

Chapter 4

Applications in Dynamical Systems

4.1 Introduction

In this chapter we examine data from numerical simulations of some discrete dynamical systems. The goal is to illustrate how our computational topology tools apply in this context. Previous applications of computational topology in dynamical systems focus on flows and chaotic time series. Muldoon *et al.* [60] compute homology groups for embedded time series data; Mischaikow *et al.* [37, 59] also use homology computation in a numerical implementation of Conley index theory to experimental data; a number of groups [28, 57, 82] have used knot-theoretic ideas to model the dynamics of attractors in \mathbb{R}^3 . We focus on discrete dynamics rather than flows because there is less numerical error and greater efficiency in iterating a map than in solving a differential equation. Maps arise naturally from flows via Poincaré sections, for example, or through methods for their numerical solution, so their properties are closely related. While our techniques apply equally well to flows, every flow trajectory is connected, so only the higher order homology is interesting. There is a greater variety of topological structures in the phase space of a discrete map since an orbit can cover a disconnected set.

The three examples we study in detail are the Hénon attractor, the transition from invariant circle to Cantor set in the standard map, and cantori in a four-dimensional sawtooth map near the anti-integrable limit. These examples have well understood structure, so we are able to evaluate the effectiveness of our computational tools. We also show how our techniques could be applied to investigate the breakup of invariant tori in higher-dimensional symplectic maps, the structure of the chaotic region of the standard map, and for pruning outliers from embedded noisy time-series data. The examples are chosen to cover a wide range of dynamical and topological phenomena. They also require different approaches to approximating the underlying set and thereby convey the general applicability of the ideas from Chapters 2 and 3.

The Hénon map is the canonical two-dimensional quadratic map with constant Jacobian, in the sense that any other map in this class is conjugate to the Hénon map with some choice of parameters. In Section 4.2 we study the map at parameter values for which it has an attractor and use the minimal spanning tree techniques from Chapter 2 to study its connectedness properties. Our results give strong support to the common intuition that the Hénon attractor has Cantor-set cross-sections.

The standard map is an area-preserving twist map of the cylinder. It is the most commonly studied example of this class because it models a number of different physical problems, has an easily controlled perturbation from integrability, and exhibits most of the dynamical phenomena common to area-preserving twist maps. In Section 4.3 we investigate the topological change in

a quasiperiodic orbit as the perturbation parameter, k , is increased. Each quasiperiodic orbit has a critical k -value below which the orbit covers a circle, and above it, a Cantor set. The existence of invariant circles is important because they trap chaotic orbits and imply some degree of stability in the system. The transition from circle to Cantor set has been studied extensively, but this is the first time that the topology of these sets has been examined numerically. The minimal spanning tree tools of Chapter 2 successfully show the change in topology on a coarse k scale. A more precise criterion for determining the critical k -value is suggested by a closer analysis of how the longest MST edge scales with the number of points in the approximating orbit.

Symplectic maps are higher-dimensional versions of area-preserving maps and are discrete versions of Hamiltonian flows. The problem of torus breakup in nearly integrable symplectic maps is not well understood; this was one of the original motivating questions for the thesis. As a first step towards this goal, we examine a four-dimensional piecewise linear symplectic map near its anti-integrable limit in Section 4.4. It is known that at this extreme, orbits with incommensurate rotation frequencies cover Cantor sets called cantori since they are the remnants of invariant tori. We show that the cantori exhibit logarithmic rather than polynomial growth in the number of components, which is related to the fact that their Hausdorff dimension is zero. We discuss the potential of the computational topology tools from Chapter 3 to help understand torus breakup of nearly-integrable systems in Section 4.5.1.

The material in Sections 4.2 and 4.4 is published in [72].

4.2 The Hénon attractor

Figure 4.1 shows the much-studied Hénon attractor Λ , for the map H :

$$\begin{aligned} x_{n+1} &= y_n + 1 - ax_n^2 \\ y_{n+1} &= bx_n \end{aligned} \tag{4.1}$$

with parameter values $a = 1.4$ and $b = 0.3$. See Robinson [73] for a review of the properties of this map and its attractor. The first observation we make is that the set has a topological dimension of one and must be connected. This follows from the fact that the attractor is the closure of the unstable manifold of a fixed point for H . The attractor has fractal structure nonetheless, and is often described as having a Cantor set cross-section [73]. We use the minimal spanning tree techniques from Chapter 2 to investigate the above topological properties.

To generate finite-point approximations to the attractor, we compute the orbit, S , of a single initial point (x_0, y_0) , from the trapping region for the Hénon attractor. For our connectedness data to be valid, we need some guarantee that this gives a good approximation, i.e., that the Hausdorff distance, $d_H(S, \Lambda) < \rho$. Firstly, this orbit must converge to the attractor, so given $\rho > 0$, there is an integer n , such that for $j > n$, $d(H^j(x_0, y_0), \Lambda) < \rho$. This means we have to iterate the map a few hundred times before we start recording the points visited by the orbit. The length of these transients depends on the strength of contraction along the stable manifolds of the attractor. If an attractor is topologically transitive, then a typical orbit will fill out the entire attractor. There is no proof that this is the case for the Hénon attractor with the parameter values $a = 1.4$ and $b = -0.3$, but numerical experiments suggest that there is a dense orbit; see [73] for more discussion.

The graphs in Figure 4.2 show the number of connected components, $C(\epsilon)$, number of isolated points, $I(\epsilon)$, and the largest component diameter, $D(\epsilon)$, for two orbits, one with 10^4 points and one with 5×10^4 . These graphs are exactly what we expect to see for data that cover a connected set in a slightly nonuniform fashion. For the orbit with 10^4 iterates, isolated points

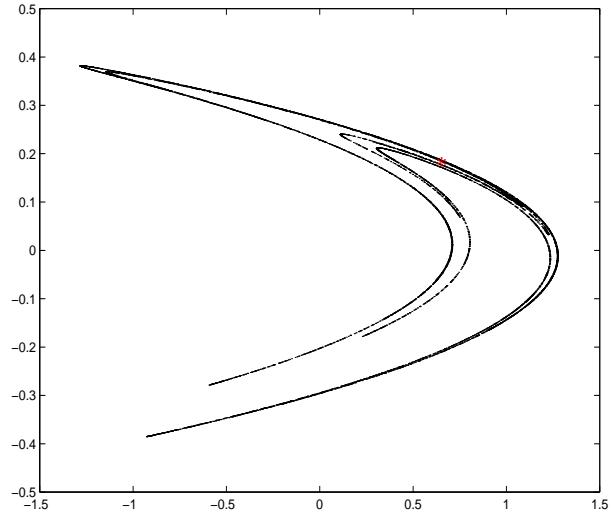


Figure 4.1: An orbit on the Hénon attractor. The red star is the unstable fixed point of H with coordinates $(x, y) \approx (0.631, 0.189)$.

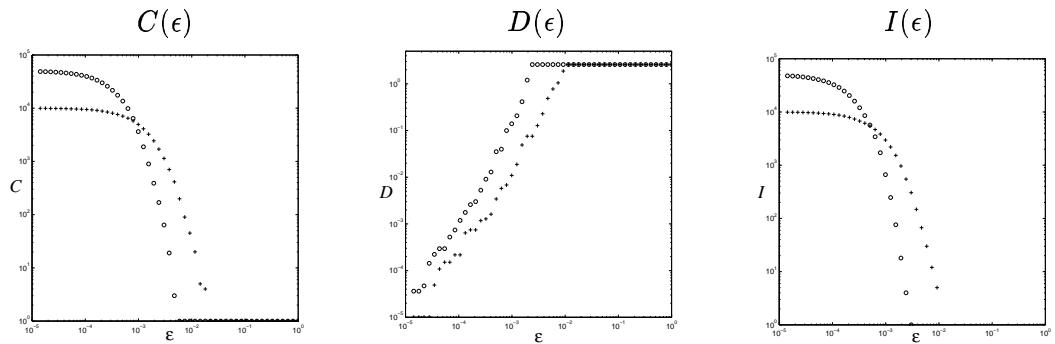


Figure 4.2: $C(\epsilon)$, $D(\epsilon)$, and $I(\epsilon)$ data for two orbits on the Hénon attractor. The crosses, $+$, represent calculations for the orbit of 10^4 iterates and the circles, \circ , are for an orbit with 5×10^4 points. All axes are logarithmic. The horizontal axis range is $10^{-5} < \epsilon < 1$.

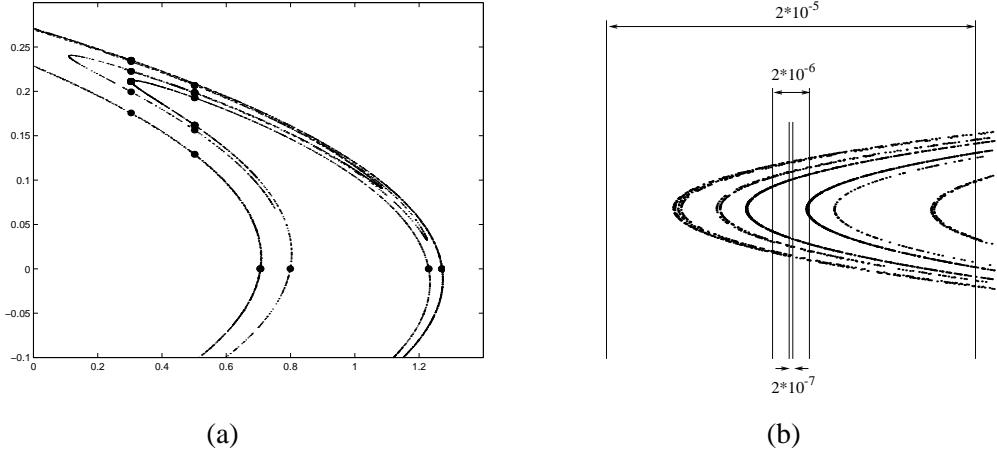


Figure 4.3: (a) A close-up of the Hénon attractor. The bold face dots are points in the three cross-sections considered in the text: slices at $x = 0.302435$, $x = 0.5$ and $y = 0$. (b) A small part of the slice at $x = 0.302435$, $y = 0.22$ that shows the folding of the attractor. The pairs of vertical lines in (b) are the boundaries of the different sub-slices of widths 2×10^{-5} , 2×10^{-6} and 2×10^{-7} .

are resolved at $\rho \approx 10^{-2}$. The number of components at this resolution is actually around ten, so this value of ρ underestimates the true cutoff resolution. The graph of $D(\epsilon)$ shows that the largest component diameter does not change until $\epsilon < \rho$, and then it decreases sharply to zero, which is what we see for other connected sets. The orbit with 10^5 points has a smaller cutoff resolution, $\rho \approx 3 \times 10^{-3}$, and qualitatively similar graphs, which increases our confidence in the numerical results.

