



ELSEVIER

Physica D 74 (1994) 254–267

PHYSICA D

Transient measures in the standard map

James D. Meiss

Program in Applied Mathematics, University of Colorado, Boulder, CO 80309, USA

Received 23 September 1993; revised 10 February 1994; accepted 13 February 1994

Communicated by H. Flaschka

Abstract

For an area preserving map, each chaotic orbit appears numerically to densely cover a region (an *irregular component*) of nonzero area. Surprisingly, the measure approximated by a long segment of such an orbit deviates significantly from a constant on the irregular component. Most prominently, there are spikes in the density near the boundaries of the irregular component resulting from the stickiness of its bounding invariant circles. We show that this phenomena is transient, and therefore numerical ergodicity on the irregular component eventually obtains, though the times involved are extremely long $\sim 10^{10}$ iterates. A Markov model of the transport shows that the density spikes cannot be explained by the stickiness of a bounding circle of a single class – for example, a rotational circle. However, the density spikes do occur in a Markov tree model that includes the effects of islands-around-islands.

1. Ergodicity and invariant measures

As is well known [1], typical area preserving maps are neither integrable nor ergodic; instead, phase space is divided into a complicated mixture of regular and irregular components. A regular component is an elliptic periodic orbit or an invariant circle. An *irregular component* is a minimal invariant set in the complement of the regular components. It is bounded by invariant circles and is composed of hyperbolic periodic orbits and “chaotic” orbits. By definition an irregular component is ergodic; that is, it can not be split into invariant subregions with nonzero area. A reasonable conjecture is that an irregular component is the closure of the unstable manifold of a hyperbolic point [2]. Numerical evidence supports the general belief that an irregu-

lar component has nonzero area [3]; however, this remains one of the most intriguing unproven conjectures about these mappings.

In this paper we study the way in which a single “chaotic” orbit densely covers an irregular component of the standard map. If our single trajectory is not special, it must eventually uniformly cover the component since Birkhoff’s ergodic theorem implies that the time average of a function along almost any orbit is equal to its spatial average over the measure defined by the irregular component.

Numerical evidence in apparent opposition to this was presented by Smith and Spiegel [4]. In their paper the phase space density of a chaotic orbit of the standard and Hénon maps exhibited concentrations, or spikes, near the boundary of chaotic zones. These high densities conform with

the notion that the boundary of a chaotic zone appears to be sticky: orbits beginning close to the boundary will stay close for a long time. However, the orbit studied by Smith and Spiegel was not initialized close to the boundary, but in the midst of the chaotic sea. Thus the spikes appear to violate ergodicity.

We will repeat the calculations of Smith and Spiegel for the standard map. Similar computations have been performed countless times by untold thousands since the birth of the workstation. Our computations differ only in the number of iterations being large, and in the analysis that we perform. We compare our observations with expectations from simple models of transport, such as the Markov model [5].

The Markov model has been previously used to successfully predict the critical exponent for escape time through a cantorus [5] as well as the algebraic decay of correlation functions in area preserving mappings with elliptic regions [6–8]. Here we will show that it can also account for the existence of the density spikes.

2. Markov model

In general, Markovian dynamics is defined on a discrete set of states with occupation numbers $n_i(t)$ at time t . The evolution is given by

$$n_j(t + 1) = \sum_i n_i(t) P_{ij}, \quad (1)$$

where P_{ij} is the transition probability from i to j . We assume that the total occupation number $\sum_i n_i = 1$ is conserved and normalized to unity; this requires $\sum_j P_{ij} = 1$. Furthermore, we assume that the Markov chain is *irreducible* (it is possible to move from a state i to any other state, i.e. $(P^n)_{ij} \neq 0$ for some n). In this case the Perron–Frobenius theorem implies that there is a unique equilibrium $n_i = A_i$, and that almost every initial condition $n_i(t) \rightarrow A_i$ as $t \rightarrow \infty$. We call A_i the “area” of the i th state. In equilibrium, we

assume that the transition probabilities satisfy the *detailed balance* condition

$$A_i P_{ij} = A_j P_{ji} \equiv \Delta W_{ij}. \quad (2)$$

This defines the “flux” ΔW_{ij} , a symmetric matrix. Eq. (2) implies that the density $\rho_i = n_i/A_i$ obeys the transpose of (1)

$$\rho_i(t + 1) = \sum_j P_{ij} \rho_j(t). \quad (3)$$

In terms of the density, the equilibrium state is $\rho_i = 1$.

To model transport in area preserving maps by (3), we interpret the states as pieces of an irregular component separated by partial barriers – either cantori or broken separatrices [5]. Then A_i is the area of the irregular component contained between the partial barriers forming state i and ΔW_{ij} is the flux of area through the partial barrier between state i and state j (area preservation implies that ΔW_{ij} is symmetric). Thus P_{ij} is exactly the probability that an initial condition placed randomly in the i th state escapes to the j th state in one step. The major assumption of the model is that P_{ij} correctly describes the transition probability for an initial condition that is not randomly chosen but that corresponds to, say, an iterate of such an initial condition. The assumption is that when a particle enters a state, memory of its past trajectory is rapidly lost. This approximation appears reasonable if the mixing time (the Lyapunov time) in a state is short compared to the residence time in the state.

The Markov model is clearly approximate, as was reiterated by Rom–Kedar and Wiggins [9]; however, in some cases it can provide an exact description [10]. More generally, the model captures the qualitative notion of the stickiness of the boundary of a chaotic zone, even though it is not quantitatively correct. There is, however, no replacement model that can be easily applied. An exact description, such as that of [9], would require following the evolution of the lobes forming the turnstiles for all time.

In the simplest version of the Markov model we assume that each state in the Markov chain corresponds to the (p, q) resonance associated with a pair of rotational periodic orbits of frequency p/q ; the boundaries of the state are either the stable and unstable manifolds forming the resonance boundary or the minimum flux cantorus between two neighboring resonances. In this case only nearest neighbor states are connected, and P_{ij} is tridiagonal. Near a critical invariant circle with frequency ω , a natural choice for the states is the set of resonances corresponding to the continued fraction convergents of ω . We will consider the region under the golden circle with frequency $1/\gamma^2 = (3 - \sqrt{5})/2$. The sequence of approximating frequencies with $\omega < 1/\gamma^2$ is $\{0/1, 1/3, 3/8, 8/21, \dots\}$ corresponding to states $0, 1, 2, \dots$. The model becomes especially simple in the case of the renormalization fixed point map with a critical golden circle [11], when the transition probabilities scale as

$$\begin{aligned} P_{i,i+1} &= \varepsilon_0 P_{i-1,i}, & P_{i,i+1} &= \mu_0 P_{i+1,i}, \\ \varepsilon_0 &= \gamma^{-2} \approx 0.381966, & \mu_0 &= (\gamma/\xi)^2 \approx 0.139045. \end{aligned} \quad (4)$$

Here ε_0 is the time scaling, the rate at which the local frequencies approach the irrational, and μ_0 is the area scaling, the rate at which the area of the states approaches zero (ξ is the area scaling eigenvalue of the renormalization group) [6].

