-wondering set of the flow, *finite* times for which the flow is computable. A particular prescription for picking the orbit-labeling points is called a *Poincaré section*. In introductory texts Poincaré sections are treated as pretty visualizations of a chaotic flows, akin to plastic surgery and Botox, but their dynamical significance is much deeper than that. Once a section is defined, a ‘Lagrangian’ description of the flow (discussed above, page 47) is replaced by the ‘Eulerian’ formulation, with the trajectory-tangent velocity field $v(\hat{x})$, $\hat{x} \in \mathcal{P}$ enabling us to go freely between the time-quotiented space \mathcal{P} and the full state space \mathcal{M} . The dynamically important *transverse dynamics* –description of how nearby trajectories attract / repel each other– is encoded in mapping of $\mathcal{P} \rightarrow \mathcal{P}$ induced by the flow - dynamics along orbits is of secondary importance.

chapter 10

Successive trajectory intersections with a Poincaré section, a $(d-1)$ -dimensional hypersurface embedded in the d -dimensional state space \mathcal{M} , figure 3.1, define the *Poincaré return map* $P(\hat{x})$, a $(d-1)$ -dimensional map of form

$$\hat{x}' = P(\hat{x}) = f^{\tau(\hat{x})}(\hat{x}), \quad \hat{x}', \hat{x} \in \mathcal{P}. \tag{3.1}$$

Here the *first return function* $\tau(\hat{x})$ —sometimes referred to as the *ceiling function*—is the time of flight to the next section for a trajectory starting at \hat{x} . The choice of the section hypersurface \mathcal{P} is altogether arbitrary. It is rarely possible to define a single section that cuts across all trajectories of interest. Fortunately, one often needs only a local section, a finite hypersurface of codimension 1 intersected by

a swarm of trajectories near to the trajectory of interest (the case of several sections is discussed in sect. 3.4). Such hypersurface can be specified implicitly by a single condition, through a function $U(x)$ that is zero whenever a point x is on the Poincaré section,

$$\hat{x} \in \mathcal{P} \quad \text{iff} \quad U(\hat{x}) = 0. \quad (3.2)$$

The gradient of $U(x)$ evaluated at $\hat{x} \in \mathcal{P}$ serves a two-fold function. First, the flow should pierce the hypersurface \mathcal{P} , rather than being tangent to it. A nearby point $\hat{x} + \delta x$ is in the hypersurface \mathcal{P} if $U(\hat{x} + \delta x) = 0$. A nearby point on the trajectory is given by $\delta x = v\delta t$, so a traversal is ensured by the *transversality condition*

$$(v \cdot \nabla U) = \sum_{j=1}^d v_j(\hat{x}) \partial_j U(\hat{x}) \neq 0, \quad \partial_j U(\hat{x}) = \frac{\partial}{\partial \hat{x}_j} U(\hat{x}), \quad \hat{x} \in \mathcal{P}. \quad (3.3)$$

Second, the gradient ∇U defines the orientation of the hypersurface \mathcal{P} . The flow is oriented as well, and a periodic orbit can pierce \mathcal{P} twice, traversing it in either direction, as in figure 3.1. Hence the definition of Poincaré return map $P(\hat{x})$ needs to be supplemented with the orientation condition

$$\begin{aligned} \hat{x}_{n+1} &= P(\hat{x}_n), & U(\hat{x}_{n+1}) &= U(\hat{x}_n) = 0, & n &\in \mathbb{Z}^+ \\ & & \sum_{j=1}^d v_j(\hat{x}_n) \partial_j U(\hat{x}_n) &> 0. & & \end{aligned} \quad (3.4)$$

In this way the continuous time t flow $x(t) = f^t(x)$ is reduced to a discrete time n sequence \hat{x}_n of successive *oriented* trajectory traversals of \mathcal{P} .

chapter 17

With a sufficiently clever choice of a Poincaré section or a set of sections, any orbit of interest intersects a section. Depending on the application, one might need to convert the discrete time n back to the continuous flow time. This is accomplished by adding up the first return function times $\tau(\hat{x}_n)$, with the accumulated flight time given by

$$t_{n+1} = t_n + \tau(\hat{x}_n), \quad t_0 = 0, \quad x_n \in \mathcal{P}. \quad (3.5)$$

Other quantities integrated along the trajectory can be defined in a similar manner, and will need to be evaluated in the process of evaluating dynamical averages.

A few examples may help visualize this.

Example 3.1 A template and the associated hyperplane Poincaré section:
The simplest choice of a Poincaré section is a plane \mathcal{P} specified by a ‘template’ point

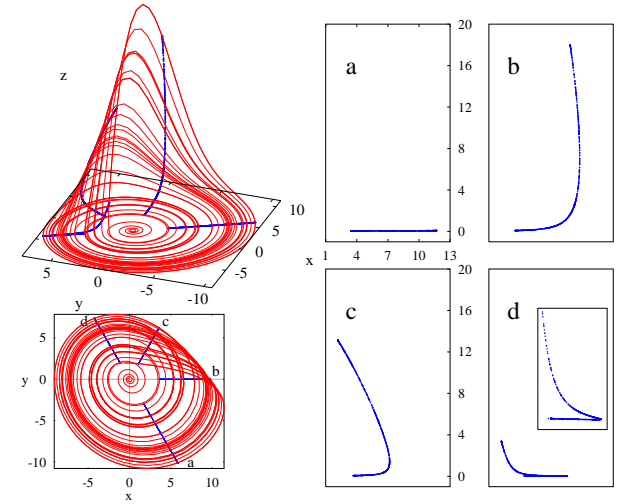


Figure 3.2: (Right:) a sequence of Poincaré sections of the Rössler strange attractor, defined by planes through the z axis, oriented at angles (a) -60° (b) 0° , (c) 60° , (d) 120° , in the x - y plane. (Left:) side and x - y plane view of a typical trajectory with Poincaré sections superimposed. (R. Paškauskas)

(located at the tip of the vector \hat{x}') and a normal vector \hat{n} perpendicular to the plane. A point \hat{x} is in this plane if it satisfies the linear condition

$$U(\hat{x}) = (\hat{x} - \hat{x}') \cdot \hat{n} = 0 \quad \text{for } \hat{x} \in \mathcal{P}. \quad (3.6)$$

Consider a circular periodic orbit centered at \hat{x}' , but not lying in \mathcal{P} . It pierces the hyperplane twice; the $v \cdot \hat{n} > 0$ transversal orientation condition (3.4) ensures that the first return time is the full period of the cycle. (continued in example 12.1)

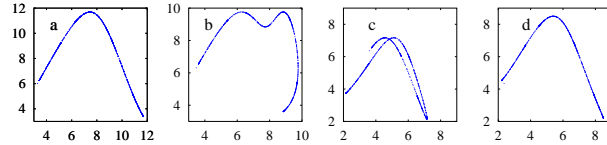
What about smooth, continuous time flows, with no obvious surfaces that would be good Poincaré sections?