To investigate the Cantor set cross-sections, we compute $C(\epsilon)$, $D(\epsilon)$, and $I(\epsilon)$ for thin slices taken through the attractor at three different places: $x = 0.302435$, $x = 0.5$ and $y = 0$ (the bold dots in Figure 4.3(a)). This process is related to taking the Poincaré section — a common technique for visualizing the structure of attractors from flows. To create a Poincaré section, a surface of codimension one is chosen and points on the section are recorded whenever the trajectory pierces the surface. The Hénon attractor is generated by iterating a map, so it is not possible to find many points on a given section. Instead, the sections are generated by recording points that fall within an interval of the given section coordinate. This means that the slices have a finite width and the data are still two-dimensional. To be confident that the observed scaling behavior is approximating that of a one-dimensional Cantor set, we compute $C(\epsilon)$, $D(\epsilon)$, and $I(\epsilon)$ for four successively narrower slices at each cross-section. The thinnest slice in each case has a width of 2×10^{-7} .

The section at $x = 0.302435$ is interesting because it cuts through a fold in the attractor. This folding of the Hénon attractor is the source of its nonuniform hyperbolic structure. Note that a one-dimensional cross-section that touches a fold will have an isolated point and therefore cannot be a Cantor set. A close-up of such a fold is shown in Figure 4.3(b). This figure shows that slices of different widths taken at this x value capture different folding structure. This is reflected in the $C(\epsilon)$ data in figure 4.4. The data for different slices at $x = 0.302435$ does not coincide exactly for $\epsilon > \rho$, as it does for the other two sections at $x = 0.5$ and $y = 0$, which show no folding at these resolutions. Note that we can reverse this observation and use the inconsistency to detect cross-sections that touch a fold. The sections at $x = 0.5$ and $y = 0$ have simpler structure. The graphs of $D(\epsilon)$ show the now-familiar staircase structure of a Cantor set.

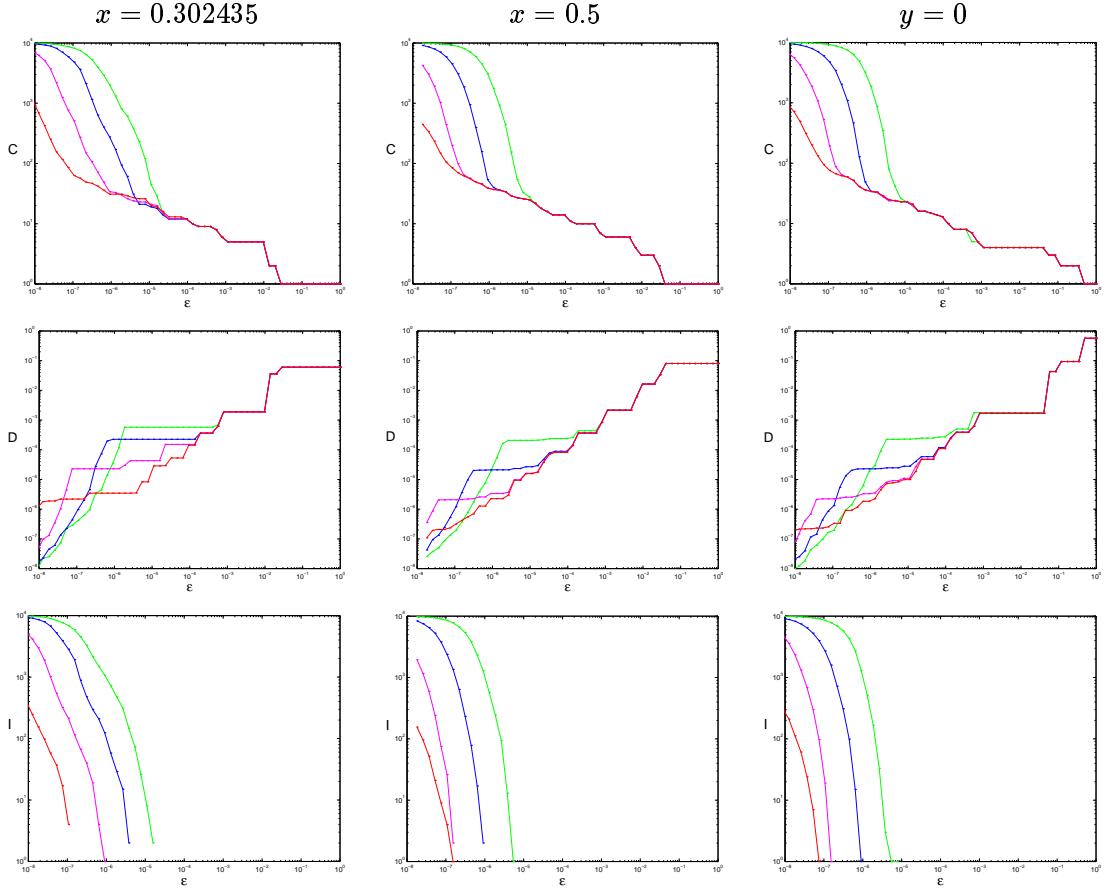


Figure 4.4: $C(\epsilon)$, $D(\epsilon)$, and $I(\epsilon)$ data for three sections of the Hénon attractor. The left column shows data for $x = 0.32435$, the middle column for $x = 0.5$, and the right for $y = 0$. The curves are colored according to the width of the slice: 2×10^{-4} is green, 2×10^{-5} is blue, 2×10^{-6} is magenta, and 2×10^{-7} is red. All axes are logarithmic. The horizontal axis range is $10^{-8} < \epsilon < 1$.

Table 4.1: Values of γ and δ for the three sections of the Hénon attractor shown in Figure 4.3.

section	γ	δ
$x = 0.302435$	0.25 ± 0.01	0.9 ± 0.1
$x = 0.5$	0.26 ± 0.01	0.85 ± 0.04
$y = 0$	0.27 ± 0.01	0.88 ± 0.02

The flat segments in each graph of $D(\epsilon)$ are due to the finite width of each slice, making the data a Cantor set of line segments.

Values of γ and δ are calculated from the $C(\epsilon)$ and $D(\epsilon)$ data for the thinnest slice at each section. The results are summarized in Table 4.1. The multifractal nature of the Hénon attractor [86] means that we expect the dimension to vary for different cross-sections. For the three examples given here, though, the variation is not significant. It is possible that the multifractal nature would become apparent if we computed γ from thinner slices with more data. However, the multifractal spectrum is determined by choosing points on the attractor and computing local scaling rates. The sections are not local and this may obscure the multifractal nature of the attractor.

The above results give strong numerical support for the common belief that cross-sections of the Hénon attractor are Cantor sets. The box-counting dimension of the Hénon attractor is estimated to be about 1.27 [86, 63]. Results on the dimension of intersections of sets [23] imply that the dimension of a cross section through the Hénon attractor should be $1.27 - 1 = 0.27$. The values of γ given in Table 4.1 are in close agreement with this value, providing further support for our conjecture that Cantor sets of zero measure have γ equal to the box-counting dimension.

4.3 Circle breakup in the standard map

In this section we examine orbits from the standard map as a parameter controlling the non-linear perturbation is increased. This map is a popular example of the class of area-preserving twist maps of the cylinder, which are closely related to Hamiltonian flows with two degrees of freedom. The properties of these maps are reviewed in [55].

Our goal is to see how the computational topology tools from Chapter 2 perform in detecting a transition from circle to Cantor set. This transition is interesting dynamically because the invariant circles trap chaotic orbits and therefore imply some degree of stability. When an invariant circle is destroyed, it is possible for nearby chaotic orbits to diffuse through the gaps in the remaining Cantor set (see Figure 4.17). If no circles exist, then it is possible for a single chaotic orbit to access most of the phase space. Many approaches to detecting the existence of invariant circles have been explored in the last twenty or so years, [8, 31, 44, 49, 53]. These techniques exploit properties such as the stability, flux, and frequency of orbits; our work is the first time the topology has been used directly.

For our computational tools to give valid results, we need good finite point approximations to the underlying circle or Cantor set. The quasiperiodic orbits which cover these sets are very difficult to find numerically because their position in phase space is uncertain and they do not attract other orbits. Instead, we use long periodic orbits that are close to quasiperiodic ones, since symmetry properties of the standard map make these relatively easy to find. This method

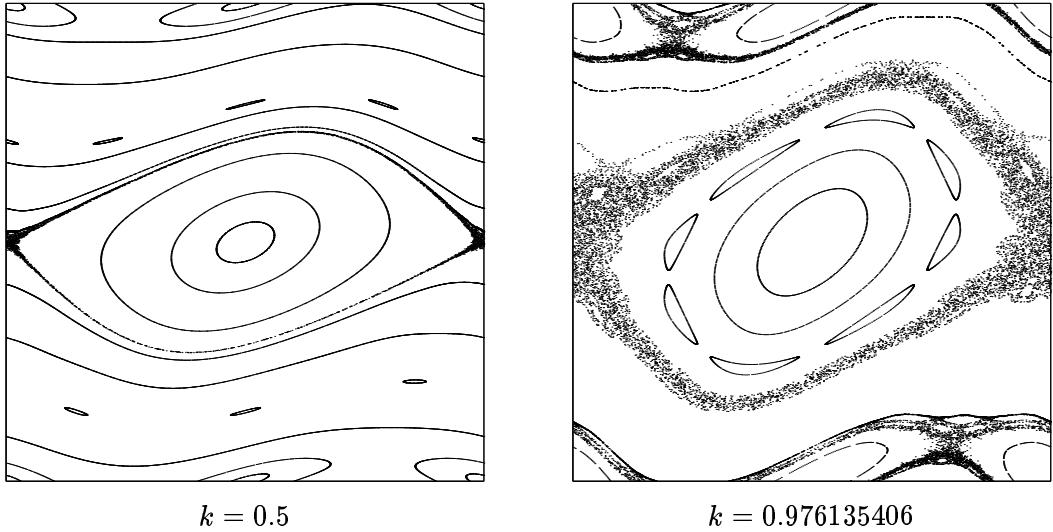


Figure 4.5: Phase space portraits for the standard map at two different k values.

of approximation was pioneered by Greene [31], and is justified by Aubry-Mather theory [46].

We begin this section by describing some of the basic phenomenology of the standard map, and reviewing the essential details of Aubry-Mather theory and the method of approximating quasiperiodic orbits by symmetric periodic orbits. For the purposes of comparison, we also summarize some previous numerical work on detecting the breakup of invariant circles. We then present connectedness data from the minimal spanning tree techniques of Chapter 2 applied to periodic orbit approximations of an invariant circle. The results obtained are not as sensitive at detecting the transition as are some previous techniques, so we investigate scaling properties of the $C(\epsilon)$ distributions in more detail. This work raises many questions about the connections between the distribution of points and the distribution of edge-lengths in their minimal spanning tree, and about the renormalization of phase space near a critical circle. This is part of a larger project that is outside the scope of this thesis.