What is of interest here is that the Markov chain (3) with nearest neighbor connections will never give the observations of Smith and Spiegel. In fact, it is not difficult to prove the following.

Lemma. Let $\rho_j(t)$ evolve under a nearest neighbor Markov chain (3) and assume that

$$P_{j-1,j} + P_{j,j-1} \leq 1, \quad (5)$$

and that $\rho_j(0)$ is a monotone function of j . Then $\rho_j(t)$ is monotone for all t .

Condition (5) is always satisfied for our system because the transition probabilities are small. Thus if one begins with all particles in state 0, corresponding to being far from the invariant circle, then ρ_j is nonincreasing as a function of j for all time and there can be no density spikes.

We observe that the lemma is true in a simple Monte Carlo numerical simulation of (3). The density relaxes to a constant, but is always monotone.

To help resolve the conflict between the model and the observations of Smith and Spiegel, we turn to numerical experiment.

3. Numerical experiments

The simplest numerical attempt to test ergodicity is to partition phase space into N^2 boxes with sides of size $1/N$, and record the number of times a trajectory visits each box during t iterations. To study a particular irregular component, we start our initial condition close to a hyperbolic periodic point; thus it approximately follows the unstable manifold of this point.

As a model, we use the standard mapping

$$\begin{aligned} y' &= y - \frac{k}{2\pi} \sin(2\pi x) \bmod 1, \\ x' &= x + y' \bmod 1. \end{aligned}$$

We study this system near $k = k_{cr} \approx 0.971635406$ where the last rotational invariant circle, the golden circle with frequency γ^{-2} , is destroyed (for a review see [12]).

Numerical iterations using floating point arithmetic suffer from the rapid loss of accuracy due to the chaotic instability, and are also not one-to-one. In fact as was noted in [13], each floating point trajectory is eventually periodic. Even worse, these trajectories can drift across what ought to be invariant circles, becoming trapped on low period orbits. One can attempt to circumvent these problems by discretizing the system on an integer lattice of L^2 points, letting $x = i/L$, and $y = j/L$, giving a one-to-one map that can be

computed with no loss of accuracy [14]. Every orbit in a discrete mapping is necessarily periodic. Rannou has shown that the average period is $\mathcal{O}(L)$ for a random permutation (with the symmetries of the standard map), and numerical computations for the standard map on lattices with $L \leq 800$ confirmed this dependence. However, Percival and Vivaldi noted that the distribution of orbit periods will vary depending upon number theoretic properties of L [15,16]. For example choosing L to be a highly composite number such as the ever popular 2^n is expected to give anomalous results. We will compare several different L in our computations.

If one naively iterates using floating point arithmetic, as we do for the most part, it is important to note at least one result from above: the effective number of lattice points is $\mathcal{O}(\text{precision}^{-2})$, and if one iterates the map longer than $\mathcal{O}(\text{precision}^{-1})$ a typical orbit would be periodic. This has observable effects, for example in single precision calculations of the diffusion coefficient. We use IEEE double precision arithmetic, and therefore require $t < 10^{16}$ – a limit that is well beyond our computational resources.

Sometimes one can appeal to *shadowing* to evade the criticism of the loss of accuracy due to chaotic instability, i.e. even if one is not iterating a true orbit of the map, shadowing would imply there is a true orbit nearby [17,18]. However, we believe that the very phenomena that make our system interesting – long time correlations near the boundaries of chaotic zones – invalidates employment of such methods; this is similar to the *glitch* discussed in [19].

3.1. Chaotic area

Fig. 1 shows the area ($A \equiv$ the number of visited cells/ N^2) visited by a single chaotic trajectory as a function of time for $N = 500$ using floating point arithmetic. The solid curve corresponds to a trajectory started at $(x,y) =$

$(-0.5, 0.0)$ which is within 10^{-16} of the unstable fixed point or $(0,1)$ orbit. For $k = k_{cr}$ the visited area converges as $t^{-0.45}$ to its asymptotic value of $A_c(k_{cr}) = 0.408$. This corresponds to the trajectory filling the chaotic zone that is bounded by the critical invariant circles with rotation numbers $\pm\gamma^{-2}$ – we expect the covered set to be a “fat fractal,” a fractal with nonzero area, following [3,20]. The dashed curve shows an orbit started at a point on the hyperbolic $(377,910)$ orbit slightly below the critical invariant circle (i.e. state 7 of the Markov chain). The numerical trajectory remains trapped in the neighborhood of the periodic orbit, thus visiting only a small region of phase space until $t \approx 5 \times 10^8$. At this time it escapes and rapidly spreads over the same domain as the first orbit.

For $k > k_{cr}$, the accessible area suddenly jumps because the golden circles become cantori. For example at $k = 0.981$ the accessible area is $A_c(0.981) = 0.534$; however, it takes an extremely long time for an orbit to escape across the golden cantorus. In the figure this escape occurs at $t = 7 \times 10^7$. Using transport theory [5], the average transit time for trajectories for initial conditions in the irregular component between the golden cantori with frequencies $\omega = \pm\gamma^{-2}$ is exactly

$$t_{ave} = \frac{A_c}{2 \Delta W_\gamma} = \frac{0.408}{2 \times 1.66(10)^{-8}} = 1.2(10)^7, \quad (6)$$

where ΔW_γ is the flux through each of the golden cantori. This agrees reasonably well with the observed time for this single trial, which after all was placed at the “maximum” distance from the cantorus.

3.2. Density distribution

Though the chaotic trajectory appears to densely cover a nonzero area, the coverage is not uniform for finite t . This was discussed by Smith and Spiegel, who observed that the density distribution on phase space exhibits large fluctua-