Example 3.2 Pendulum: The phase space of a simple pendulum is 2-dimensional: momentum on the vertical axis and position on the horizontal axis. We choose the Poincaré section to be the positive horizontal axis. Now imagine what happens as a point traces a trajectory through this phase space. As long as the motion is oscillatory, in the pendulum all orbits are loops, so any trajectory will periodically intersect the line, that is the Poincaré section, at one point.

Consider next a pendulum with friction, such as the unforced Duffing system plotted in figure 2.4. Now every trajectory is an inward spiral, and the trajectory will intersect the Poincaré section $y = 0$ at a series of points that get closer and closer to either of the equilibrium points; the Duffing oscillator at rest.

Motion of a pendulum is so simple that you can sketch it yourself on a piece of paper. The next example offers a better illustration of the utility of visualization of dynamics by means of Poincaré sections.

Figure 3.3: Return maps for the $r_n \rightarrow r_{n+1}$ radial distance Poincaré sections of figure 3.2. The ‘multi-valuedness’ of (b) and (c) is only apparent: the full return map is 2-dimensional, $\{r', z'\} = P\{r, z\}$. (R. Paškauskas)



Example 3.3 Rössler flow: (continued from example 2.3) Consider figure 2.6, a typical trajectory of the 3-dimensional Rössler flow (2.17). The strange attractor wraps around the z axis, so one choice for a Poincaré section is a plane passing through the z axis. A sequence of such Poincaré sections placed radially at increasing angles with respect to the x axis, figure 3.2, illustrates the ‘stretch & fold’ action of the Rössler flow, by assembling these sections into a series of snapshots of the flow. A line segment in (a), traversing the width of the attractor at $y = 0, x > 0$ section, starts out close to the x - y plane, and after the stretching (a) \rightarrow (b) followed by the folding (c) \rightarrow (d), the folded segment returns (d) \rightarrow (a) close to the initial segment, strongly compressed. In one Poincaré return the interval is thus stretched, folded and mapped onto itself, so the flow is expanding. It is also mixing, as in one Poincaré return a point from the interior of the attractor can map onto the outer edge, while an edge point lands in the interior.

Once a particular Poincaré section is picked, we can also exhibit the return map (3.1), as in figure 3.3. Cases (a) and (d) are examples of nice 1-to-1 return maps. While (b) and (c) appear multimodal and non-invertible, they are artifacts of projecting a 2-dimensional return map $(r_n, z_n) \rightarrow (r_{n+1}, z_{n+1})$ onto a 1-dimensional subspace $r_n \rightarrow r_{n+1}$. (continued in example 3.5)



fast track:
sect. 3.3, p. 63

The above examples illustrate why a Poincaré section gives a more informative snapshot of the flow than the full flow portrait. For example, while the full flow portrait of the Rössler flow figure 2.6 gives us no sense of the thickness of the attractor, we see clearly in the Poincaré sections of figure 3.2 that even though the return maps are 2-dimensional \rightarrow 2-dimensional, the flow contraction is so strong that for all practical purposes it renders the return maps 1-dimensional. (We shall quantify this claim in example 4.6.)

3.1.1 Section border

How far does the neighborhood of a template extend along the hyperplane (3.6)? A section captures faithfully neighboring orbits as long as it cuts them transversally; it fails the moment the velocity field at point \hat{x}^* fails to pierce the section. At this location the velocity is tangent to the section and, thus, orthogonal to the template normal \hat{n} ,

$$\hat{n} \cdot v(\hat{x}^*) = 0, \quad \hat{x}^* \in S, \quad (3.7)$$

i.e., $v_{\perp}(\hat{x})$, component of the $v(\hat{x})$ normal to the section, vanishes at \hat{x}^* . For a smooth flow such points form a smooth $(d-2)$ -dimensional section border $S \subset \mathcal{P}$, encompassing the open neighborhood of the template characterized by qualitatively similar flow. We shall refer to this region of the section hyperplane as the (maximal) chart of the template neighborhood for a given hyperplane (3.6).

If the template point is an equilibrium x_q , there is no dynamics exactly at this point as the velocity vanishes ($v(x_q) = 0$ by the definition of equilibrium) and cannot be used to define a normal to the section. Instead, we use the local linearized flow to construct the local Poincaré section \mathcal{P} . We orient \mathcal{P} so the unstable eigenvectors are transverse to the section, and at least the slowest contracting eigenvector is tangent to the section, as in figure 4.7. This ensures that the flow is transverse to \mathcal{P} in an open neighborhood of the template x_q .

exercise 3.7

Visualize the flow as a smooth 3-dimensional steady fluid flow cut by a 2-dimensional sheet of light. Lagrangian particle trajectories either cross, are tangent to, or fail to reach this plane; the 1-dimensional curves of tangency points define the section border. An example is offered by the velocity field of the Rössler flow of figure 4.6. Pick a Poincaré section hyperplane so it goes through both equilibrium points. The section might be transverse to a large neighborhood around the inner equilibrium x_- , but dynamics around the outer equilibrium x_+ is totally different, and the competition between the two types of motion is likely to lead to vanishing of $v_{\perp}(\hat{x})$, component of the $v(\hat{x})$ normal to the section, someplace in-between the two equilibria. A section is good up to the section border, but beyond it an orbit infinitesimally close to \hat{x}^* generically does not cross the section hyperplane, at least not infinitesimally close to S .

For 3-dimensional flows, the section border S is a 1-dimensional closed curve in the section 2-dimensional \mathcal{P} , and easy to visualize. In higher dimensions, the section border is a $(d-2)$ -dimensional manifold, not easily visualized, and the best one can do is to keep checking for change of sign (3.4) at Poincaré section returns of nearby trajectories close to the section border hypersurface S ; (3.7) will be positive inside, negative immediately outside S .

Thus for a nonlinear flow, with its complicated curvilinear invariant manifolds, a single section rarely suffices to capture all of the dynamics of interest.

3.1.2 What is the best Poincaré section?

In practice, picking sections is a dark and painful art, especially for high-dimensional flows where the human visual cortex falls short. It helps to understand why we need them in the first place.