4.3.1 Background

The standard map

$$\begin{aligned} y_{n+1} &= y_n - \frac{k}{2\pi} \sin(2\pi x_n) \\ x_{n+1} &= x_n + y_{n+1} \mod 1. \end{aligned} \tag{4.2}$$

appears in a number of physical applications. It models charged particle motion in a simple particle accelerator, the cyclotron, and is equivalent to the Frenkel-Kontorova model of a chain of atoms in a one dimensional periodic potential. The standard map is also the result of applying a first order numerical solver to the pendulum equation $\ddot{\theta} - K \sin \theta = 0$, with $x = \theta/(2\pi)$, $y = \dot{x}$, and $t = n$. The x variable is an angle in $[0, 1)$ and y is analogous to momentum. Clearly y is also periodic, so the phase space is in fact homeomorphic to the torus and we need only consider orbits in the unit square.

We label the orbits by their rotational frequency, ω , the average change in x per iteration of

the map:

$$\omega = \lim_{n \rightarrow \infty} \frac{x_n - x_0}{n}. \quad (4.3)$$

To calculate ω , we iterate $x_{n+1} = x_n + y_{n+1}$, without taking the fractional part (this corresponds to taking a lift of the map of the cylinder (4.2) and making it a map of the covering space, \mathbb{R}^2). The twist condition,

$$\frac{\partial x_{n+1}}{\partial y_n} \geq C > 0, \quad (4.4)$$

means that orbits with larger y have larger rotational frequency. This is compatible with the physical interpretation of y as a momentum variable. Periodic orbits satisfying $y_q = y_0$, $x_q = x_0 + p$, have rational rotation frequency p/q and are called (p, q) -periodic orbits. Quasiperiodic orbits have irrational rotation frequency. Note that the limit in (4.3) may not exist, for example, in the case the orbit is chaotic.

When the nonlinearity parameter k is zero, the map is integrable. Its orbits lie on invariant circles $y = y_0$ and evolve as rigid rotations of the angle variable, $x_{n+1} = x_n + y_0$, so that $\omega = y_0$. When $k > 0$, the geometric structure of the orbits becomes dramatically more complicated; this is illustrated in Figure 4.5. The Poincaré-Birkhoff theorem shows that every invariant circle with rational rotation frequency is destroyed, leaving a pair of periodic orbits with that frequency, one elliptic and one hyperbolic. Circles with irrational rotation frequency are guaranteed to persist for a small interval of k values by the Kolmogorov-Arnold-Moser (KAM) theorem. See Arrowsmith and Place [3] for more details about these results.

Periodic and quasiperiodic orbits

The existence of a special class of periodic and quasiperiodic orbits in any area-preserving twist map follows from Aubry-Mather theory; see [55] for an overview. These orbits minimize a quantity called the action and are monotone (the order of points is preserved under iteration). An extremely useful result from Aubry-Mather theory is that if a sequence of rationals converges to an irrational, $p_n/q_n \rightarrow \omega$, then the minimizing (p_n, q_n) -periodic orbits converge (in the Hausdorff metric) to a set of minimizing orbits with rotation number ω . Aubry-Mather theory also shows that this limit set is either an invariant circle or a Cantor set. Since it is relatively easy to find minimizing periodic orbits in the standard map, this result means we can use sequences of these orbits to approximate a quasiperiodic one. This approach, first used by Greene [31], will form the basis of approximation in our numerical investigation of invariant circles and Cantor sets in Sections 4.3.2 and 4.3.3.

It is the symmetry properties of the standard map that make minimizing (p, q) -period orbits easy to find (see [47] for more details). Each such orbit must have a point on two of the following four symmetry lines:

$$a = \{z = (x, y) \mid x = 0\} \quad (4.5)$$

$$b = \{z = (x, y) \mid x = 1/2\} \quad (4.6)$$

$$c = \{z = (x, y) \mid x = y/2\} \quad (4.7)$$

$$d = \{z = (x, y) \mid x = (y + 1)/2\}. \quad (4.8)$$

The pair of lines is determined by p and q . For action-minimizing orbits in the standard map the observed pattern is:

$$p \text{ odd, } q \text{ even} \Rightarrow z_0 \in c \text{ and } z_{q/2} \in d \quad (4.9)$$

$$p \text{ odd, } q \text{ odd} \Rightarrow z_0 \in b \text{ and } z_{(q+1)/2} \in c \quad (4.10)$$

$$p \text{ even, } q \text{ odd} \Rightarrow z_0 \in b \text{ and } z_{(q+1)/2} \in d. \quad (4.11)$$

These properties mean the search for an orbit is made along a single line, rather than over the whole of phase space. We use these types of orbits in Section 4.3.2.

In practice, we want the sequence of rationals p_n/q_n to converge to the irrational ω as fast as possible. This is achieved by using the continued fraction convergents of ω (see [55] for a description of how to generate continued fraction expansions). The continued fraction convergents are the “best approximants” in the sense that if p/q is a convergent of ω , then any other rational p'/q' , with $q' \leq q$ is further from ω , i.e., $|\omega - p'/q'| > |\omega - p/q|$. The number with the slowest convergence of its best approximants is the golden mean,

$$\mu = \frac{1 + \sqrt{5}}{2} = 1 + \cfrac{1}{1 + \cfrac{1}{1 + \cfrac{1}{1 + \dots}}} = 1.61803\dots \quad (4.12)$$

The convergents, p_n/q_n , of μ are related to the Fibonacci sequence: $F_{n+1} = F_n + F_{n-1}$, with $F_0 = F_1 = 1$, by setting $p_n = F_{n+1}$, and $q_n = F_n$.

The invariant circle with golden mean rotation frequency plays an important role in the standard map, since it is the last one to be destroyed, i.e., its transition from circle to Cantor set occurs at a k -value larger than that of any other quasiperiodic orbit. This stems from the status of the golden mean as the number that is hardest to approximate by rationals, and therefore the orbit for which the KAM small divisor problems are minimal.

Renormalization near a critical circle

The phase space near an invariant circle at its critical k -value (i.e. the point of breakup) has a remarkable degree of local self-similarity [31, 75]. MacKay [47] studies this extensively and quantifies the scaling relations using renormalization techniques. We observe related scaling properties in the MST data, as discussed in Section 4.3.3.

As an illustrative example, we describe results for the standard map and an orbit with rotation frequency of $1/\mu^2 = 0.38196\dots$. This number is closely related to the golden mean and it has continued fractions convergents, p_n/q_n with $p_n = F_n$, $q_n = F_{n+2}$; it also has the same critical value of $k = k_c = 0.971635406$.

Near the point where the golden circle crosses $x = 0$, the phase space is asymptotically self-similar with a scale factor of $\alpha_0 = -1.4148360$ in x and $\beta_0 = -3.0668882$ in y . The negative signs indicate a flip in orientation which is caused by the oscillation of successive convergents: $p_n/q_n < \omega < p_{n+1}/q_{n+1}$.

At $x = 1/2$, the scaling is a little more complicated. Here, the structure of the period orbits cycles with period 3. Shenker and Kadanoff [75] explain this in terms of the symmetry properties of the orbits. The Fibonacci numbers,

$$F_i = 1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, 377\dots$$

follow a repeating pattern of odd, odd, even. This implies that the minimizing (p_n, q_n) -periodic orbits that converge to the $1/\mu^2$ circle cycle through the three symmetry patterns of (4.9). The renormalization must take this into account and MacKay [47] computes 3-step ratios of $\alpha_{.5} = -4.84581$ in x and $\beta_{.5} = -16.8597$. Since area is preserved, these values are related to α_0 and β_0 via:

$$4.3391 = \alpha_0\beta_0 = \sqrt[3]{(\alpha_{.5}\beta_{.5})} = (1.6922)(2.5642) = 4.3391. \quad (4.13)$$

Although we described this scaling in terms of a particular critical circle, the scaling holds for any circle with a rotation frequency that has a continued fraction expansion that ends in ones [47].

Nonexistence of invariant circles

Invariant circles that go around the cylinder (i.e. ones that are homotopic to $y = 0$) are called rotational invariant circles. Their existence is important because they confine chaotic orbits to a small region of phase space and imply that the system is essentially stable. If no rotational invariant circles exist for any rotation frequency, it is possible for a single chaotic orbit to cover most of the phase space. This means that the dynamics is inherently unstable; in particular, the momentum variable, y , can increase without bound. Invariant circles are observed to exist for much larger values of k than those given by KAM theory. Many methods exist for estimating the critical value, k_c , such that for $k > k_c$, no KAM circles remain; we outline three computational approaches below. The most accurate numerical determination of k_c is due to MacKay [47, p.199], who uses Greene's criterion and renormalization techniques to estimate $k_c = 0.971635406$. The best rigorous bound on k_c for the standard map is due to MacKay and Percival [50] who show that there are no invariant circles for $k > 63/64 = 0.984375$.

A very direct method due to Chirikov [8] estimates k_c by measuring diffusion times of chaotic orbits. For different values of k , and initial conditions, (y_0, x_0) with $y_0 \approx 10^{-3}$, he records the number of iterations, N , it takes for the orbit to reach $y \approx 0.5$. There is a sharp increase in N for values of k near k_c ; the longest orbits considered have $N = 10^7$. By fitting the data obtained for $N(k)$, he estimates $k_c \approx 0.989$. This is surprisingly close to the above values, given the inaccuracies of these numerical techniques.

More recently, Laskar [44] has developed a general method for detecting the destruction of invariant circles by computing the frequency of orbits with different initial conditions. Recall that the twist condition implies that along a vertical line, $x = 0$ say, orbits with larger initial y coordinate have larger frequency. Laskar therefore uses initial conditions along such a vertical line and looks for non-monotonicity in the computed frequencies. For the standard map he finds k_c to five significant figures using orbits of around 10^6 points. This makes it computationally intensive, but the approach generalizes readily to higher-dimensional maps.