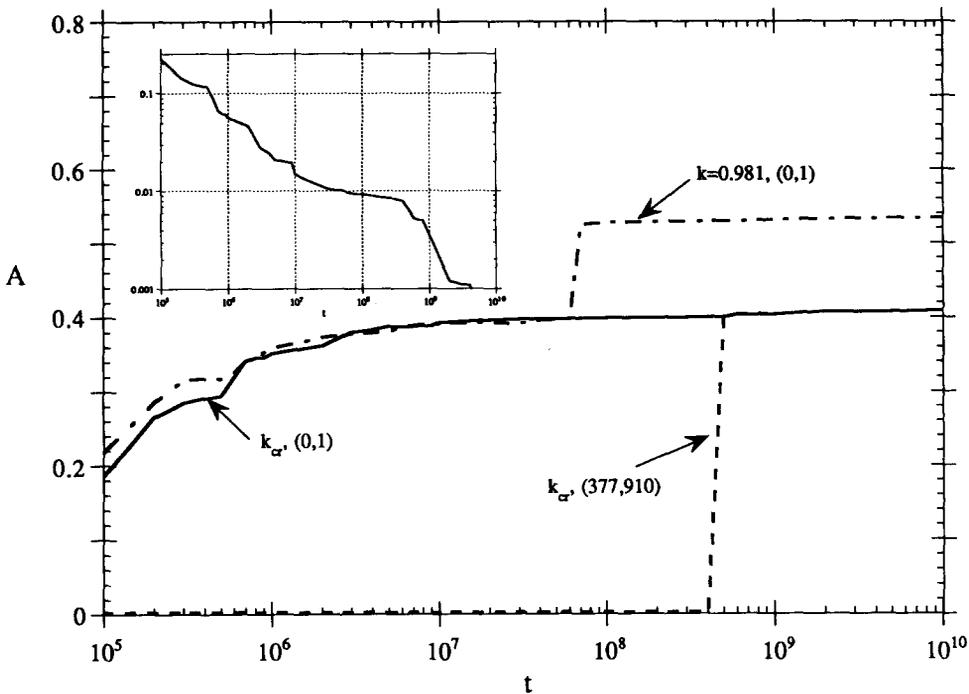


Fig. 1. Occupied area vs. time. Area is measured by the number of occupied cells in a 500×500 cell grid. Initial conditions are labeled by the frequency of the nearby hyperbolic orbit. For the dot-dashed curve, the orbit remains trapped between the golden cantori until $t \approx 7 \times 10^7$. For the dotted curve, the trajectory is trapped in a narrow region around the $(377,910)$ orbit until $t \approx 5 \times 10^8$. The inset shows $A_c - A \sim t^{-0.45}$ for the $k = k_{cr}$ trajectory beginning at $(0,1)$.

tions [4]. Here we analyze this phenomena in more detail.

The asymptotic, average occupation number is expected to be t divided by the number of accessible cells $N^2 A_c$, so we define the normalized density

$$\rho = k \frac{N^2 A_c}{t} \tag{7}$$

for occupation number k . Fig. 2 is a color density plot of the trajectory beginning at the hyperbolic fixed point after 6×10^9 iterates. The majority of cells, colored red and blue, have $1 < \rho < 1.17$. A few cells have small density (19.9% of the occupied cells have $\rho < 1$). As indicated by the yellow cells in the figure, these tend to be on the boundary of islands – they represent boxes that intersect the boundary of the chaotic zone and are therefore partially filled naturally. In support of this interpretation, the percentage of cells with $\rho < 1$ stays constant at $20 \pm 0.5\%$ for $t > 2 \times$

10^8 . There are also a small number of cells that have a density far in excess of the mean – the maximum density at $t = 6 \times 10^9$ is $\rho = 5.23$. The 0.57% of the cells that have density larger than 1.17 are white in the figure; these are also to be found on the edges of several island chains, as shown in Fig. 3. We postpone discussion of these for the moment.

A plot of the probability distribution of the density is shown in Fig. 4. For $t = 10^6$ about 14% of A_c is not yet explored and the probability distribution is very broad. The mode of this distribution decreases with time, tending asymptotically to $\rho = 1.13 \pm 0.005$. The decrease is caused by the fraction of occupied cells growing with time towards A_c . The mode is not unity as would be expected because of the low density tail of partially occupied cells. The mean density is unity once all A_c cells are occupied (this is due to our normalization).