Whenever a system has a continuous symmetry G , any two solutions related by the symmetry are equivalent. We do not want to keep recomputing these over and over. We would rather replace the whole continuous family of solutions by one solution in order to be more efficient. This approach replaces the dynamics (M, f) with dynamics on the quotient state space $(M/t, \hat{f})$. For now, we only

chapter 10

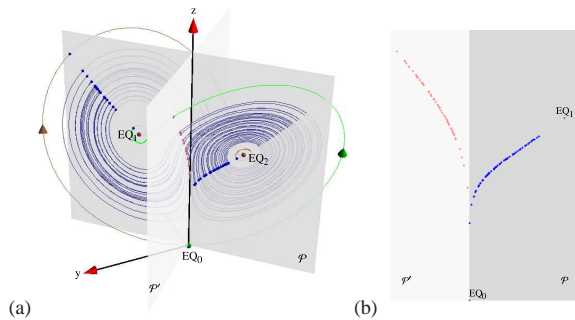


Figure 3.4: (a) Lorenz flow figure 2.5 cut by $y = x$ Poincaré section plane \mathcal{P} through the z axis and both $E_{Q_{1,2}}$ equilibria. Points where flow pierces into section are marked by dots. To aid visualization of the flow near the E_{Q_0} equilibrium, the flow is cut by the second Poincaré section, \mathcal{P}' , through $y = -x$ and the z axis. (b) Poincaré sections \mathcal{P} and \mathcal{P}' laid side-by-side. The singular nature of these sections close to E_{Q_0} will be elucidated in example 4.7 and figure 11.8 (b). (E. Siminos)

remark that constructing explicit quotient state space flow \hat{f} is either extremely difficult, impossible, or generates unintelligible literature. Our solution (see chapter 10) will be to resort to the method of slices.

Time evolution itself is a 1-parameter Lie group, albeit a highly nontrivial one (otherwise this book would not be much of a doorstop). The invariants of the flow are its infinite-time orbits; particularly useful invariants are compact orbits such as equilibrium points, periodic orbits, and tori. For any orbit it suffices to pick a single state space point $x \in \mathcal{M}_p$, the rest of the orbit is generated by the flow.

Choice of this one “labeling” point is utterly arbitrary; in dynamics this is called a “Poincaré section,” and in theoretical physics this goes by the exceptionally uninformative name of “gauge fixing.” The price is that one generates “ghosts,” or, in dynamics, increases the dimensionality of the state space by additional constraints (see sect. 13.4). It is a commonly deployed but inelegant procedure where symmetry is broken for computational convenience, and restored only at the end of the calculation, when all broken pieces are reassembled.

With this said, there are a few rules of thumb to follow: (a) You can pick as many sections as convenient, as discussed in sect. 3.4. (b) For ease of computation, pick linear sections (3.6) when possible. (c) If equilibria play important role in organizing a flow, pick sections that go through them (see example 3.4). In that case, try to place contractor eigenvectors inside the hyperplane, see Lorenz figure 3.4. Remember, the stability eigenvectors are never orthogonal to each other, unless that is imposed by some symmetry. (d) If you have a global discrete or continuous symmetry, pick sections left invariant by the symmetry (see example 9.14). For example, setting the normal vector \hat{n} in (3.6) at x to be the velocity $v(x)$ is natural and locally transverse. (e) If you are solving a local problem, like finding a periodic orbit, you do not need a global section. Pick a section or a set of (multi-shooting) sections on the fly, requiring only that they are locally transverse to the flow. (f) If you have another rule of thumb dear to you, let us know.

chapter 9

Example 3.4 Sections of Lorenz flow: (continued from example 2.2) The plane \mathcal{P} fixed by the $x = y$ diagonal and the z -axis depicted in figure 3.4 is a natural choice of a Poincaré section of the Lorenz flow of figure 2.5, as it contains all three equilib-

ria, $x_{EQ_0} = (0, 0, 0)$ and the (2.13) pair x_{EQ_1}, x_{EQ_2} . A section has to be supplemented with the orientation condition (3.4): here points where flow pierces into the section are marked by dots.

Equilibria x_{EQ_1}, x_{EQ_2} are centers of out-spirals, and close to them the section is transverse to the flow. However, close to E_{Q_0} trajectories pass the z -axis either by crossing the section \mathcal{P} or staying on the viewer's side. We are free to deploy as many sections as we wish: in order to capture the whole flow in this neighborhood we add the second Poincaré section, \mathcal{P}' , through the $y = -x$ diagonal and the z -axis. Together the two sections, figure 3.4 (b), capture the whole flow near E_{Q_0} . In contrast to Rössler sections of figure 3.2, these appear very singular. We explain this singularity in example 4.7 and postpone construction of a Poincaré return map until example 9.14. (E. Siminos and J. Halcrow)

3.2 Computing a Poincaré section



(R. Mainieri)

For almost any flow of physical interest a Poincaré section is not available in analytic form, so one tends to determine it crudely, by numerically bracketing the trajectory traversals of a section and iteratively narrowing the bracketing time interval. We describe here a smarter method, which you will only need when you seriously look at a strange attractor, with millions of points embedded in a high(er)-dimensional Poincaré section - so skip this section on the first reading.

remark 3.2

Consider the system (2.6) of ordinary differential equations in the vector variable $x = (x_1, x_2, \dots, x_d)$

$$\frac{dx_i}{dt} = v_i(x, t), \tag{3.8}$$

where the flow velocity v is a vector function of the position in state space x and the time t . In general, the map $f^{t_n}(x_n) = x_n + \int d\tau v(x(\tau))$ cannot be integrated analytically, so we will have to resort to numerical integration to determine the trajectories of the system. Our task is to determine the points at which the numerically integrated trajectory traverses a given hypersurface. The hypersurface will be specified implicitly through a function $U(x)$ that is zero whenever a point x is on the Poincaré section, such as the hyperplane (3.6).

If we use a tiny step size in our numerical integrator, we can observe the value of U as we integrate; its sign will change as the trajectory crosses the hypersurface. The problem with this method is that we have to use a very small integration time step. However, there is a better way to land exactly on the Poincaré section.