The most sensitive numerical technique is probably Greene's mean residue criterion [31]. This approach uses the stability of long periodic orbits that are close to a quasiperiodic one to indicate the existence of an invariant circle. The residue of a (p, q) -periodic orbit is a measure of stability related to the eigenvalues: $R = \frac{1}{4}(2 - \text{Trace } M)$, where M is the matrix linearization of the q -fold composition of the map. The mean residue is a type of geometric mean, $f = (|R|/\beta)^{1/q}$, where β is a scale factor used to improve convergence. To extend the definition of f to quasiperiodic orbits, Greene uses the continued fraction convergents p_n/q_n of an irrational ω . The value of $f(\omega)$ is found by taking the limit: $f(\omega) = \lim_{n \rightarrow \infty} f(p_n/q_n)$. For the standard map, $\beta = 1/4$, and Greene's criterion is that the orbit with rotation number ω covers a circle if and only if $f(\omega) < 1$. An attempt to generalize this criterion to higher-dimensional symplectic

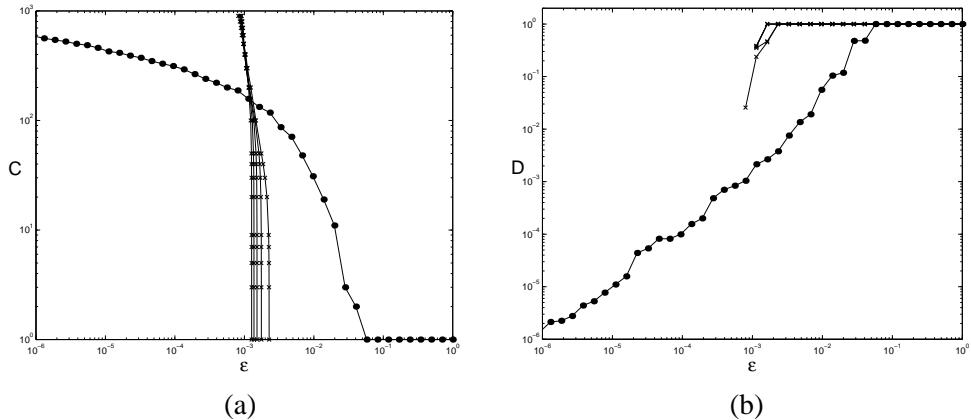


Figure 4.6: Connected component data for a $(377, 987)$ -periodic orbit approximation to the golden $1/\mu^2$ orbit at different values of k . In both graphs, the data marked with solid circles, \cdot , is for $k = 1.0$. The crosses, \times , mark data for the other k values, $k = 0.5, 0.6, 0.7, 0.8, 0.9$. The cut-off resolution for the $k = 1.0$ data is $\rho \approx 6 \times 10^{-5}$. (a) Number of components as a function of resolution, ϵ . (b) Diameter of the largest component as a function of ϵ .

twist maps is made in [70]. A formal justification for this approach is hampered by the fact that there is no corresponding Aubry-Mather theory in higher dimensions. Greene's criterion implies that the golden invariant circle is destroyed at $k = k_c = 0.971635$. This value of k_c is determined in [31] using orbits of length around 10^4 .

4.3.2 MST analysis of the transition from circle to Cantor set

In this section we investigate the transition from invariant circle to Cantor set using the computational tools based on the minimal spanning tree from Chapter 2. As one might suspect, these techniques cannot make as precise a determination of k_c as Greene's residue criterion does, for example. This is because there is a continuous metric change from invariant circle to Cantor set; the results of Chapter 3 imply that the ϵ -neighborhoods have correspondingly close topology. However, our connectedness techniques do show the transition from circle to Cantor set on a coarse k -scale.

We start by examining the structure of a single orbit in the sequence of periodic approximations to the $1/\mu^2$ orbit as k is changed. The data in Figure 4.6 is for the minimizing (377,987)-periodic orbit approximation for $k = 0.5, 0.6, \dots, 0.9, 1.0$. There is a significant difference between the $C(\epsilon)$ and $D(\epsilon)$ graphs for $k \leq 0.9$ and $k = 1.0$, consistent with the fact that $0.9 < k_c < 1.0$. For $k \leq 0.9$, the $C(\epsilon)$ and $D(\epsilon)$ data look like that of a connected set. For $k = 1.0$, the cutoff resolution drops to 6×10^{-5} and $D(\epsilon) \rightarrow 0$, which suggests a totally disconnected set. We estimate δ as the slope of the least squares linear fit for $10^{-4} \leq \epsilon \leq 10^{-2}$, and find $\delta \approx 1.15 \pm 0.05$. The growth in the number of components, $C(\epsilon)$, resembles that of a Cantor set with $\gamma = 0$ (recall Section 2.2.3).

Similar results are obtained with a finer k scale, as shown in Figure 4.7. Here, we present connected component data for the period-987 orbit with $k = 0.90, 0.91, \dots, 1.0$. Again, we see a significant change in the $C(\epsilon)$ and $D(\epsilon)$ curves as k increases. This time, for $k \leq 0.97$, the orbit appears to approximate a connected set, and for $k \geq 0.98$, a Cantor set. The above results give an estimate of k_c to two significant figures using an orbit of just 987 points. To obtain higher precision we need to use longer orbits; the different topological structure of the golden

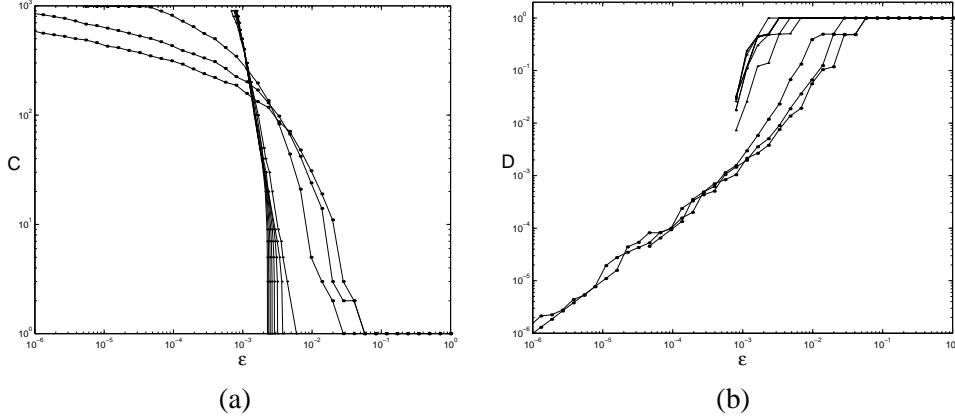


Figure 4.7: Connected component data for a $(377, 987)$ -periodic orbit approximation to the golden orbit for $k = 0.90, 0.91, \dots, 0.99, 1.0$. (a) Number of components as a function of resolution, ϵ . (b) Diameter of the largest component as a function of ϵ . The change in the shape of the curves suggests that the transition from circle to cantorus occurs between $k = 0.97$ and 0.98 .

orbit at $k = 0.971$ and $k = 0.972$, for example, is revealed for approximations of around 10^4 points.

The graphs in Figure 4.8 show the ambiguity that occurs when the approximation is not good enough. In each column, we fix k and plot connectedness data for a sequence of successively longer periodic orbit approximations ($q_n = 377, \dots, 4181$) to the golden orbit. The data for $k = 0.97$ show that the orbits are probably converging to a connected set, as expected. For $k = 0.98$, it is clear from the persistence of the long MST edges, and the trend $D(\epsilon) \rightarrow 0$, that the orbits are converging to a Cantor set. At $k = 0.972$, however, the data is misleading since the orbits appear to be converging to a connected circle even though the $1/\mu^2$ orbit covers a Cantor set. The problem is that the largest gap in the Cantor set is shorter than the cutoff resolution for the data; see Figure 4.9.

The graphs of $C(\epsilon)$ for $k = 0.97$ and $k = 0.972$ in Figure 4.8 appear to keep the same form as q_n increases. The graph of $C(\epsilon)$ is really a cumulative distribution of edge lengths from the MST. It may be possible, therefore, to obtain more accurate information about the underlying topology by analyzing the relative distribution of edge lengths. We examine how distributions of MST edges change with k and the period of the orbit in the following section.

4.3.3 Scaling of the gaps

To explore the gap distributions, we again use the minimizing (p_n, q_n) -periodic orbit approximations to the $1/\mu^2$ orbit. We can simplify the numerical computation of the MST by using a result from Aubry-Mather theory which shows that a monotone invariant set (i.e., minimizing periodic and quasiperiodic orbits) must be a graph over x [55]. This lets us order the q points in an orbit $z_i = (x_i, y_i)$ by their x -coordinate, i.e., $x_i < x_j$ if $i < j$. The lengths of the MST edges are then given by the ‘‘gaps,’’ $g_i = d(z_{i+1}, z_i)$ for $i = 0, \dots, q - 1$, and $g_q = d(z_0, \tilde{z}_q)$, where $\tilde{z}_q = (x_q - 1, y_q)$. In fact, the q gaps, g_i , give us the distance between every neighboring pair of points, whereas the MST, by its definition, omits the largest gap. We start by examining how the maximum, minimum, and mean of the gap lengths scale with the length of the periodic orbit approximations at different values of k . When $k = k_c$, our results are closely related to the local

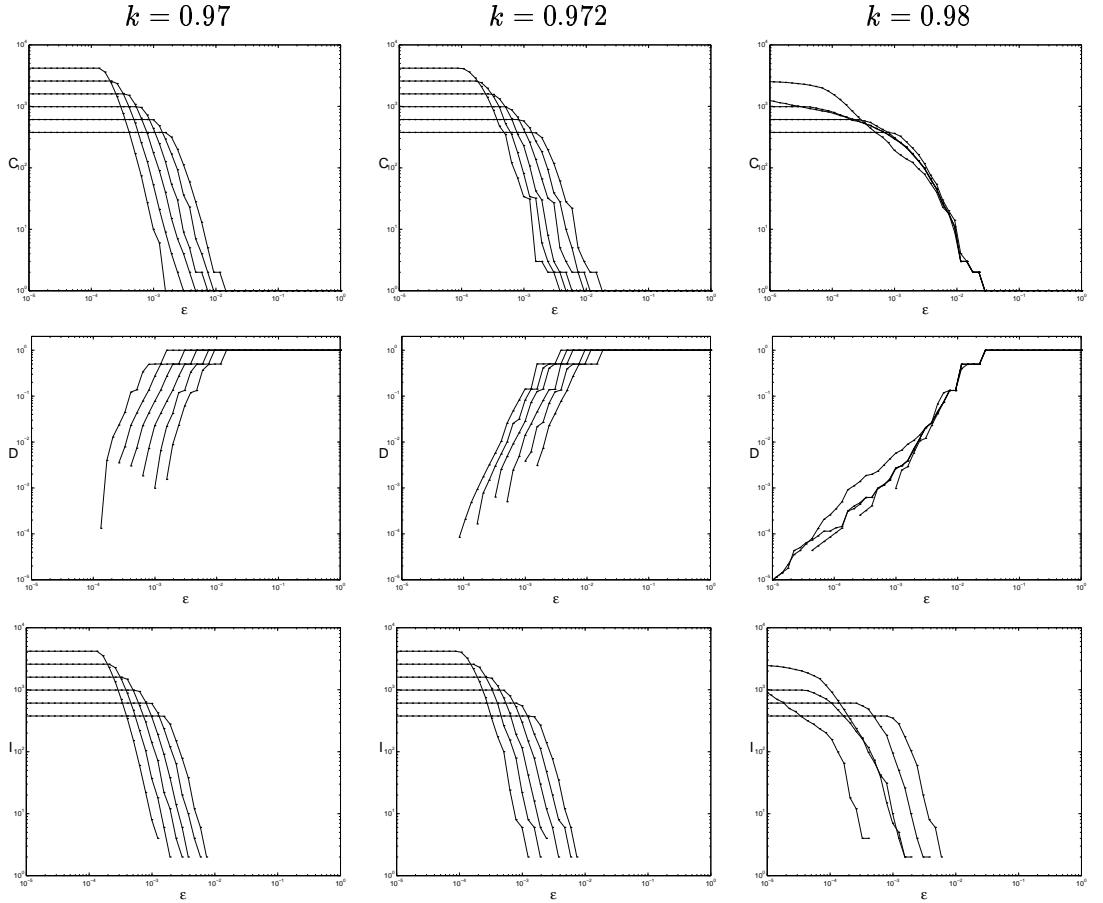


Figure 4.8: $C(\epsilon)$, $D(\epsilon)$, and $I(\epsilon)$ data for three values of k in the standard map. The left column has $k = 0.97$, the middle column, $k = 0.972$, and the right one, $k = 0.98$. In each case the data is obtained for a sequence of (p_n, q_n) -periodic orbit approximations to the $1/\mu^2$ orbit, with $q_n = 377, 610, 987, 1597, 2584, 4181$. For $k = 0.97$, the data suggests that the orbits are approximating a circle, and for $k = 0.98$, a Cantor set. When $k = 0.972$ is just larger than the critical value $k_c = 0.971635$, the golden orbit is a Cantor set, but the data given suggests the periodic orbits are converging to a circle. This is because the orbits are not long enough to allow us to resolve the gaps.