For intermediate times the density distribution

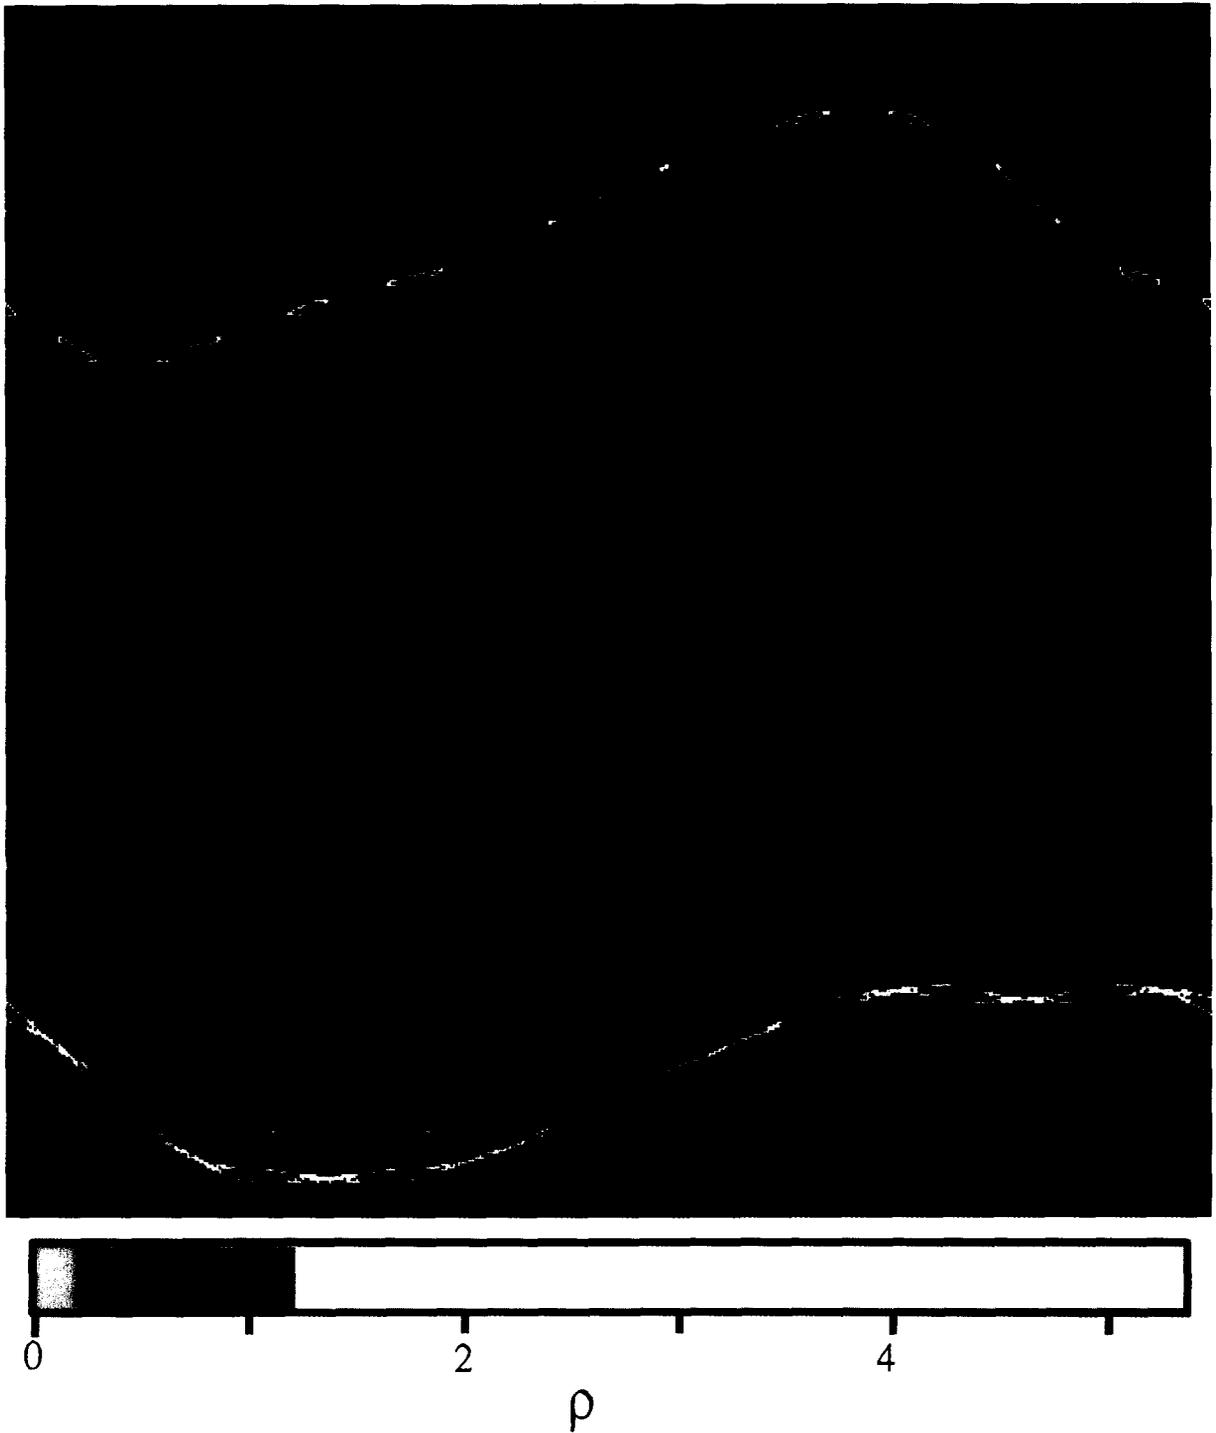


Fig. 2. Color Density plot of a single trajectory at $k = k_c$, for $t = 6 \times 10^9$, $N = 500$. The color-density scale is shown at the bottom.

is narrow and Gaussian-like; however it has a secondary peak at lower density (e.g. at $t = 10^8$, this peak is at $\rho = 0.84$). This secondary peak slowly moves upwards; it appears as a shoulder at $\rho = 1.05$ for $t = 2.5 \times 10^8$ (see also Fig. 6

below) and by 10^9 it has been virtually absorbed. The low density cells corresponding to this peak are mostly in the (1,3) resonance zone: they are the red cells in Fig. 2. The existence of this secondary peak appears to depend sensitively on

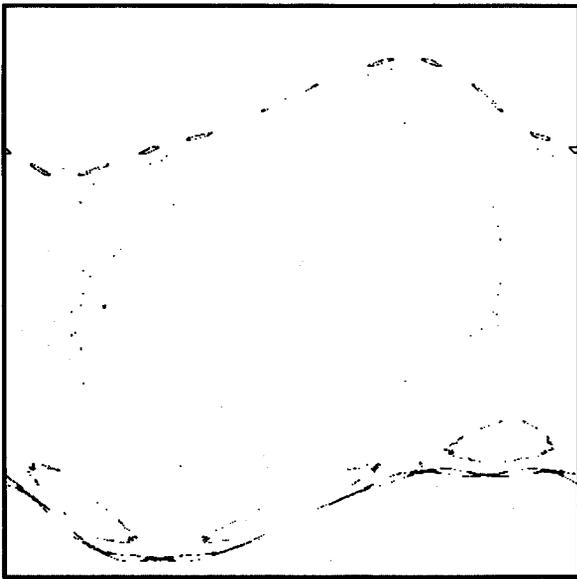


Fig. 3. The cells from Fig. 2 that have $\rho > 1.17$.

the choice of initial condition (see discussion below).

If the density distribution were obtained from a random process (randomly put t balls in $A_c N^2$ urns) then the distribution of the expected number of cells with a given occupation number would be binominal. In the limit of large time this would become the Gaussian about $\rho = 1$ with the variance

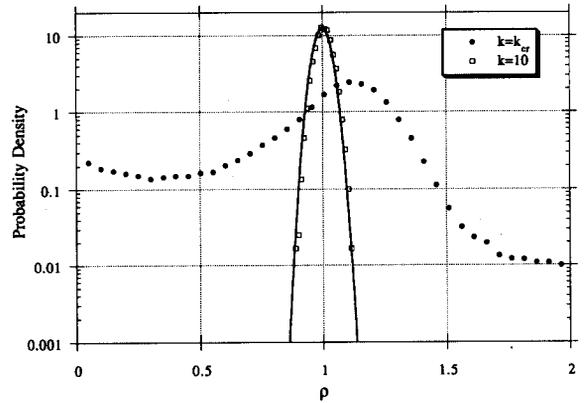


Fig. 5. Comparison of Gaussian with computations for $t = 10^7$. Shown are data for $k = 10, N = 100$ and $k = k_{cr}, N = 500$.

$$\sigma^2 = (A_c N^2 - 1)/t. \tag{8}$$

Indeed, computations at large values of k (e.g. $k = 10$, see Fig. 5) where the entire phase space appears chaotic, give results for the density distribution virtually indistinguishable from the Gaussian. As Figs. 5 and 6 show, the agreement is not as good for $k = k_{cr}$. In the first place, the mode of the Gaussian is 1.00, while the numerical results give a larger value. This is due to the 20% of the cells that are cut by the boundary of the irregular component. In Fig. 6 we shift the

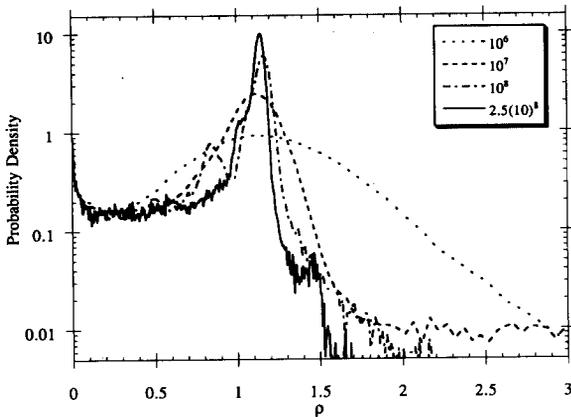


Fig. 4. Histogram of the probability distribution of the density at $k = k_{cr}$.