Let t_a be the time just before U changes sign, and t_b the time just after it changes sign. The method for landing exactly on the Poincaré section will be to

convert one of the space coordinates into an integration variable for the part of the trajectory between t_a and t_b . Using

$$\frac{dx_k}{dx_1} \frac{dx_1}{dt} = \frac{dx_k}{dx_1} v_1(x, t) = v_k(x, t) \quad (3.9)$$

we can rewrite the equations of motion (3.8) as

$$\frac{dt}{dx_1} = \frac{1}{v_1}, \dots, \frac{dx_d}{dx_1} = \frac{v_d}{v_1}. \quad (3.10)$$

Now we use x_1 as the ‘time’ in the integration routine and integrate it from $x_1(t_a)$ to the value of x_1 on the hypersurface, determined by the hypersurface intersection condition (3.6). This is the end point of the integration, with no need for any interpolation or backtracking to the surface of section. The x_1 -axis need not be perpendicular to the Poincaré section; any x_i can be chosen as the integration variable, provided the x_i -axis is not parallel to the Poincaré section at the trajectory intersection point. If the section crossing is transverse (3.3), v_1 cannot vanish in the short segment bracketed by the integration step preceding the section, and the point on the Poincaré section.

Example 3.5 Computation of Rössler flow Poincaré sections. (continued from example 3.3) Convert Rössler equation (2.17) to cylindrical coordinates:

$$\begin{aligned} \dot{r} &= v_r = -z \cos \theta + ar \sin^2 \theta \\ \dot{\theta} &= v_\theta = 1 + \frac{z}{r} \sin \theta + \frac{a}{2} \sin 2\theta \\ \dot{z} &= v_z = b + z(r \cos \theta - c). \end{aligned} \quad (3.11)$$

Poincaré sections of figure 3.2 are defined by the fixing angle $U(x) = \theta - \theta_0 = 0$. In principle one should use the equilibrium x_+ from (2.18) as the origin, and its eigenvectors as the coordinate frame, but here original coordinates suffice, as for parameter values (2.17), and (x_0, y_0, z_0) sufficiently far away from the inner equilibrium, θ increases monotonically with time. Integrate

$$\frac{dr}{d\theta} = v_r/v_\theta, \quad \frac{dt}{d\theta} = 1/v_\theta, \quad \frac{dz}{d\theta} = v_z/v_\theta \quad (3.12)$$

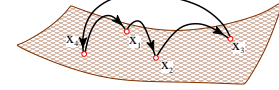
from (r_n, θ_n, z_n) to the next Poincaré section at θ_{n+1} , and switch the integration back to (x, y, z) coordinates. (continued in example 4.1) (Radford Mitchell, Jr.)

3.3 Mappings



Do it again!
—Isabelle, age 3

Figure 3.5: A flow $x(t)$ of figure 3.1 represented by a Poincaré return map that maps points in the Poincaré section \mathcal{P} as $\hat{x}_{n+1} = f(\hat{x}_n)$. In this example the orbit of \hat{x}_1 is periodic and consists of the four periodic points $(\hat{x}_1, \hat{x}_2, \hat{x}_3, \hat{x}_4)$.



Though we have motivated discrete time dynamics by considering sections of a continuous flow and reduced the continuous-time flow to a family of maps $P(\hat{x})$ mapping points \hat{x} from a section to a section, there are many settings in which dynamics is inherently discrete, and naturally described by repeated iterations of the same map

remark 3.1

$$f : \mathcal{M} \rightarrow \mathcal{M},$$

or sequences of consecutive applications of a finite set of maps, a different map, f_A, f_B, \dots , for points in different regions $\{\mathcal{M}_A, \mathcal{M}_B, \dots, \mathcal{M}_Z\}$,

$$\{f_A, f_B, \dots, f_Z\} : \mathcal{M} \rightarrow \mathcal{M}, \quad (3.13)$$

for example maps relating different sections among a set of Poincaré sections. The discrete ‘time’ is then an integer, the number of applications of the map or maps. As writing out formulas involving repeated applications of a set of maps explicitly can be awkward, we streamline the notation by denoting the (non-commutative) map composition by ‘o’

$$f_Z(\dots f_B(f_A(x))) \dots = f_Z \circ \dots \circ f_B \circ f_A(x), \quad (3.14)$$

and the n th iterate of map f by

$$f^n(x) = f \circ f^{n-1}(x) = f(f^{n-1}(x)), \quad f^0(x) = x.$$

section 2.1

The trajectory of x is the finite set of points

$$\{x, f(x), f^2(x), \dots, f^n(x)\},$$

traversed in time n , and \mathcal{M}_x , the orbit of x , is the subset of all points of \mathcal{M} that can be reached by iterations of f . A periodic point (cycle point) x_k belonging to a periodic orbit (cycle) of period n is a real solution of

$$f^n(x_k) = f(f(\dots f(x_k) \dots)) = x_k, \quad k = 0, 1, 2, \dots, n - 1. \quad (3.15)$$

For example, the orbit of \hat{x}_1 in figure 3.5 is a set of four cycle points, $(\hat{x}_1, \hat{x}_2, \hat{x}_3, \hat{x}_4)$.

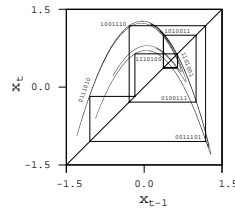


Figure 3.6: The strange attractor and an unstable period 7 cycle of the Hénon map (3.17) with $a = 1.4$, $b = 0.3$. The periodic points in the cycle are connected to guide the eye. (from K.T. Hansen [3.2])

The functional form of such Poincaré return maps P as figure 3.3 can be approximated by tabulating the results of integration of the flow from \hat{x} to the first Poincaré section return for many $\hat{x} \in \mathcal{P}$, and constructing a function that interpolates through these points. If we find a good approximation to $P(\hat{x})$, we can get rid of numerical integration altogether, by replacing the continuous time trajectory $f^t(\hat{x})$ by iteration of the Poincaré return map $P(\hat{x})$. Constructing accurate $P(\hat{x})$ for a given flow can be tricky, but we can already learn much from approximate Poincaré return maps. Multinomial approximations

$$P_k(\hat{x}) = a_k + \sum_{j=1}^d b_{kj} \hat{x}_j + \sum_{i,j=1}^d c_{kij} \hat{x}_i \hat{x}_j + \dots, \quad \hat{x} \in \mathcal{P} \quad (3.16)$$

to Poincaré return maps

$$\begin{pmatrix} \hat{x}_{1,n+1} \\ \hat{x}_{2,n+1} \\ \dots \\ \hat{x}_{d,n+1} \end{pmatrix} = \begin{pmatrix} P_1(\hat{x}_n) \\ P_2(\hat{x}_n) \\ \dots \\ P_d(\hat{x}_n) \end{pmatrix}, \quad \hat{x}_n, \hat{x}_{n+1} \in \mathcal{P}$$

motivate the study of model mappings of the plane, such as the Hénon map.