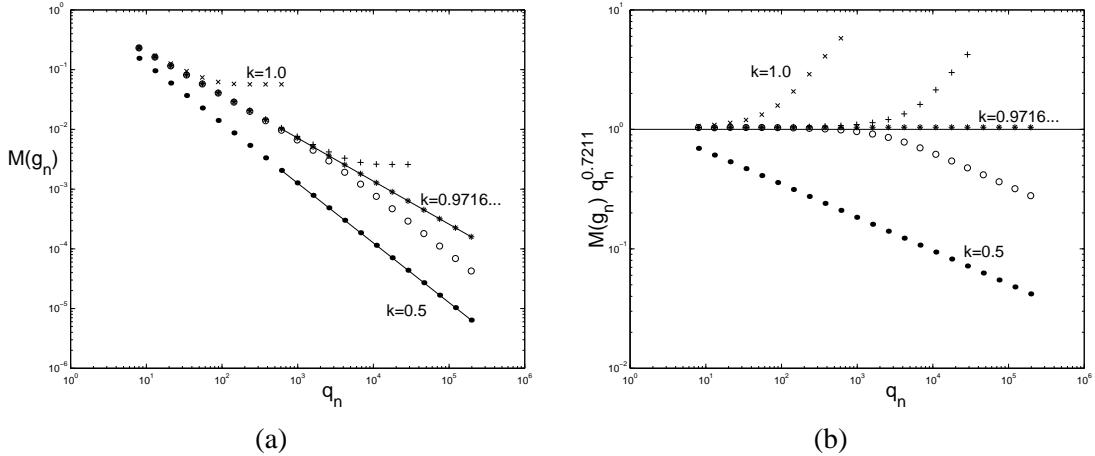


Figure 4.9: (a) The largest gap, $M(g_n)$, as a function of the length of the periodic orbit approximation, q_n , to the golden orbit. The solid circles \cdot are data for $k = 0.5$, the circles \circ for $k = 0.971$, the stars $*$ for $k = k_c = 0.971635406$, the pluses $+$ for $k = 0.972$, and the crosses \times for $k = 1.0$. The lines through the data for $k = 0.5$ and $k = k_c$ are least square fits for $610 \leq q \leq 196418$; the slopes are -1.0000 and -0.7211 respectively. (b) The same data plotted as $M(g_n)q_n^{0.7211}$ versus q_n , to exaggerate the difference between orbits converging to a circle and those converging to a Cantor set.

renormalization properties of the standard map [47]. In the final paragraphs of this section, we make a preliminary investigation of the relationship between the distribution of points and the distribution of gaps.

The largest gap

Figure 4.9(a) summarizes a study of the relationship between the largest gap, $M(g_n)$, and the period, q_n , of orbits converging to the $1/\mu^2$ orbit. There are three scaling regimes in these graphs, defined by $k < k_c$, $k = k_c$, and $k > k_c$. For $k < k_c$, $M(g_n)$ tends to zero because the orbits are converging to an invariant circle. The same is true at the critical value k_c , but the rate of convergence is different because the density of points goes to zero at every point in the orbit of $x = 0$. For $k > k_c$, $M(g_n)$ decreases and then flattens at a value that corresponds to the length of the largest gap in the Cantor set. For $k = 0.972$, this occurs when $q_n > 10^4$, and for $k = 1.0$, when $q_n > 10^2$.

We now describe the linear scaling for $k \leq k_c$ in more detail. The data for $k = 0.5$ show a clear linear relationship between $\log M(g_n)$ and $\log q_n$, which implies $M(g_n) = C/q_n^\nu$. From the least squares linear fit of the data for $q_n > 100$ we find that $\nu = 1.0000$ and $C = 1.256$. The slope, ν , equals one because the points are fairly evenly distributed around this circle (see Figure 4.12). There is also simple linear scaling at the critical value, $k_c = 0.971635406$. In this case the least squares fit for $q_n > 100$ gave $\nu = 0.7211$ and $C = 1.042$. The rate at which $M(g_n) \rightarrow 0$ is closely related to the renormalization scheme for a critical circle at $x = 0$ because the largest gap always occurs around $x = 0$. The scaling described in Section 4.3.1 implies that $M(g_n) \sim \alpha_0^{-n}$. Our results relate $M(g_n)$ and the period q_n as $M(g_n) \sim q_n^{-\nu}$. Since the periods are Fibonacci numbers, their growth is governed by the golden mean, $q_n \sim \mu^n$. Therefore, we must have $\alpha_0 = \mu^\nu$, or $\nu = \log(\alpha_0)/\log(\mu) = 0.72107$ (using the value of α_0 from [47]), which is exactly what we computed from the graph in Figure 4.9(a). When k is

just less than the critical value, e.g., $k = 0.971$, there appear to be two scaling regions. This is the result of the nonuniform distribution of points in the periodic orbit approximations to the invariant circle. For $q_n < 1000$, the orbits are close to those approximating the critical circle, so $M(g_n)$ appears to scale as $C/q_n^{0.73}$. However, for $q_n > 10^4$, the slope is very close to one; we compute $\nu = 0.998$ and $C = 8.2$.

The graph in Figure 4.9(a) suggests a new, purely geometric, criterion for determining the transition from circle to Cantor set in area-preserving maps. Given (p_n, q_n) -periodic orbits with p_n/q_n converging to an irrational, the quantity to examine is $M(g_n)q_n^\nu$ as $q_n \rightarrow \infty$. The parameter ν should be chosen to make the convergence as follows: for $k < k_c$, $M(g_n)q_n^\nu \rightarrow 0$, and for $k > k_c$, $M(g_n)q_n^\nu \rightarrow \infty$. The motivation for this is analogous to the choice of the parameter β in Greene's criterion. For the standard map, we could choose $0.7211 \leq \nu < 1$. Setting $\nu = 0.7211$ gives the best results because this balances the critical scaling exactly. This data is shown in Figure 4.9(b). Based on the data given in that graph, we estimate that the value of k_c could be determined to about five or six significant figures using this criterion.

Much more numerical work is needed to test the validity of this idea. We need to look at circles with other irrational frequencies and different area-preserving maps. The advantage of this criterion is that it does not require any information about the stability of an orbit — an expensive computation involving eigenvalues of large matrices. The new criterion is unlikely to generalize to maps in higher dimensions, however, since the topological transition is very different.

The mean and minimum gap

We can summarize some information about the distribution of gap lengths by plotting the maximum, $M(g_n)$, mean, $\langle g_n \rangle$, and minimum, $m(g_n)$, values as functions of the period, q_n on the same axes. This data is given in Figure 4.10 for the same values of k as in Figure 4.9.

In all cases, we see linear scaling in the mean gap size. This follows from the fact that

$$\langle g_n \rangle = \frac{\sum_{j=1}^{q_n} g_j}{q_n} \approx \frac{L}{q_n}. \quad (4.14)$$

The quasiperiodic orbit covers an invariant circle or a Cantor set. In either case, the orbit lies on a Lipschitz graph, so its length L , is finite. We are computing the distance between points that are close to the Lipschitz graph, so the sum of the gaps converges L . This convergence is quick, so the approximation in (4.14) is a reasonable one. If we instead compute g_j as the difference in x -coordinate, then we would have $\sum_{j=1}^{q_n} g_j = 1$ for all n and k .

We now look at the scaling of the shortest gaps. For $k = 0.5$, the relationship between $\log m(g_n)$ and $\log q_n$ is clearly linear, implying $m(g_n) = c/q_n^\eta$. From the data for $q_n > 100$, we compute $c = 0.893$ and $\eta = 1.0000$. Again, the slope is one because the points are fairly evenly distributed around the circle. For $k = k_c$ there is also linear scaling. Initially, we estimate $\eta = 1.0928$ from the data for $q_n > 100$. Since the density of points is highest at $x = 0.5$, we expect the smallest gaps to occur there; this is corroborated by our observations. It follows that we should see the period-3 oscillation described in Section 4.3.1, caused by the cycling symmetry patterns of the approximating orbits. Because the oscillations are small, it is not possible to see this in the graph of Figure 4.10. We evaluate the difference between the least-squares linear fit and the data; the graph of this in Figure 4.11 shows a very definite period-3 oscillation. We therefore compute the slope of $\log m(g_n)$ versus $\log q_n$ using every third data point and find an average value of $\eta = 1.0931$. This number is related to the renormalization scaling at $x = 0.5$ via a similar argument to that for maximum gap. Using the values from