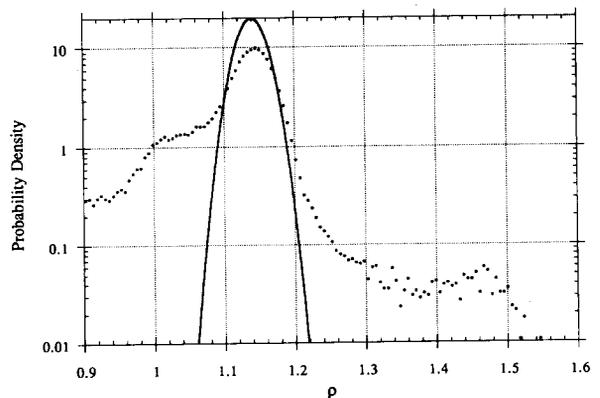


Fig. 6. Comparison of computations with Gaussian for $t = 2.5 \times 10^8, k = k_{cr}, N = 500$. The solid curve is a Gaussian artificially centered at $\rho = 1.14$, but with variance given by (8).

mode of the Gaussian, to show that its variance agrees well with the numerical distribution. For other times the numerical σ also agrees with the Gaussian; for example, from the data at $t = 10^9$ we compute $\sigma = 0.01 \pm 0.003$ in precise agreement with (8). The low density shoulder on the numerical distribution is quite visible in Fig. 6; it is absorbed into the main peak for $t > 10^9$. Our conclusion is that for $t > 10^9$ the computations are consistent with ergodicity on the irregular component – with the exception of the $\frac{1}{2}\%$ of high density cells. We discuss this in the next section.

Before leaving this section, we pause to consider how our results depend upon the choice of initial condition and the use of floating point arithmetic. We will find that the transient features of the density distribution do vary, however, by $t = 10^9$ most of these transients are gone.

In the previous computations, we used an initial condition as close as possible to the hyperbolic fixed point, letting $(x, y) = (-0.5, 0.0)$ – this point is not fixed numerically since $\sin(\pi) \approx 1.22 \times 10^{-16}$ to double precision accuracy. A prominent feature of the distribution for this initial condition was the low density shoulder caused by sparse occupation of the (1,3) resonance. This feature is no longer present in a computation for the nearby initial point $(-0.5 + 10^{-12}, 0.0)$; this is shown as the solid curve in Fig. 7. The distribution now takes the Gaussian form much earlier. Never-the-less, by $t = 10^9$ the low density shoulder in the first computation is absorbed and the two give comparable results.

To compare the double precision calculations with those of an integer mapping we set $x = i/L - 0.5$, $y = j/L$, and iterate the bijective map.

$$f_L: \begin{cases} j' = j + S(i) \\ i' = i + j' \end{cases}, \quad S(i) = \left\lfloor \frac{kL}{2\pi} \cos\left(\frac{2\pi i}{L}\right) \right\rfloor,$$

where $[x]$ is the largest integer less than x (i.e. floor(x)). Using the initial condition $(i, j) = (0, 1)$ we obtain results, shown in Fig. 7, that exhibit

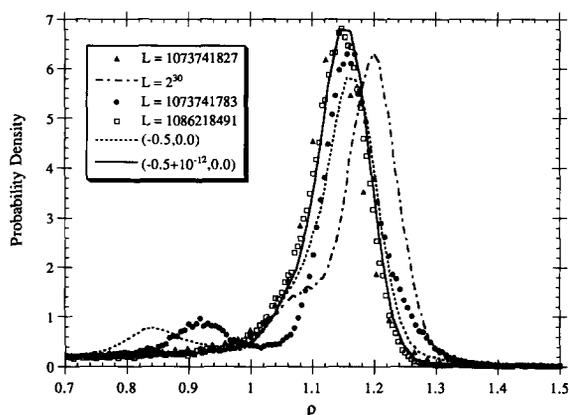


Fig. 7. Comparison of density distributions for $k = k_{cr}$, $N = 500$ and $t = 10^8$. The solid line is a double precision calculation for the initial condition $(-0.5 + 10^{-12}, 0.0)$. Integer calculations for the initial condition $(i, j) = (0, 1)$ with various prime lattices are shown with the symbols noted. The two values of $L > 2^{30} = 1073741824$ give almost identical results to the double precision calculation. However a double precision calculation for $(0.5, 0.0)$ (dotted) and ones using prime $L < 2^{30}$, as well as for $L = 2^{30}$ show low density, transient features.

some sensitivity to the integer L . For our computations changing L caused variations in the result that were similar to changing the initial condition in the floating point calculation; indeed, since we fixed $j = 1$ we were in fact changing the initial condition as we varied L . We were unable to determine if the bijective nature of the integer map made it superior to the floating point map in any way.

One issue of interest is the possible dependence of the results on whether L is composite or prime [15]. We compared the computation for $L = 2^{30}$ with several nearby prime numbers. Interestingly, for those prime $L > 2^{30}$ that we checked, the integer calculations are nearly Gaussian just like the solid curve in Fig. 7. However for $L = 2^{30}$ or a prime slightly smaller than this, the calculations differed significantly – often exhibiting a low density peak showing that some region of phase space was under-occupied. Some examples are displayed in Fig. 7.

In all cases, however, these differences appear to be transient. The probability distribution

becomes increasingly narrow and Gaussian-like as t increases, taking into account the flat tail of low density cells caused by the boundary of the irregular component.

3.3. Density spikes

We have seen that most of the anomalous features of the density distribution appear to be transient, implying numerical ergodicity on the irregular component. However, the existence of a small number of cells with excess density is still troubling. As shown in Fig. 2, at $t = 6 \times 10^9$ about 0.57% of the cells have density larger than 1.17. According to (8), the statistical standard deviation at this time is 0.004. Thus, since 1.17 is 10 standard deviations above 1.13, only $\sim 10^{-23}$ of the cells should have $\rho > 1.17$. Indeed, such an observation led Smith and Spiegel to entitle their paper “Strange Accumulators” – as if the islands could act as attracting sets. Of course this cannot happen for a Hamiltonian system (as they realized), so we must look elsewhere for an explanation. One might think that a simple Markov model of the process could show strong peaks in the density caused by the fact the states approaching an invariant circle have escape probabilities that tend to zero, and thus a particle that lands in such a state will be trapped for times that approach infinity. However, this must be balanced by the fact that it is increasingly difficult to enter these states. As we showed in section 2, a nearest neighbor Markov model will *never* show these spikes.