Example 3.6 Hénon map: The map

$$\begin{aligned} x_{n+1} &= 1 - ax_n^2 + by_n \\ y_{n+1} &= x_n \end{aligned} \quad (3.17)$$

is a nonlinear 2-dimensional map frequently employed in testing various hunches about chaotic dynamics. The Hénon map is sometimes written as a 2-step recurrence relation

$$x_{n+1} = 1 - ax_n^2 + bx_{n-1}. \quad (3.18)$$

An n -step recurrence relation is the discrete-time analogue of an n th order differential equation, and it can always be replaced by a set of n 1-step recurrence relations.

The Hénon map is the simplest map that captures the 'stretch & fold' dynamics of return maps such as Rössler's, figure 3.2. It can be obtained by a truncation of a polynomial approximation (3.16) to a Poincaré return map (3.16) to second order.

A quick sketch of the long-time dynamics of such a mapping (an example is depicted in figure 3.6), is obtained by picking an arbitrary starting point and iterating (3.17) on a computer. We plot here the dynamics in the (x_n, x_{n+1}) plane, rather than in the (x_n, y_n) plane, because we think of the Hénon map as a model return map $x_n \rightarrow x_{n+1}$. As we shall soon see, periodic orbits will be key to understanding the long-time dynamics, so we also plot a typical periodic orbit of such a system, in this case an unstable period 7 cycle. Numerical determination of such cycles will be explained in sect. 29.1, and the periodic point labels 0111010, 1110100, ... in sect. 12.2.

Example 3.7 Lozi map: Another example frequently employed is the Lozi map, a linear, 'tent map' version of the Hénon map (3.17) given by

$$\begin{aligned} x_{n+1} &= 1 - a|x_n| + by_n \\ y_{n+1} &= x_n. \end{aligned} \quad (3.19)$$

Though not realistic as an approximation to a smooth flow, the Lozi map is a very helpful tool for developing intuition about the topology of a large class of maps of the 'stretch & fold' type.

What we get by iterating such maps is—at least qualitatively—not unlike what we get from Poincaré section of flows such as the Rössler flow figure 3.3. For an arbitrary initial point this process might converge to a stable limit cycle, to a strange attractor, to a false attractor (due to roundoff errors), or diverge. In other words, mindless iteration is essentially uncontrollable, and we will need to resort to more thoughtful explorations. As we shall explain in due course, strategies for systematic exploration rely on stable/unstable manifolds, periodic points, saddle-saddle methods and so on. exercise 17.1

Example 3.8 Parabola: For sufficiently large value of the stretching parameter a , one iteration of the Hénon map (3.17) stretches and folds a region of the (x, y) plane centered around the origin, as will be illustrated in figure 12.4. The parameter a controls the amount of stretching, while the parameter b controls the thickness of the folded image through the '1-step memory' term bx_{n-1} in (3.18). In figure 3.6 the parameter b is rather large, $b = 0.3$, so the attractor is rather thick, with the transverse fractal structure clearly visible. For vanishingly small b the Hénon map reduces to the 1-dimensional quadratic map

$$x_{n+1} = 1 - ax_n^2. \quad (3.20)$$

exercise 3.6

By setting $b = 0$ we lose determinism, as on reals the inverse of map (3.20) has two real preimages $\{x_{n-1}^+, x_{n-1}^-\}$ for most x_n . If Bourbaki is your native dialect: the Hénon map is injective or one-to-one, but the quadratic map is surjective or many-to-one. Still, this 1-dimensional approximation is very instructive. (continued in example 11.5)

As we shall see in sect. 11.3, an understanding of 1-dimensional dynamics is indeed the essential prerequisite to unraveling the qualitative dynamics of many

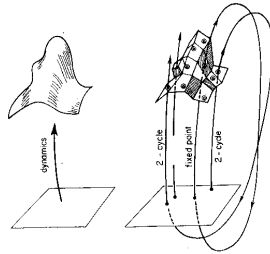
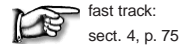


Figure 3.7: Reduction of a continuous-time flow (left frame) to a set of Poincaré maps (right frame), with a point on 1-cycle and the two cycle points of a 2-cycle used as template points.

higher-dimensional dynamical systems. For this reason many expositions of the theory of dynamical systems commence with a study of 1-dimensional maps. We prefer to stick to flows, as that is where the physics is.

appendix H.8



fast track:
sect. 4, p. 75

3.4 Charting the state space

In simple examples, such as the Rössler example 3.3, a single Poincaré section suffices, but this is rarely the case for flows of physical interest. In this section (skip it on first reading) we commence a discussion of the general case.



A Poincaré section is constructed by picking a ‘template’ point \hat{x}' within a state space region of interest, and defining a hypersurface (3.2) that goes through the template point. In theory, this Poincaré section could be any $(d-1)$ -dimensional manifold. In practice, a hyperplane (3.6) is the most convenient, the natural choice for the vector normal to the section being $\hat{n} = v(\hat{x}')$, the velocity field at the template point \hat{x}' . This Poincaré section $\hat{x} \in \mathcal{P}$ is a hyperplane,

appendix 13.4

$$v' \cdot (\hat{x} - \hat{x}') = 0, \quad v' = v(\hat{x}'), \tag{3.21}$$

normal to the flow direction v' at the template point \hat{x}' . Such section cuts the nearby trajectories transversally, and is a good description of solutions similar to the given template.

So one hyperspace \mathcal{P} will, in general, not suffice. A more insightful picture of the dynamics is obtained by partitioning the state space into N qualitatively distinct regions $\{\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_N\}$ and constructing a Poincaré section per region, global *atlas* of the state space composed of N local Poincaré sections $\mathcal{P}^{(j)}$ or *charts*, each one capturing a neighborhood of a qualitatively prominent state $\hat{x}^{(j)} \in S$. We shall refer to these states as *templates*, each represented in the state space \mathcal{M} of the system by a *template point* $\{\hat{x}^{(1)}, \hat{x}^{(2)}, \dots, \hat{x}^{(N)}\}$.

section 11.1

Our Poincaré section is a hyperplane. If we pick another template point $\hat{x}^{(2)}$, it comes along with its own section hyperplane. The $(d-1)$ -dimensional Poincaré sections for an adjacent pair of template intersects in a ‘ridge’ (‘boundary,’ ‘edge’), a $(d-2)$ -dimensional hyperplane, easy to compute. Follow an ant (the sequence of Poincaré map iterates) as it progresses along the Poincaré section $\mathcal{P}^{(1)}$. The moment $(\hat{x}^{(1)}(\tau) - \hat{x}^{(2)}) \cdot \hat{n}^{(2)}$ changes sign, the ant has crossed the ridge, we switch the Poincaré section, and the ant continues its merry stroll now confined to the $\mathcal{P}^{(2)}$ section. Each Poincaré section $\mathcal{P}^{(j)}$, provides a local chart at $\hat{x}^{(j)}$ for a neighborhood of an important, qualitatively distinct class of solutions; together they ‘Voronoi’ tessellate the curved manifold in which the reduced dynamics is replaced by a finite set of mappings between hyperplane tiles. An example is the periodic-orbit implementation of the idea of state space tessellation by neighborhoods of recurrent points, so dear to professional cyclists, illustrated in figure 3.7.