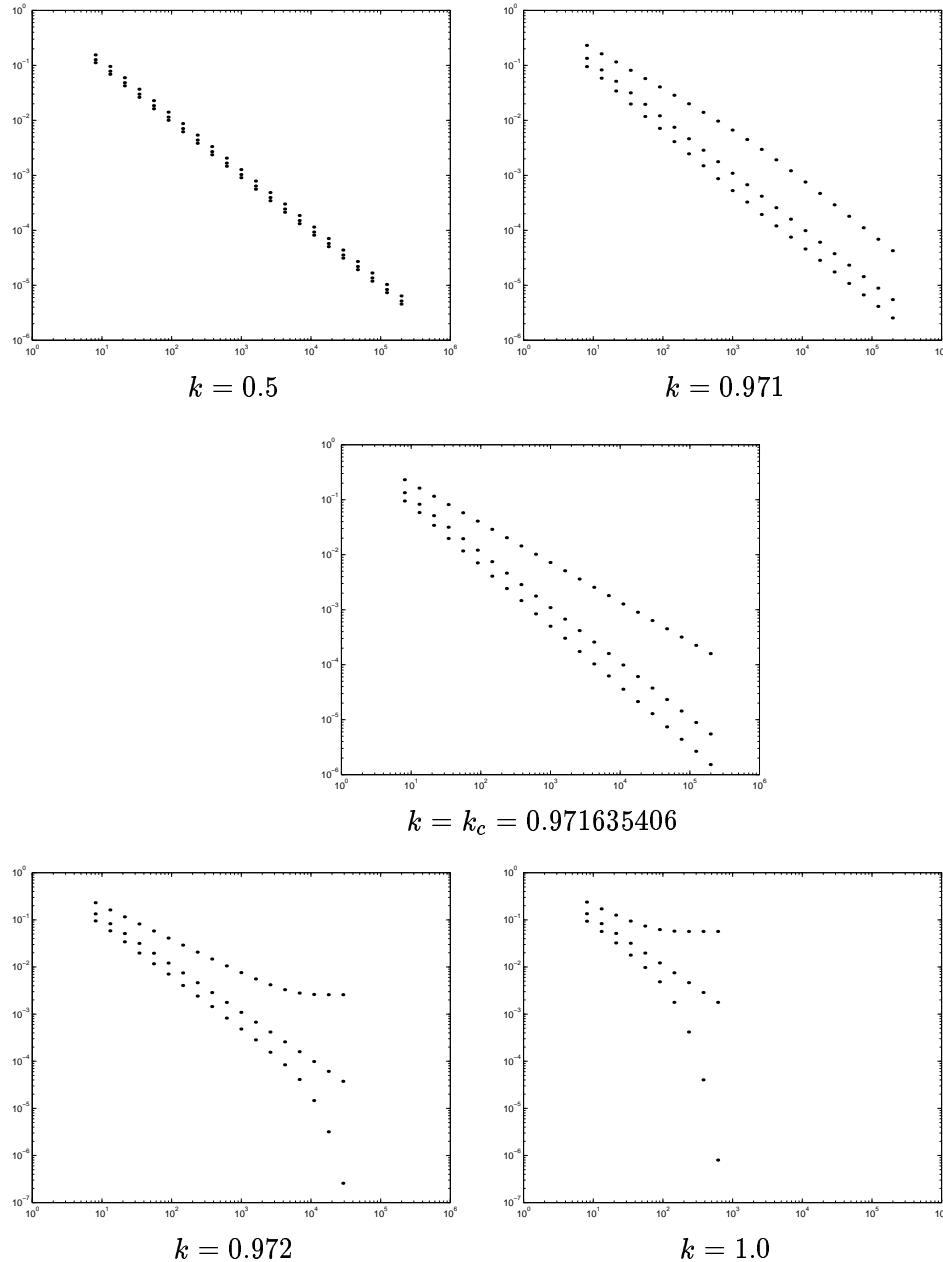


Figure 4.10: Maximum, minimum, and mean of the gap lengths for periodic orbit approximations to the golden orbit at different values of k .

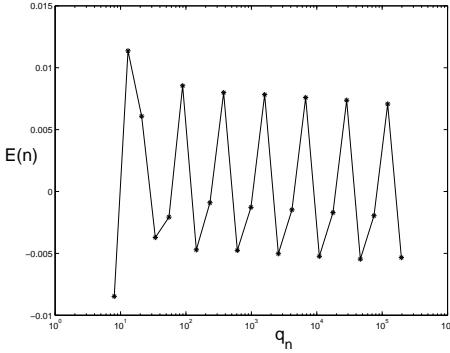


Figure 4.11: Error in the least squares linear fit to the minimum gap data for $k = k_c$ (see Figure 4.10). Specifically, we use $100 < q_n < 2 \times 10^5$ to fit the line and then compute the error as $e(n) = \log_{10}(m(g_n)) - (-1.0928 \log_{10}(q_n) - 0.0268)$. The period-3 oscillation is due to the periodic orbit symmetry properties.

(4.13), we therefore expect $\eta = \log(\sqrt[3]{\alpha_5}) / \log(\mu) = 1.09315$, which is exactly what we found. At $k = 0.971$ we see the same type of transitional behavior as before, with the data close to that of the critical circle for small q_n and slopes that converge to 1.00 as q_n gets larger.

For approximations to the Cantor sets, we see $m(g_n)$ drop off exponentially. A possible explanation for this starts with a model that gaps in the Cantor set decrease as $g_j = c\lambda^j$. A finite approximation to the Cantor set should then have smallest gap $m(g_n) \approx c\lambda^{q_n}$. In the log-log plot we therefore expect to see

$$\log m(g_n) = \log c + e^{\log q_n} \log \lambda.$$

For $k = 0.972$ we estimate $\lambda = 0.9997$ and $c = 0.0002$ and for $k = 1.0$, $\lambda = 0.983$ and $c = 0.02$. These numbers are calculated from least-square fits for data from the longest four approximations at each k , since these are orbits for which the largest gaps are resolved properly. Since such few data points are used to fit the line, these values of c and λ are not likely to be very accurate.

The linear scaling of the mean suggests it is the appropriate number to rescale by to get a limiting density distribution of gaps in the golden orbit. It is not obvious that this should exist or be well defined, but it might help give an expected form for the $C(\epsilon)$ curve which could then help distinguish between connected and disconnected sets. In the following section we pose some questions about the relationship between the density of points on an invariant circle and the distribution of gap sizes. This leads us into the problem of how to relate the density-measure for set of points to the statistics of the minimal spanning tree — a much larger project that is outside the scope of this thesis.

Renormalized gap distributions

We start by describing the distribution of points around an invariant circle; see Figure 4.12 for some example densities. We generate these graphs as histograms of x -coordinate values from very long minimizing periodic orbit approximations to the $1/\mu^2$ orbit. Sussman *et al.* [79] generate similar graphs using the linearization of the standard map and a highly accurate estimate of the tangent direction at a single point on the circle. For $k = 0$, the points of each periodic orbit are perfectly evenly spaced around the circle. This implies that the limiting

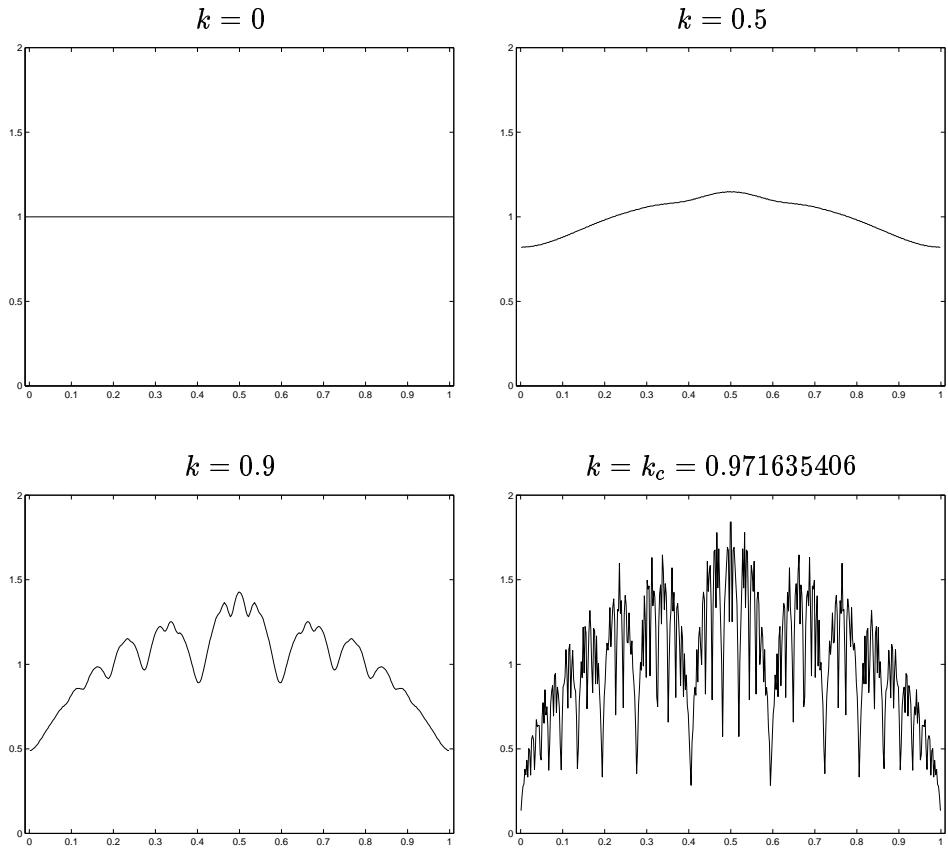


Figure 4.12: Relative density distribution of x -coordinate values for long periodic orbits ($q \sim 10^5$) approximating the golden circle at different choices of $k \leq k_c$. The number of bins in the histogram is the square-root of the number of points.

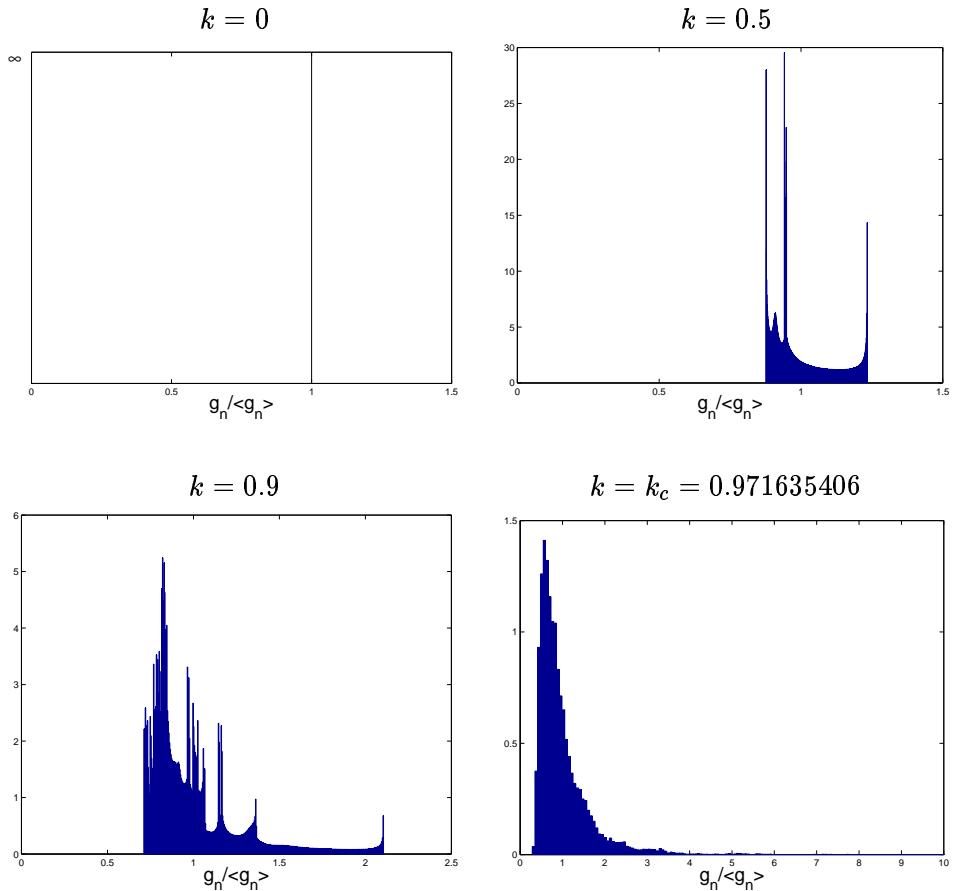


Figure 4.13: Distribution of gap lengths normalized by their mean for the same orbits and k values used in Figure 4.12. For $k = 0$, the distribution is a delta-function since the points in the orbit are evenly spaced around the circle.