In the calculations, the maximum density and the location of the over-dense cells in phase space depend sensitively on the choice of initial condition and, if we discretize, on L . In all cases however, the observed maximum density declines toward the mean as t increases, see Fig. 8. This decay is punctuated by sudden increases when the orbit is trapped for long segments. Following these trapping episodes, ρ_{\max} decays precisely as t^{-1} because the maximum occupation number is constant. The overall decay is

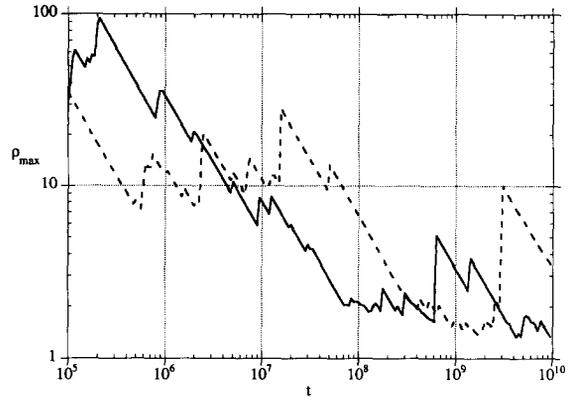


Fig. 8. Maximum density $k = k_{cr}$, $N = 500$. The solid curve is for the initial condition $(0.5 + 10^{-12}, 0)$, and the dashed curve for $(0.5, 0, 0)$. The straight segments decay exactly as t^{-1} . The spikes are due to trapping episodes when the maximum occupation number jumps up.

roughly $t^{-0.2}$, though it is difficult to extract a rate from the data. The computations of Smith and Spiegel were for relatively short times, $t \approx 10^8$, and they observed $\rho_{\max} \sim 20$; this is consistent with our observations. In any case, the spikes are a transient phenomena, and at least we can conclude that there are no “accumulators” in our system.

3.4. Reduction to linear chain

To further elucidate the occurrence of the density spikes we compare the numerical experiments with a one-dimensional Markov model. Thus it is necessary to determine the state containing each phase point $z_i = (x_i, y_i)$. The correct way to do this would be to choose a set of cantori, discretize the phase space into states separated by these cantori, and use these for comparison. This would be computationally intensive, however. Instead of this procedure we use the width function introduced in [21]. This has the advantage that the algorithm for checking whether a point is “in” a particular state is trivial; the disadvantage is that we are an additional step removed from a direct comparison between the Markov model and the map.

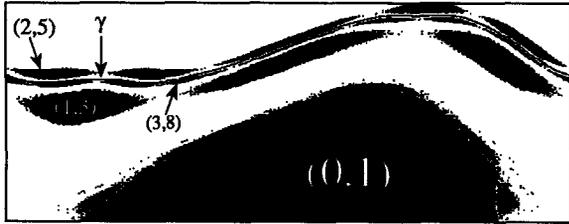


Fig. 9. Points that have widths less than one for 500 iterates for several values of ω and $k = k_{cr}$. Shown are the (0,1), (1,3), (2,5), (3,8) and γ states.

The *width* of a point relative to a given frequency ω is defined as

$$w_\omega[z_0] = \sup_k (x_k - \omega k) - \inf_j (x_j - \omega j), \quad (9)$$

thus it measures the deviation of the orbit from rigid rotation with frequency ω . All points on “ordered” or “monotone” orbits have $w_\omega \leq 1$, relative to their rotation frequency, and all rotational invariant circles and cantori are monotone [12]. Thus it is natural to consider the set of points whose orbits have width less than one:

$$W_\omega = \{z: w_\omega[z] \leq 1\}. \quad (10)$$

With several reasonable, geometric assumptions, we can prove that $W_{p/q}$ is contained in the resonance of frequency p/q [21]. We show several of the sets $W_{p/q}$ in Fig. 9. Thus the width distinguishes between the states in the Markov chain. Furthermore, w_ω is easily computed from the orbit: no complicated geometrical intersections need be computed.

For the standard map, we use the reflection

symmetry about the origin to reduce our phase space to the orbifold $\{-0.5 \leq x < 0.5, 0 \leq y \leq 0.5\}$. For $k = k_{cr}$, we choose the states to be the continued fraction convergents of the golden mean; when $\omega < 1/\gamma^2$, the set of states is $\{0/1, 1/3, 3/8, 8/21, 21/55, 55/144, \dots\}$ for $i = 0, 1, 2, \dots$. Note, however that the union of these states does not cover the phase space below the critical circle: one would need to include all rationals to do this [22].

For our computations we begin with a point z_0 close to the hyperbolic point of the p/q orbit and assume that it is in fact in this state. We compute w_ω along this orbit until it first exceeds one, declaring that it has now made a transition to a neighboring state. If the width exceeds unity because the supremum has increased we declare the new state to have larger frequency. Conversely if the infimum has decreased the new state has a lower frequency. This gives a set of nearest neighbor transitions. The occupation number T_i is simply the number of iterates in state i and the transition probabilities are the fraction of the T_i that lead to a transition on the next step.

In Table 1, we show the results of this calculation for $t_{tot} = 10^{10}$ and an orbit started at the (0,1) hyperbolic point, state 0. Remarkably this orbit never moves above state 4! On the other hand, if we initialize the orbit in a higher state, e.g. $s_0 = 4$, then it remains trapped there for a long initial segment, as we saw before. However, by $t = 10^{10}$, the density distribution and the

Table 1

Occupation probabilities and transition probabilities using the width calculation for $t = 10^{10}$ using floating point arithmetic, compared with theoretical results based on the resonance areas and turnstiles

State	$s_0 = 0$			$s_0 = 4$			Theoretical		
	T_i/T_{tot}	P_{i-1}	P_{i+1}	T_j/T_{tot}	P_{j-1}	P_{j+1}	A_i/A_{tot}	$\Delta W_{i-1}/A_j$	$\Delta W_{i+1}/A_i$
0	4.17×10^{-1}		8.35×10^{-2}	4.16×10^{-1}		8.33×10^{-2}	8.02×10^{-1}		2.85×10^{-3}
1	5.11×10^{-1}	6.81×10^{-2}	1.27×10^{-3}	5.11×10^{-1}	6.77×10^{-2}	1.25×10^{-3}	1.71×10^{-1}	8.50×10^{-4}	5.47×10^{-4}
2	6.40×10^{-2}	1.01×10^{-2}	1.82×10^{-4}	6.15×10^{-2}	1.04×10^{-2}	1.90×10^{-4}	2.29×10^{-2}	2.84×10^{-4}	1.97×10^{-4}
3	7.39×10^{-3}	1.58×10^{-3}	2.12×10^{-5}	7.93×10^{-3}	1.47×10^{-3}	4.24×10^{-5}	3.16×10^{-3}	9.96×10^{-5}	7.32×10^{-5}
4	3.92×10^{-4}	4.00×10^{-4}		2.89×10^{-3}	1.17×10^{-4}	7.75×10^{-6}	4.39×10^{-4}	3.80×10^{-5}	2.79×10^{-5}
5				4.10×10^{-4}	5.46×10^{-5}		6.10×10^{-5}	1.45×10^{-5}	1.07×10^{-5}

transition probabilities for this calculation agree well with those for the previous one, except for the excessive population in states 4 and 5.