For a given dynamical flow, the physical task is to pick a minimal set of qualitatively distinct templates. The state space might be filled by all kinds of highly unstable, never revisited equilibria and relative periodic orbits. The choice of templates should reflect the dynamically prominent states seen in the long-time simulations of system’s dynamics. We have only vague advice on how to pick a single Poincaré section (see sect. 3.1.2), and no advice on how to systematically pick a set of ‘good’ templates, other than that the associated section tiles should be as large as possible, but still sufficiently small to exclude orbit tangencies, i.e., stop before crossing their section borders (3.7). Ideally, one would like to pick as few templates as possible in figure 3.7. Once templates are picked, the rest is geometry of hyperplanes, so checking whether the section border is on the far side of the tile edge (ridge between two sections) is a fast, linear computation.

There is a rub, though - you need to know how to pick the neighboring templates. Perhaps a glance at figure 3.7 helps visualize the problem; imagine that the tiles belong to the Poincaré sections through template points on these orbits. One could slide templates along their trajectories until the pairs of straight line segments connecting neighboring template points are minimized, but that seems a bit arbitrary. At this time we have no advice as how to ‘synchronize’ the templates relative to each other. The astute reader will instantly recognize this as the problem of ‘local gauge invariance’ or ‘gauge fixing’ of Quantum Field Theory and General Relativity.

3.4.1 Navigating the Poincaré-charted state space

Our goal now is to replace the continuous-time dynamics by a set of Poincaré maps between a set of hyperplane sections, as in figure 3.7. The flat hyperplane (3.6) is an *ad hoc* construct; one Poincaré section rarely suffices to capture all of the dynamics of interest. Instead we chart the state space by partitioning it into N qualitatively distinct regions $\{\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_N\}$. Successive trajectory intersections with the set of $(d-1)$ -dimensional hypersurfaces \mathcal{P}_s embedded in the d -dimensional state space \mathcal{M} , define the set of $(d-1) \rightarrow (d-1)$ Poincaré maps

section 11.1

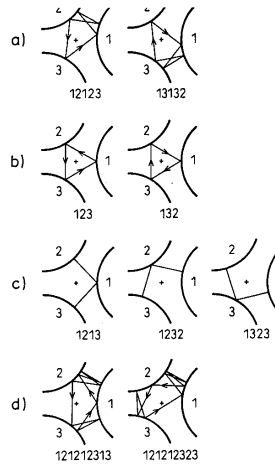


Figure 3.8: Some examples of 3-disk cycles: (a) 12123 and 13132 are mapped into each other by the flip across 1 axis. Similarly (b) 123 and 132 are related by flips, and (c) 1213 , 1232 and 1323 by rotations. (d) The cycles 121212313 and 121212323 are related by rotation and time reversal. These symmetries are discussed in chapter 9. (From ref. [3.1])

$$\hat{x}_{n+1} = P_{s_{n+1}s_n}(\hat{x}_n) = f^{\tau(\hat{s}_n)}(\hat{x}_n) \tag{3.22}$$

$$\hat{x}_{n+1} \in \mathcal{P}^{s_{n+1}}, \hat{x}_n \in \mathcal{P}^{s_n}, \quad s \in \{1, 2, \dots, N\}.$$

The d -dimensional continuous time flow is thus reduced to discrete time composition

$$P_{s_0 s_1 \dots s_n} = P_{s_n s_{n-1}} \circ \dots \circ P_{s_2 s_1} \circ P_{s_1 s_0}$$

of a set of Poincaré maps (3.22) that map the coordinates of Poincaré section \mathcal{P}_{s_n} to those of $\mathcal{P}_{s_{n+1}}$, the next section traversed by a given trajectory.

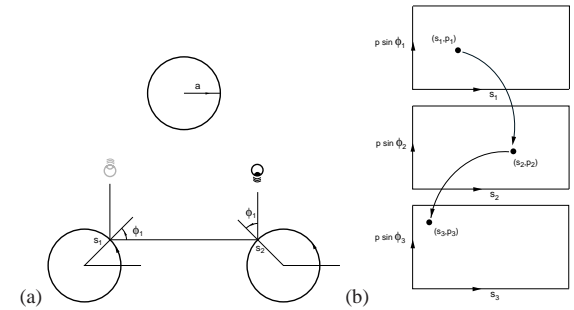
If a trajectory traverses regions $M_{s_0} \rightarrow M_{s_1} \rightarrow \dots \rightarrow M_{s_n}$, the sequence $s_0 s_1 \dots s_n = s_n \leftarrow \dots \leftarrow s_1 \leftarrow s_0$ is said to be *admissible*. The return map P_{s_0} from section \mathcal{P}_{s_0} to itself has a contribution from any admissible returning (periodic, $s_n = s_0$) sequence of compositions section 11.6

$$P_{s_0 s_1 \dots s_{n-1} s_0} = P_{s_0 s_{n-1}} \circ \dots \circ P_{s_2 s_1} \circ P_{s_1 s_0} \tag{3.23}$$

The next example offers an unambiguous set of such Poincaré sections which do double duty, providing us both with an exact representation of dynamics in terms of maps, and with a symbolic dynamics, a subject that we will return to in chapter 11. chapter 11

Example 3.9 Pinball game, Poincaré dissected. (continued from sect. 1.4) A phase-space orbit is fully specified by its position and momentum at a given instant, so no two distinct phase-space trajectories can intersect. The configuration space

Figure 3.9: (a) Poincaré section coordinates for the 3-disk game of pinball. (b) Collision sequence $(s_1, p_1) \mapsto (s_2, p_2) \mapsto (s_3, p_3)$ from the boundary of a disk to the boundary of the next disk is coded by the Poincaré maps sequence $P_{3 \leftarrow 2} P_{2 \leftarrow 1}$.



trajectories, however, can and do intersect, in rather unilluminating ways, as e.g. in figure 3.8 (d), and it can be rather hard to perceive the systematics of orbits from their configuration space shapes. The problem is that we are looking at the projections of 4-dimensional state space trajectories onto a 2-dimensional configuration subspace. A much clearer picture of the dynamics is obtained by constructing a set of Poincaré sections.