distribution of points for an invariant circle is uniform. As k increases, this distribution becomes slightly less uniform with a minimum at $x = 0$ and maximum at $x = 0.5$. For k close to the critical value, the density becomes highly nonuniform, though still nonzero everywhere. At k_c the density drops to zero at $x = 0$, and at every point in the orbit of $x = 0$. The histogram approximation does not show this because of discretization effects. When $k > k_c$, the orbit approximates a Cantor set, so there are intervals of x -values where the density is zero. Since the cantori have Lebesgue measure zero, the support of the density has measure zero; it is therefore very difficult to generate a graph of this function using histograms.

The gap distributions graphed in Figure 4.13 are computed from the same orbits and k -values as Figure 4.12. Here, we have rescaled the gap lengths to make the mean equal to one in each case. Since the mean, $\langle g_n \rangle \sim L/q_n$, and $L \approx 1$, this rescaling amounts to multiplying the gaps by the number of points in the orbit. For $k = 0$, the gaps are exactly the same length and the distribution is therefore a delta function. When $k < k_c$, the gap distributions have a sharp maximum and minimum which correspond to the constants C and c computed above. A possible explanation for the spikes in the distribution is given in the next paragraph. For $k = k_c$, we know that $M(g_n) \sim C/q_n^{0.7211}$ and $m(g_n) \sim c/q_n^{1.093}$. We therefore expect the rescaled distributions to have a minimum that goes to zero and a maximum that goes to infinity as the number of points is increased. The rescaled gap distributions for the Cantor sets (i.e. $k > k_c$) should also have a minimum that goes to zero and a maximum that goes to infinity. The difference between a Cantor set and the critical circle is that the relative number of gaps less than the mean should diverge for the Cantor set.

It is not obvious how to analytically relate the density of points to the distribution of gap lengths. It seems intuitively reasonable that the longest gap will occur at the minimum in density and the shortest gap at the maximum and this is what we observe. In fact, the local rescaled gap size, $g(x)$, should be inversely proportional to the density of the points, i.e. $g(x) \sim 1/f(x)$. The relative number of gaps of a particular size, y , should then be related to the number of points where $1/f(x) = y$. This explains the spikes in the gap distributions for $k < k_c$ as the places where the density, $f(x)$, has a turning point. A formal derivation of these relationships requires some work but should be fairly straightforward. The results should be related to order statistics [10].

4.4 Cantori from sawtooth maps

The next set of examples we examine are orbits from 4d symplectic sawtooth maps. A cantorus is an invariant Cantor set of a symplectic mapping on which the dynamics is semi-conjugate to an incommensurate rotation. It is known that close enough to an “anti-integrable” limit, symplectic maps have cantori for all such rotations [48]. As we saw in the previous section, cantori in area-preserving twist maps arise from the destruction of invariant circles. For higher-dimensional maps, it is not known if the breakup of an invariant torus leads immediately to a cantorus, or if there are some intermediate stages; this question is one we intend to explore further (see Section 4.5.1). Below, we look at cantori from the anti-integrable limit.

A simple model for which cantori are analytically computable is the following piecewise linear map, called a sawtooth map:

$$\begin{aligned} p_{n+1} &= p_n + \nabla V(x_n) \\ x_{n+1} &= x_n + p_{n+1} \mod 1. \end{aligned}$$

Here $x \in \mathbb{T}^2$ (the 2-torus) and $p \in \mathbb{R}^2$ and $V = \frac{1}{2}x^t Ax$ is a quadratic potential. By choosing

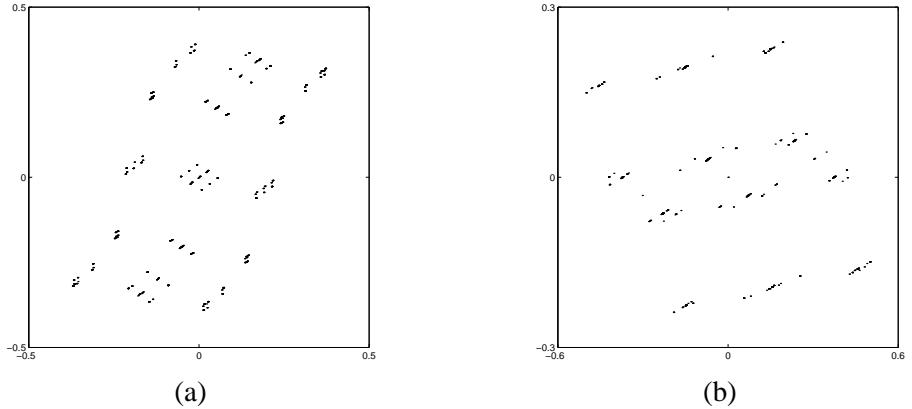


Figure 4.14: Two examples of cantori generated by symplectic sawtooth maps. Each orbit has 10^4 points.

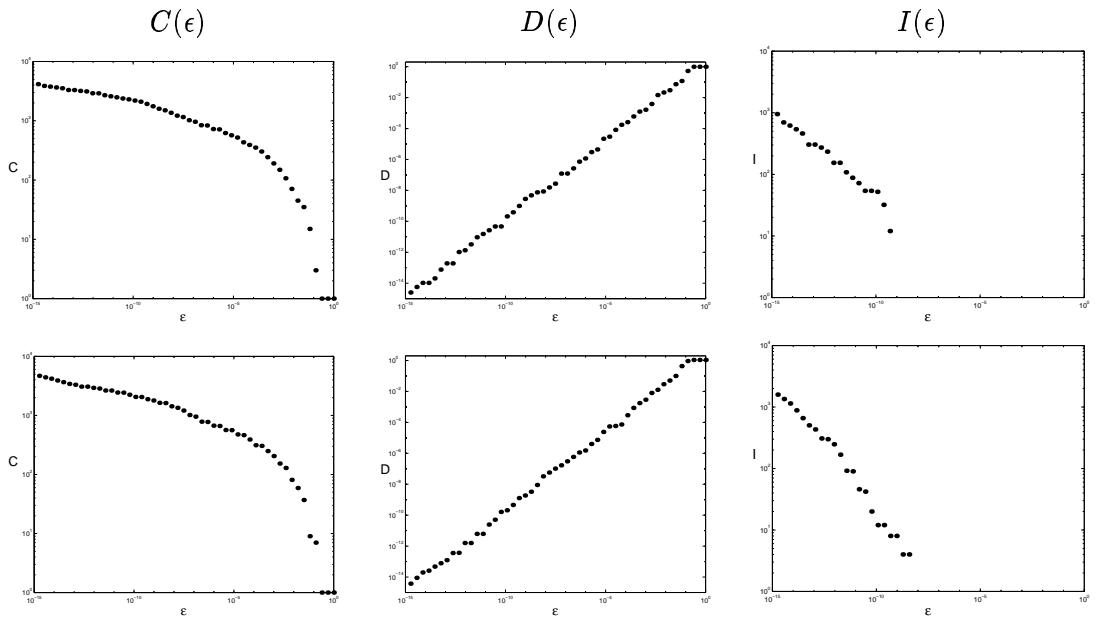


Figure 4.15: $C(\epsilon)$, $D(\epsilon)$ and $I(\epsilon)$ data for the two cantori. Top row: data for the cantorus in Figure 4.14(a). Bottom row: data for the cantorus in Figure 4.14(b). All axes are logarithmic. The horizontal axis range is $10^{-15} < \epsilon < 1$.

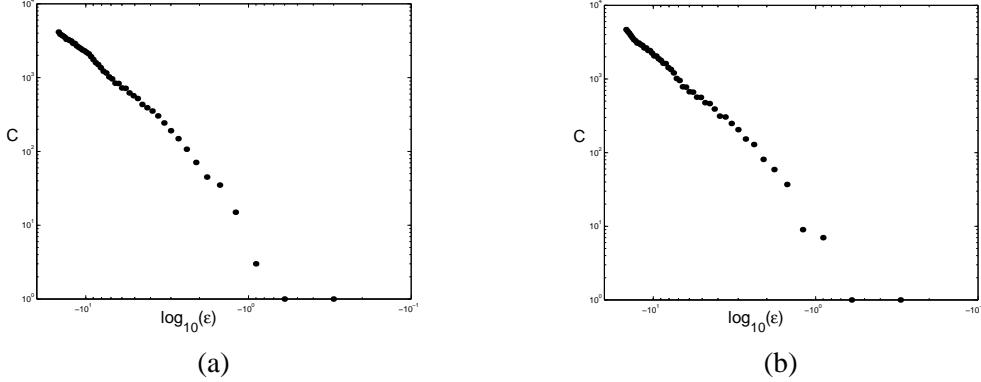


Figure 4.16: $C(\epsilon)$ versus $\log(\epsilon)$ for the two cantori of Figure 4.14(a) and Figure 4.14 (b) respectively. All axes are logarithmic. The horizontal range is $-20 < \log(\epsilon) < -0.1$

an incommensurate pair of irrational rotation numbers, for example, the golden mean $\frac{\sqrt{5}-1}{2}$ and $\frac{1}{\sqrt{2}}$, we can find an orbit analytically by looking for the semi-conjugacy, $x_t = X(\theta + \omega t)$. In Figure 4.14, we show the projection of the cantorus onto the configuration plane (x coordinates) for two choices of ω : for Figure 4.14(a) $\omega = (\frac{1}{\sqrt{2}}, \frac{\sqrt{5}-1}{2})$, and for Figure 4.14(b) ω is the spiral mean, (τ^{-1}, τ^{-2}) where τ is the real root of $\tau^3 - \tau - 1$. These orbits are generated by the same dynamical system since in both cases the matrix for the quadratic form is:

$$A \approx \begin{pmatrix} 1.9152 & -2.0358 \\ 0.5214 & 0.0847 \end{pmatrix}.$$

The graphs of $D(\epsilon)$ in Figure 4.15 are similar to those for previous Cantor sets. They tell us that the cantori are totally disconnected, because $D(\epsilon) \rightarrow 0$. We estimate δ to be very close to one: for the cantorus of Figure 4.14(a) $\delta = 0.99 \pm 0.05$; for that of Figure 4.14(b) $\delta = 1.03 \pm 0.04$. The graphs of $C(\epsilon)$ exhibit scaling behavior similar to that of the Cantor sets in the standard map (see Figure 4.6). There is no linear scaling region from which to determine γ , and this suggests that our assumption that $C(\epsilon) \sim \epsilon^{-\gamma}$ does not hold here. Since the underlying sets are Cantor sets, we must still see $C(\epsilon) \rightarrow \infty$ as $\epsilon \rightarrow 0$, but the growth rate is possibly logarithmic, rather than polynomial, in ϵ . To test this hypothesis, we plot $C(\epsilon)$ versus $\log(\epsilon)$ on logarithmic axes in Figure 4.16. The observed linear relationship tells us that $C(\epsilon) \sim (\log \epsilon)^{-\nu}$ near $\epsilon = 0$. We estimate the value of ν using a least-squares fit of the slope of the line, and find that $\nu = 2.14 \pm 0.05$ for both cases. This also implies that $\gamma = 0$. A result due to A. Fathi [25] implies that these cantori have Hausdorff dimension equal to zero, so again we see that γ and the dimension are equal for a zero-measure Cantor set.