If the simple Markov chain model were valid, the transition probability would be given by Eq. (2), where ΔW_{ii+1} (ΔW_{ii-1}) is the area of the turnstile in the upper (lower) separatrix of the i th resonance, and A_i is the area of the connected chaotic component of the resonance. We can compute ΔW and the total resonance area (including the regular regions) using the action function [22]; the results are given in Table 2. We use these values in the last three columns of Table 1 to obtain the expected equilibrium occupation probability A_i/A_{tot} , as well as the expected transition probabilities. While there is certainly not quantitative agreement, there is order of magnitude agreement – the computations tend to be within a factor of 3 of the predictions. With the exception of state 0, the occupation probabilities of the states tend to be larger, and the downward probabilities are always larger, than expected. We can estimate the state density by $\rho_i = T_i/T_{\text{tot}} \times A_{\text{tot}}/A_i$. The values in Table 1 give

$$\rho_i = \{0.52, 2.99, 2.79, 2.34, 0.89\}$$

for the $s_0 = 0$ calculation. Thus there are no density spikes: the density decreases with state with the exception of state 0. For $s_0 = 4$, on the other hand $\rho_4 = 6.58$ and $\rho_5 = 6.72$, but these excesses are due to the initialization of the orbit in this state 4. We conclude that the observed density spikes would not occur if there were only

Table 2

Areas of resonances and their turnstiles for k_{cr} . The italicized entries are not calculated, but are estimated according to the renormalization scaling, Eq. (4).

ω	Area	ΔW_{ii+1}	ΔW_{ii-1}
(0,1)	1.9654×10^{-1}	5.5950×10^{-4}	
(1,3)	4.1981×10^{-2}	2.2945×10^{-5}	3.5690×10^{-5}
(3,8)	5.6110×10^{-3}	1.1059×10^{-6}	1.5913×10^{-6}
(8,21)	7.7411×10^{-4}	5.6646×10^{-8}	7.7131×10^{-8}
(21,55)	1.0764×10^{-4}	3.0085×10^{-9}	4.0964×10^{-9}
(55,144)	1.4901×10^{-5}	1.5978×10^{-10}	2.1756×10^{-10}

a single invariant circle bounding the irregular component.

4. Markov tree

As observed in the numerical experiments, the highest densities occur near the boundary of small elliptic islands; for example in Fig. 2 these correspond to the island chains of the (3,10) and the (-3,8) orbits. The density spikes occur in a small region near the outermost invariant circles of these islands; in fact, probably around small islands which are themselves encircling these islands.

A similar phenomena was observed by Karney [23] in numerical experiments on the first escape time distribution for an orbit that began near an island. He found that the longest trapped segments corresponded to the orbit being trapped around an island-around-an-island-around-an. . . .

A Markov model can be constructed which takes islands-around-islands into account [7]. Each state still corresponds to a resonance zone whose frequency is a convergent of a boundary circle. These we call the *levels*. Each boundary circle gives a new *class* of states [24]. For example the rotational boundary circle with $\omega = \gamma^{-2}$ corresponds to class zero, and its convergents give the levels that we treated before – the Markov chain. Within the (p,q) state corresponding to one of these levels is the (p,q) rotational elliptic periodic orbit. It is surrounded by circles invariant under the q th iterate of the map – these are class 1. The outermost of these forms a boundary circle, and its convergents provide a new infinite set of levels at class 1. Within each of these levels there can be class 2 islands chains, and so forth.

The resulting set of states has the topology of a tree. Following [7], we denote a state on the tree by a symbol sequence $S = \{abc...wx\}$ where the letters denote the sequence of branches leading to S from the root state R . Each state

has several progeny, say $S\alpha = \{abc...wx\alpha\}$, etc. but only a single parent state, $DS = \{abc...w\}$. There are transition probabilities $P_{S,S\alpha}$ to each daughter and $P_{S,DS}$ to the parent. An example of this structure for a single new class at each level is shown in Fig. 10. A left transition on the tree is designated with “0”, and a right transition with a “1”. Left transitions correspond to approaching a particular boundary circle, while right transitions correspond to becoming trapped around a higher class island. In general, the number of branches coming out of any junction corresponds to the number of distinct elliptic orbits within each level. All but a small number of the islands will be too small to be of any practical importance on realizable time scales.

What is surprising about the Markov dynamics on the tree is that the density does not necessarily remain monotone (i.e. with the density of the parent larger than that of each of its progeny). Fig. 11 displays the result of a Monte Carlo calculation for which the particle is initially in the root state and is iterated 10^6 times. We used the self-similar transition rates [7], analogous to (4):

$$\begin{aligned} P_{S,DS} &= p_0 \epsilon_0^{\lambda[S]} \epsilon_1^{\rho[S]}, \\ P_{S,S0} &= \mu_0 P_{S0,S}, \\ P_{S,S1} &= \mu_1 P_{S1,S}, \end{aligned} \tag{11}$$

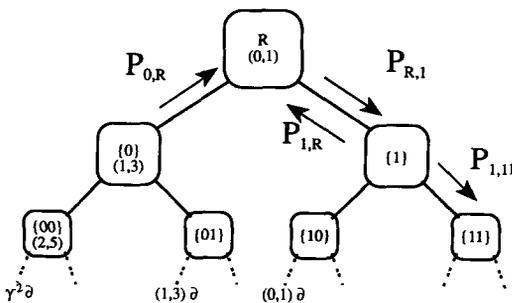


Fig. 10. Sketch of the states and transitions for a two branch Markov tree. The leftmost path corresponds to the sequence of levels in the Markov chain model. Each transition to the right corresponds to becoming trapped around an island within a level. If the root state corresponds to the (0,1) resonance, then the state {1000...} approaches the boundary of the island around the (0,1) elliptic point, and {01000...} approaches the boundary of the (1,3) island.

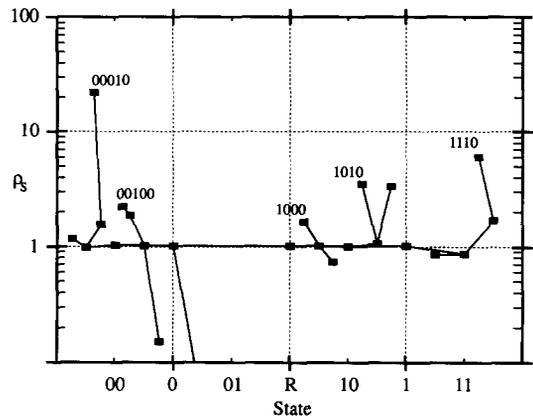


Fig. 11. State density for the Markov tree model at $t = 10^6$ with parameters from (4) and (12). The abscissa is given by the symbol sequence for the state interpreted as a binary decimal with an appended “1”. States not shown were unoccupied.

where $\lambda[S]$ is the number of 0’s in S , and $\rho[S]$ is the number of 1’s. Thus there are two time scaling constants, ϵ_0 and ϵ_1 , and two area scaling constants, μ_0 and μ_1 . We require $\mu_0 + \mu_1 < 1$ in order that the total area be finite. Reasonable values for the new parameters are [7]

$$\epsilon_1 = \frac{1}{7}, \quad \mu_1 = 0.099105. \tag{12}$$

The ordinate in Fig. 11 is the density ρ_S in each state and the abscissa represents the tree structure as in Fig. 10. If the density were monotone each of the branches in Fig. 11 would hang down, yielding a tree of the willow variety. Nonmonotonicity occurs at a number of states. As t is increased, the density in most of the states tends towards unity; however depending upon the seed in the random number generator a few of the states acquire a large density, which then slowly relaxes as t^{-1} similar to the standard map data shown in Fig. 8.