Suppose that the pinball has just bounced off disk 1. Depending on its position and outgoing angle, it could proceed to either disk 2 or 3. Not much happens in between the bounces—the ball just travels at constant velocity along a straight line—so we can reduce the 4-dimensional flow to a 2-dimensional map $P_{\sigma_n \leftarrow \sigma_{n-1}}$ that maps the coordinates (Poincaré section \mathcal{P}_1) of the pinball from one disk edge to another. Just after the moment of impact the trajectory is defined by s_n , the arc-length position of the n th bounce along the billiard wall, and $p_n = p \sin \phi_n$ the outgoing momentum component parallel to the billiard wall at the point of impact, figure 3.9 (a). These coordinates (due to Birkhoff) are smart, as they conserve the phase-space volume. Trajectories originating from one disk can hit either of the other two disks, or escape without further ado. We label the survivor state space regions $\mathcal{P}_{12}, \mathcal{P}_{13}$. In terms of the three Poincaré sections, one for each disk, the dynamics is reduced to the set of six maps example 8.6

$$(s_{n+1}, p_{n+1}) = P_{\sigma_{n+1} \leftarrow \sigma_n}(s_n, p_n), \quad \sigma \in \{1, 2, 3\} \tag{3.24}$$

from the boundary of a disk to the boundary of the next disk, figure 3.9 (b). The explicit form of this map is easily written down, see example 8.1, but much more economical is the symmetry quotiented version of chapter 9 which replaces the above 6 forward maps by a return map pair P_0, P_1 . (continued in chapter 8) chapter 8

Billiard dynamics is exceptionally simple - free flight segments, followed by specular reflections at boundaries, with billiard boundaries the obvious choice as Poincaré sections. For a general flow one is never so lucky. Also, so far we have discussed only flows with a 1 continuous parameter (the time). The general case of N -parameter continuous symmetries we postpone to chapter 10.

Résumé

In recurrent dynamics a trajectory exits a region in state space and then reenters it infinitely often, with finite return times. If the orbit is periodic, it returns after

a full period. So, on average, nothing much really happens along the trajectory—what is important is behavior of neighboring trajectories transverse to the flow. This observation motivates a replacement of the continuous time flow by iterative mapping, the Poincaré maps. A visualization of a strange attractor can be greatly facilitated by a felicitous choice of Poincaré sections, and the reduction of flow to Poincaré maps. This observation motivates in turn the study of discrete-time dynamical systems generated by iterations of maps.

A particularly natural application of the Poincaré section method is the reduction of a billiard flow to a boundary-to-boundary return map, described in chapter 8. As we shall show in chapter 6, further simplification of a Poincaré return map, or any nonlinear map, can be attained through rectifying these maps locally by means of smooth conjugacies.

In truth, as we shall see in chapter 10, the reduction of a continuous time flow by the method of Poincaré sections is not a convenience, but an absolute necessity - to make sense of an ergodic flow, all of its continuous symmetries must be reduced, evolution in time being one of these symmetries.

Commentary

Remark 3.1 Functions, maps, mappings. In mathematics, “mapping” is a noun, “map” is a verb. Nevertheless, “mapping” is often shortened to “map” and is often used as a synonym for “function.” “Function” is used for mappings that map to a single point in \mathbb{R} or \mathbb{C} , while a mapping which maps to \mathbb{R}^d would be called a “mapping,” and not a “function.” Likewise, if a point maps to several points and/or has several pre-images, this is a “many-to-many” mapping, rather than a function. In the sense used here, in the theory of dynamical systems, dynamical evolution from an initial state to a state finite time later is a (time-forward) map.

Remark 3.2 Determining a Poincaré section. The trick described in sect. 3.2 is due to Hénon [3.3, 3.4, 3.5]. The idea of changing the integration variable from time to one of the coordinates, although simple, avoids the alternative of having to interpolate the numerical solution to determine the intersection.

Remark 3.3 Hénon, Lozi maps. The Hénon map is of no particular physical import in and of itself—its significance lies in the fact that it is a minimal normal form for modeling flows near a saddle-node bifurcation, and that it is a prototype of the stretching and folding dynamics that leads to deterministic chaos. It is generic in the sense that it can exhibit arbitrarily complicated symbolic dynamics and mixtures of hyperbolic and non-hyperbolic behaviors. Its construction was motivated by the best known early example of ‘deterministic chaos,’ the Lorenz equation, see example 2.2 and remark 2.3. Y. Pomeau’s studies of the Lorenz attractor on an analog computer, and his insights into its stretching and folding dynamics motivated Hénon [3.6] to introduce the Hénon map in 1976. Hénon’s and Lorenz’s original papers can be found in reprint collections refs. [3.7, 3.8]. They are

a pleasure to read, and are still the best introduction to the physics motivating such models. Hénon [3.6] had conjectured that for $(a, b) = (1.4, 0.3)$ Hénon map a generic initial point converges to a *strange attractor*. Its existence has never been proven. While for all practical purposes this is a strange attractor, it has not been demonstrated that long time iterations are not attracted by some long attracting limit cycle. Indeed, the pruning front techniques that we describe below enable us to find stable attractors arbitrarily close by in the parameter space, such as the 13-cycle attractor at $(a, b) = (1.39945219, 0.3)$. A rigorous proof of the existence of Hénon attractors close to 1-dimensional parabola map is due to Benedicks and Carleson [3.9]. A detailed description of the dynamics of the Hénon map is given by Mira and coworkers [3.10, 3.11, 3.12], as well as very many other authors. The Lozi map (3.19) is particularly convenient in investigating the symbolic dynamics of 2-dimensional mappings. Both the Lorenz and Lozi [3.13] systems are uniformly expanding smooth systems with singularities. The existence of the attractor for the Lozi map was proven by M. Misiurewicz [3.14], and the existence of the SRB measure was established by L.-S. Young [3.15].

exercise 17.1

section 16.1

Exercises

3.1. **Poincaré sections of the Rössler flow.** (continuation of exercise 2.8) Calculate numerically a Poincaré section (or several Poincaré sections) of the Rössler flow. As the Rössler flow state space is $3D$, the flow maps onto a $2D$ Poincaré section. Do you see that in your numerical results? How good an approximation would a replacement of the return map for this section by a 1-dimensional map be? More precisely, estimate the thickness of the strange attractor. (continued as exercise 4.4)
(R. Paškauskas)

3.2. **A return Poincaré map for the Rössler flow.** (continuation of exercise 3.1) That Poincaré return maps of figure 3.3 appear multimodal and non-invertible is an artifact of projections of a 2-dimensional return map $(R_n, z_n) \rightarrow (R_{n+1}, z_{n+1})$ onto a 1-dimensional subspace $R_n \rightarrow R_{n+1}$. Construct a genuine $s_{n+1} = f(s_n)$ return map by parameterizing points on a Poincaré section of the attractor figure 3.2 by a Euclidean length s computed curvilinearly along the attractor section.