The graphs of $C(\epsilon)$ for the two cantori are almost identical, as are their scaling rates, ν . Experimental observations show that the $C(\epsilon)$ data vary most strongly with the choice of eigenvalues for the quadratic form, rather than the eigenvectors or the frequencies. This suggests that there may be some relationship between the eigenvalues and the scaling rate, ν . A relationship between these numbers is not surprising, since the eigenvalues govern the contraction along an orbit, and contraction ratios are often associated with fractal dimensions.

4.5 Further applications

In this section we describe some possible further applications of our computational topology techniques. The list is far from exhaustive!

4.5.1 Torus breakup in the Froeschlé map

Hamiltonian flows with $d \geq 3$ degrees of freedom are modelled by symplectic maps of \mathbb{R}^{2d} — higher-dimensional analogues of area-preserving maps. One example that symmetrically couples two standard maps, is called Froeschlé map:

$$\begin{aligned} y'_1 &= y_1 + \frac{a}{2\pi} \sin(2\pi x_1) + \frac{c}{2\pi} \sin(2\pi(x_1 + x_2)) \\ y'_2 &= y_2 + \frac{b}{2\pi} \sin(2\pi x_2) + \frac{c}{2\pi} \sin(2\pi(x_1 + x_2)) \\ x'_1 &= x_1 + y'_1 \mod 1 \\ x'_2 &= x_2 + y'_2 \mod 1. \end{aligned} \tag{4.15}$$

This map was originally introduced as a model of the time evolution of elliptic galaxies [27]. When $a = b = c = 0$, orbits live on two-tori embedded in four-dimensional space. These tori are labelled by the rotation frequencies of the orbit. The KAM-theorem implies that tori with incommensurate frequencies persist for small values of a , b , and c . We saw in Section 4.4 that far from integrability, the remnants of these tori are Cantor sets, but little is known about the structure of torus remnants for moderate perturbations. It is possible that these sets have the structure of a Cantor set of lines or a Sierpinski carpet [41]. We believe that the computational topology tools from Chapters 2 and 3 could be used to help investigate this problem.

The invariant tori do not have quite the same importance dynamically as the invariant circles of area-preserving maps since a two-dimensional surface cannot enclose a four-dimensional region. This observation led Arnold to reason that a chaotic orbit of any Hamiltonian system with three or more degrees of freedom can diffuse through a dense subset of the phase space. This Arnold diffusion is typically extremely slow, and one expects that transport will be more rapid when there are few invariant tori. It is therefore still desirable to have a better understanding of the structure of torus remnants.

Recall that in Section 4.3 we approximated quasiperiodic orbits by long periodic ones. We would like to use similar techniques to approximate orbits on invariant tori in the Froeschlé map. Unfortunately, many properties of area-preserving maps do not necessarily hold in higher-dimensional symplectic maps. There is no nice generalization of continued fractions, so the first problem is to generate rational approximants to pairs of incommensurate irrational numbers. One reasonably effective approach is a generalization of the Farey tree, due to Kim and Ostlund [38]. Second, although it is easy to generalize the twist condition, there has been no generalization of Aubry-Mather theory to higher-dimensional symplectic maps. However, numerical simulations suggest that symmetric periodic orbits do converge to quasiperiodic ones. It may also be possible to use frequency analysis techniques such as those of Laskar [43] or Sussman *et al.* [79] to approximate quasiperiodic orbits.

Preliminary work with periodic orbit approximations to the Froeschlé map shows that it is easy to find long orbits close to invariant tori when the perturbation parameters, a , b , and c , are small. As these parameters increase, the orbits become unstable; it becomes impossible to find orbits with more than a few hundred points. This suggests that orbit stability is a good predictor of the existence of invariant tori [70]. It also means that we cannot get enough data

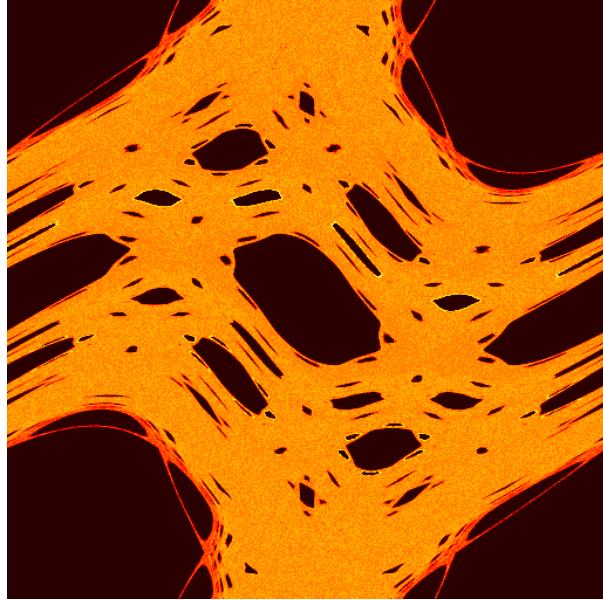


Figure 4.17: A two-dimensional histogram for a single chaotic orbit in the standard map with $k = 1.1$. We iterated the map 10^7 times, and binned the points in a 512×512 grid of boxes in the unit square $x, y \in [-0.5, 0.5]$. The color indicates the relative number of points in each box; yellow is high density, red is low density, black is zero.

to make a confident analysis of the topology — the cutoff resolution is too high. Until we can generate good approximations to quasiperiodic orbits in the Froeschlé map we will be unable to say much about their topological structure.

4.5.2 The chaotic fat fractal

As we described in Section 4.3, when $k > k_c$ all the rotational invariant circles of the standard map are destroyed and it is possible for a single chaotic orbit to cover most of phase space. The regions that remain inaccessible are the elliptic island resonance zones. This means the chaotic region contains holes at many different resolutions. Umberger and Farmer [84] show that the chaotic region is a fat fractal, and claim that the scaling they observe implies the existence of global scaling laws for the dynamics. For us, the chaotic region provides a nice example where we can test the relationship between the growth rate in the number of holes, γ_1 , and the fat fractal exponent (defined in Chapter 5).

There are a few technical issues to address here. The natural way to generate a finite set of points that cover the chaotic region is to use one or several orbits with initial conditions near an unstable point. The problem with this is that an orbit takes a very long time ($\sim 10^{10}$ iterations) to cover the chaotic region uniformly [56]. Using so many data points is impractical, but if $k \gtrsim 1$, then orbits with around 10^4 points typically cover only a fraction of the entire chaotic region. We can reduce the amount of data we store by binning — i.e., we draw an $N \times N$ grid on the phase space and record only those squares the orbit hits. This essentially digitizes the data, and means the cutoff resolution is at least $1/N$. The alpha shape approach to computational homology described in Chapter 3 is not the most efficient way find holes in gridded data. The cubical complexes proposed in Section 3.4.3 would be a more natural setting.

4.5.3 Pruning outlying data points

In [90], Zahn proposed using the minimal spanning tree to identify outlying points from experimental measurements of particle tracks in bubble chambers. We think these ideas could be extended and used to remove noisy points from embedded chaotic time series data. Regular noise filtering techniques work by deleting a band of frequencies from the signal. This is inappropriate for chaotic signals since they typically have a wide spectrum of frequencies. It is possible to compute information about the dynamics, e.g. the Lyapunov exponents, from a chaotic attractor that is reconstructed by embedding the time series data. Reducing the amount of noise in the signal is desirable, as this will improve the accuracy of such information.

Our idea is to use the minimal spanning tree and nearest neighborhood graph to identify points as “outliers at resolution ϵ .” Essentially, these will be the points isolated at resolution ϵ , so the NNG is of more interest than the MST. These ideas are very preliminary but they suggest the versatility of our techniques.

4.6 Concluding remarks

We have demonstrated the effectiveness of our computational topology algorithms by applying our techniques to a number of well-understood discrete dynamical systems. Our analysis of some sections of the Hénon attractor in Section 4.2 provides strong evidence confirming the intuition that they have Cantor-set structure. The cantori of Section 4.4 have sub-polynomial growth in the number of components —i.e., $\gamma = 0$ — even though they are Cantor sets. We resolved this problem by showing that the number of components grew logarithmically. The most interesting results of the chapter came from studying the transition from invariant circle to Cantor set in the standard map. In Section 4.3 we showed that the graphs of the number and the size of ϵ -components distinguished between approximations to invariant circles and Cantor sets at a fairly crude level. We developed a more sensitive criterion for detecting the transition by examining the largest gaps in a sequence of periodic orbit approximations. An interesting avenue for further work is to investigate the relationship between the MST edge-lengths and the underlying point distributions from orbits on invariant circles or Cantor sets. We presented numerical evidence that as the number of points increases, the gap distributions scale by their means. Since the invariant circles and Cantor sets are essentially one-dimensional objects, this problem should have parallels with the study of order statistics [10]. A better understanding of the MST edge-lengths may have implications for the renormalization of the standard map phase space.

It should be noted that our techniques use purely geometric information, no knowledge of the underlying dynamics is needed. This is an advantage in some applications where nothing is known about the dynamical system. Typically, however, there is some information that can be exploited. Thus, computational approaches, such as Mischaikow’s [59], that make explicit use of dynamical systems theory, have greater success at diagnosing the dynamics. We deliberately kept our theory and algorithms as general as possible, so that they can be adapted to different circumstances.