One might object that Fig. 11 was obtained for a single realization, and that statistics might wash out the effect. However, it is easy to show analytically that nonmonotonicity can arise on the Markov tree. Suppose for example that the density has reached equilibrium along one branch of the tree, i.e. let $\rho_S = 1$ for a set of states $DS, S, S0$ and $S00$, but that the density of

$S1$ and $S01$ is zero. Then, providing only that $P_{S \rightarrow S1} > P_{S0 \rightarrow S01}$, (which is satisfied by the model (11) if $\epsilon_0 < 1$) one iteration of (3) leads to $\rho'_{S0} > \rho'_S$. This occurs because it is more likely for particles in state S than those in $S0$ to move into the branches of the tree, thus the density in S decreases more.

Of course, even though the occupation number is transiently not monotone, it asymptotically relaxes to the constant density equilibrium, as is also consistent of our computations for the standard map.

It would be difficult to provide a more quantitative comparison of the Markov tree model with the standard map. This would require choosing the appropriate states for the map and constructing an algorithm that would decide which of state contains the phase point z_t . For the “short time” dynamics of $t \leq 10^{10}$, the effective structure of the tree will not be close to the self-similar assumption, since as we have seen the phase point does not actually get too deeply into the tree structure (four or fewer levels as we saw for the Markov chain).

5. Conclusions

We have seen that a chaotic orbit on an irregular component appears to be numerically ergodic. The measure defined by such an orbit limits to the constant, area measure but exhibits several anomalies along the way. Most conspicuous are the density spikes (up to a factor of 100) that occur around small elliptic islands embedded in the irregular component. Though these spikes are transient, they might as well, for all practical purposes, be eternal: even after 10^{10} iterates a small fraction of the phase space has density up to three times the mean. Since in most cases 10^{10} might as well be eternity, it would be better to use a nonuniform transient measure in applications, as opposed to the constant invariant measure.

The density spikes are inconsistent with a nearest neighbor Markov model of transport.

For this model, if the density begins as a monotone function of distance from a bounding invariant circle, it must remain monotone. However, a Markov tree model can account for nonmonotonicity of the density. Metaphorically, this occurs because the density on a large branch of the tree can more easily disperse into small branches than can density on the small branches disperse into the twigs.

Finally we mention another possible explanation for the density spikes – the presence of local “accelerator modes.” These could also destroy the nearest neighbor connectivity of the transport model, allowing large density to be transported immediately past intervening low density regions. We have not investigated the existence of such modes.

It remains an open problem to provide a rigorous theory of these phenomena.

Acknowledgements

This research was begun when Harvey Segur asked the right question. Helpful discussions with Robert McLachlan, Robert Easton and John Williamson are gratefully acknowledged. Partial support was obtained from the National Science Foundation under grant DMS-93-05847.

References

- [1] A.S. Wightman, The mechanisms of stochasticity in classical dynamical systems, *Perspectives in Statistical Physics*, (ed.) H.J. Raveché (North Holland, Amsterdam, 1981).
- [2] R.W. Easton, Trellises formed by stable and unstable manifolds in the plane, *Trans. Am. Math. Soc.* 294 (1986) 719–732.
- [3] D.K. Umberger and J.D. Farmer, Fat fractals on the energy surface, *Phys. Rev. Lett.* 55 (1985) 661–664.
- [4] L.A. Smith and E.A. Spiegel, Strange accumulators, *Chaotic Phenomena in Astrophysics*, (eds.) J.R. Buchler and H. Eichhorn (New York Academy of Sciences, New York, 1987).
- [5] R.S. MacKay, J.D. Meiss and I.C. Percival, Transport in Hamiltonian systems, *Physica D* 13 (1984) 55–81.

- [6] J.D. Hanson, J.R. Cary and J.D. Meiss, Algebraic decay in self-similar Markov chains, *J. Stat. Phys.* 39 (1985) 327–345.
- [7] J.D. Meiss and E. Ott, Markov tree model of transport in area preserving maps, *Physica D* 20 (1986) 387–402.
- [8] J.M. Greene, R.S. MacKay and J. Stark, Boundary circles for area-preserving maps, *Physica D* 21 (1986) 267–295.
- [9] V. Rom-Kedar and S. Wiggins, Transport in two-dimensional maps, *Arc. Ration. Mech. Anal.* 109 (1988) 239–298.
- [10] Q. Chen, I. Dana, J.D. Meiss, N. Murray and I. Percival, Resonances and transport in the Sawtooth map, *Physica D* 46 (1990) 217–240.
- [11] R.S. MacKay, A renormalisation approach to invariant circles in area-preserving maps, *Physica D* 7 (1983) 283–300.
- [12] J.D. Meiss, Symplectic maps, variational principles, and transport, *Rev. Mod. Phys.* 64 (1992) 795–848.
- [13] D.J.D. Earn and S. Tremaine, Exact numerical studies of Hamiltonian maps: iterating without roundoff error, *Physica D* 56 (1992) 1–22.
- [14] F. Rannou, Numerical study of discrete plane area-preserving mappings, *Astron. Astrophys.* 31 (1974) 289–301.
- [15] I.C. Percival and F. Vivaldi, Arithmetical properties of strongly chaotic motion, *Physica D* 25 (1987) 105–130.
- [16] F. Vivaldi, Dynamics over irreducible polynomials, *Nonlinearity* 5 (1992) 941–960.
- [17] S. Hammel, J.A. Yorke and C. Grebogi, Numerical orbits of chaotic processes represent true orbits, *Bull. Amer. Math. Soc.* 19 (1988) 465–470.
- [18] S. Hammel, J.A. Yorke and C. Grebogi, Do numerical orbits of chaotic dynamical processes represent true orbits?, *Complexity* 3 (1987) 136–145.
- [19] T. Sauer and J.A. Yorke, Rigorous verification of trajectories for the computer simulation of dynamical systems, *Nonlinearity* 4 (1991) 961–979.
- [20] R. Eykholt and D.K. Umberger, Characterization of fat fractals in nonlinear dynamical systems, *Phys. Rev. Lett.* 57 (1986) 2333–2336.
- [21] R.W. Easton, J.D. Meiss and S. Carver, Exit times and transport for symplectic twist maps, *Chaos* 3 (1993) 153–165.
- [22] R.S. MacKay, J.D. Meiss and I.C. Percival, Resonances in area preserving maps, *Physica D* 27 (1987) 1–20.
- [23] C.F.F. Karney, Long time correlations in the stochastic regime, *Physica D* 8 (1983) 360–380.
- [24] J.D. Meiss, Class renormalization: islands around islands, *Phys. Rev. A* 34 (1986) 2375–2383.