This is best done (using methods to be developed in what follows) by a continuation of the unstable manifold of the 1-cycle embedded in the strange attractor, figure 13.2 (b).

(P. Cvitanović)

3.3. **Arbitrary Poincaré sections.** We will generalize the construction of Poincaré sections so that they can have any shape, as specified by the equation $U(x) = 0$.

(a) Start by modifying your integrator so that you can change the coordinates once you get near the Poincaré section. You can do this easily by writing the equations as

$$\frac{dx_k}{ds} = \kappa f_k, \quad (3.25)$$

with $dt/ds = \kappa$, and choosing κ to be 1 or $1/f_1$. This allows one to switch between t and x_1 as the integration 'time.'

(b) Introduce an extra dimension x_{n+1} into your system and set

$$x_{n+1} = U(x). \quad (3.26)$$

How can this be used to find a Poincaré section?

3.4. **Classical collinear helium dynamics.** (continuation of exercise 2.10) Make a Poincaré surface of section by plotting (r_1, p_1) whenever $r_2 = 0$: Note that for $r_2 = 0$, p_2 is already determined by (7.8). Compare your results with figure 6.3 (b).

(Gregor Tanner, Per Rosenqvist)

3.5. **Hénon map fixed points.** Show that the two fixed points (x_0, x_0) , (x_1, x_1) of the Hénon map (3.17) are given by

$$\begin{aligned} x_0 &= \frac{-(1-b) - \sqrt{(1-b)^2 + 4a}}{2a}, \\ x_1 &= \frac{-(1-b) + \sqrt{(1-b)^2 + 4a}}{2a}. \end{aligned} \quad (3.27)$$

3.6. **Fixed points of maps.** A continuous function F is a contraction of the unit interval if it maps the interval inside itself.

(a) Use the continuity of F to show that a 1-dimensional contraction F of the interval $[0, 1]$ has at least one fixed point.

(b) In a uniform (hyperbolic) contraction the slope of F is always smaller than one, $|F'| < 1$. Is the composition of uniform contractions a contraction? Is it uniform?

3.7. **Section border for Rössler.** (continuation of exercise 3.1) Determine numerically section borders (3.7) for several Rössler flow Poincaré sections of exercise 3.1 and figure 3.2, at least for angles

(a) -60° , (b) 0° , and

(c) A Poincaré section hyperplane that goes through both equilibria, see (2.18) and figure 4.6. Two points only fix a line: think of a criterion for a good orientation of the section hyperplane, perhaps by demanding that the contracting eigenvector of the 'inner' equilibrium x_- lies in it.

(d) (Optional) Hand- or computer-draw a visualization of the section border as 3-dimensional fluid flow which either crosses, is tangent to, or fails to cross a sheet of light cutting across the flow.

As the state space is 3-dimensional, the section borders are 1-dimensional, and it should be easy to outline the border by plotting the color-coded magnitude of $v_{\perp}(\hat{x})$, component of the $v(\hat{x})$ normal to the section, for a fine grid of 2-dimensional Poincaré section plane points. For sections that go through the z -axis, the normal velocity $v_{\perp}(\hat{x})$ is tangent to the circle through \hat{x} , and vanishes for $\hat{\theta}$ in the polar coordinates (3.11), but that is not true for other Poincaré sections, such as the case (c).

(P. Cvitanović)

References

- [3.1] P. Cvitanović, B. Eckhardt, P. E. Rosenqvist, G. Russberg and P. Scherer, "Pinball Scattering," in G. Casati and B. Chirikov, eds., *Quantum Chaos* (Cambridge U. Press, Cambridge 1993).
- [3.2] K.T. Hansen, *Symbolic Dynamics in Chaotic Systems*, Ph.D. thesis (Univ. of Oslo, 1994); ChaosBook.org/projects/KTHansen/thesis.
- [3.3] M. Hénon, "On the numerical computation of Poincaré maps," *Physica D* **5**, 412 (1982).
- [3.4] N.B. Tufillaro, T.A. Abbott, and J.P. Reilly, *Experimental Approach to Non-linear Dynamics and Chaos* (Addison Wesley, Reading MA, 1992).
- [3.5] Bai-Lin Hao, *Elementary symbolic dynamics and chaos in dissipative systems* (World Scientific, Singapore, 1989).
- [3.6] M. Hénon, "A two-dimensional mapping with a strange attractor," *Comm. Math. Phys.* **50**, 69 (1976).
- [3.7] *Universality in Chaos*, P. Cvitanović, ed., (Adam Hilger, Bristol 1989).
- [3.8] Bai-Lin Hao, *Chaos* (World Scientific, Singapore, 1984).
- [3.9] M. Benedicks and L. Carleson, "The dynamics of the Hénon map," *Ann. of Math.* **133**, 73 (1991).
- [3.10] C. Mira, *Chaotic Dynamics—From one dimensional endomorphism to two dimensional diffeomorphism*, (World Scientific, Singapore, 1987).
- [3.11] I. Gumowski and C. Mira, *Recurrences and Discrete Dynamical Systems* (Springer-Verlag, Berlin 1980).
- [3.12] D. Fournier, H. Kawakami and C. Mira, *C.R. Acad. Sci. Ser. I*, **298**, 253 (1984); **301**, 223 (1985); **301**, 325 (1985).
- [3.13] R. Lozi, "Un attracteur étrange du type attracteur de Hénon," *J. Phys. (Paris) Colloq.* **39**, 9 (1978).
- [3.14] M. Misiurewicz, "Strange attractors for the Lozi mapping," *Ann. New York Acad. Sci.* **357**, 348 (1980).
- [3.15] L.-S. Young, "Bowen-Ruelle measures for certain piecewise hyperbolic maps," *Trans. Amer. Math. Soc.* **287**, 41 (1985).
- [3.16] W. S. Franklin, "New Books," *Phys. Rev.* **6**, 173 (1898); see www.ceafinney.com/chaos.
- [3.17] P. Dahlqvist and G. Russberg, "Existence of stable orbits in the x^2y^2 potential," *Phys. Rev. Lett.* **65**, 2837 (1990).