Roundoff-induced periodicity and the correlation dimension of chaotic attractors

Celso Grebogi,* Edward Ott,* and James A. Yorke†
Naval Surface Warfare Center, R-41, Silver Spring, Maryland 20903-5000
(Received 11 April 1988)

Due to roundoff, digital computer simulations of orbits on chaotic attractors will always eventually become periodic. The expected period, probability distribution of periods, and expected number of periodic orbits are investigated for the case of fractal chaotic attractors. The expected period scales with roundoff ϵ as $\epsilon^{-d/2}$, where d is the correlation dimension of the chaotic attractor.

Numerical experiments using digital computers have played the key role in many fundamental advances in the understanding of chaotic dynamics. It is therefore of great interest to investigate inherent limitations imposed by computer roundoff on numerical investigations of chaos. In particular, since computers approximate the continuum of numbers by a discrete set, any orbit on a chaotic attractor must repeat exactly, if the orbit is followed long enough. Hence, such orbits always become periodic. Since *true* (i.e., without roundoff) chaotic orbits do not repeat, computer simulations of orbits on chaotic attractors may be brought into question if the orbit is followed too long—in particular, longer than the expected computer-roundoff-induced periodicity time. It is our purpose here to investigate this phenomenon for the case of fractal chaotic attractors of D-dimensional maps. For related work see Refs. 1–4. The new feature in the present paper is the treatment of the fractal case[3] [in particular Eq. (1), below]. The principal results are as follows.

(1) The expected value of the period scales with round off ϵ as

$$\bar{m} \sim \epsilon^{-d/2} . \tag{1a}$$

Here D is the correlation dimension[5] of the chaotic attractor,

$$d = \lim_{\delta \to 0} \frac{\ln \sum_i p_i^2}{\ln \delta} , \tag{1b}$$

where we construct a D-dimensional cubic grid of edge length δ and denote by p_i the frequency with which the chaotic orbit visits the ith cube.

(2) The periods m have substantial statistical fluctuation. That is, $P(m)$, the probability that the period is m, is not strongly peaked around $\bar{m}$. This probability is given by

$$P(m) = (1/\bar{m}) F(\sqrt{\pi/8}\, m/\bar{m}) , \tag{2}$$

where

$$F(x) = \sqrt{\pi/8} \int_x^\infty \exp(-x^2/2) dx$$
$$= \sqrt{\pi/8} [1 - (\pi/2)^{1/2} \mathrm{erf}(x/\sqrt{2})] ,$$

and erf denotes the error function. See also Ref. 2 for this result.

(3) If one puts down l initial conditions, one expects, on average, to find $\bar{N}_p(l)$ different roundoff-induced periodic orbits, where

$$\bar{N}_p(l) = \sum_{j=1}^{l} \frac{1}{2j-1} . \tag{3a}$$

An excellent approximation to Eq. (3a), which is good to within better than 0.3% for $l \geq 2$, is

$$\bar{N}_p(l) \simeq \tfrac{1}{2} \ln l + 0.982 . \tag{3b}$$

Note that Eqs. (1)–(3) do not depend on the details of the map considered and that Eqs. (2) and (3) are independent of ϵ so long as $\epsilon \ll 1$.

In what follows we first present numerical experiments on Eqs. (1)–(3) for a particular two-dimensional map with a chaotic attractor. Following that, we discuss Eqs. (1)–(3).

The numerical results we present are for the example of the Ikeda map ($z \equiv x + iy$),

$$z_{n+1} = p + B z_n \exp i[\kappa - \alpha/(1 + |z_n|^2)] , \tag{4}$$

with parameters $p = 1$, $B = 0.9$, $\kappa = 0.4$, and $\alpha = 6$. This map is a model of a nonlinear laser cavity system. Figure 1 shows the resulting attractor. For this attractor, we nu-

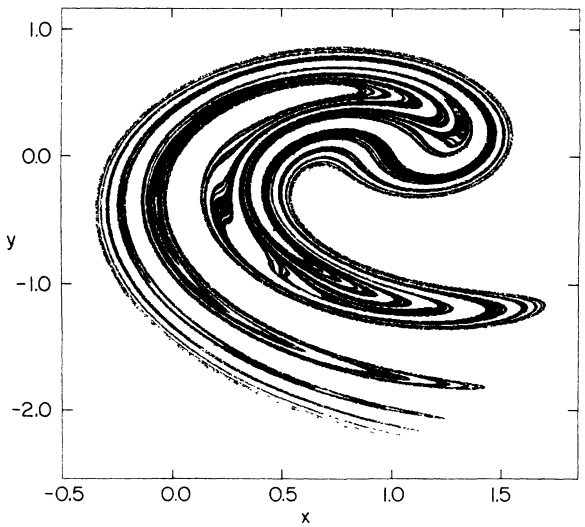

FIG. 1. Chaotic attractor for the Ikeda map.

merically determined the correlation dimension to be $d \simeq 1.67$. To simulate roundoff at a level ϵ, where ϵ is greater than the actual roundoff of the computer used, we replace all numbers after each iteration by a rounded-off number obtained as follows:

$$x \rightarrow \epsilon \, \mathrm{int}(x/\epsilon) ,$$

where $\mathrm{int}(u)$ is the integer part of u. (The same results for the average period versus ϵ were obtained when this roundoff operation was performed after each line in our computer code.)

Figure 2 shows a log-log plot of data obtained for the period m resulting from a single initial condition $z_0 = 0.1 + i0.1$ and 300 roundoff levels in the range $3 \times 10^{-7} \leq \epsilon < 10^{-2}$. In each case we iterated the rounded-off map until, at some iterate $n+1$, z repeats a previous point on its orbit (Fig. 3), $z_{n+1} = z_j$, $n \geq j$. At this point we say that the orbit has "locked in" to a periodic orbit of period $m = n - j + 1$. The data in Fig. 2 show substantial scatter. In particular, the results for nearby ϵ values appear to be essentially uncorrelated. Nevertheless, an overall power-law trend is clearly evident, as indicated by the line with the slope predicted by Eq. (1) (slope $= -d/2 = 0.835$) which we have superimposed on Fig. 2. In order to make the agreement with Eq. (1) more easily discernible and precise, it is desirable to average in such a way as to eliminate the scatter. To do this we calculate an average period $\overline{m}$ corresponding to a roundoff level ϵ_0 by taking a small interval $\Delta\epsilon$ about ϵ_0, $\epsilon_0 - \Delta\epsilon/2 \leq \epsilon \leq \epsilon_0 + \Delta\epsilon/2$, choosing a large number of ϵ values in this interval, and averaging the resulting periods over the ϵ values in this interval. Results of such calculations are shown as the dots in Fig. 4, in which we have also plotted a line of best least-squares fit. Each dot corresponds to $\Delta\epsilon/\epsilon_0 = 0.1$ with 100 values of ϵ in each interval. The slope determined from the least-squares fit is 0.832 as compared to $d/2 \simeq 0.835$.

Next we wish to numerically compute the probability distribution of m. To do this we choose a small interval in ϵ centered at $\epsilon = 10^4$, $0.95 \times 10^{-4} < \epsilon < 1.05 \times 10^{-4}$. Using 9000 evenly spaced ϵ values in this interval, we determine the m value for each ϵ. The same initial condition, $z_0 = 0.1 + i0.1$ is used each time. We then make a

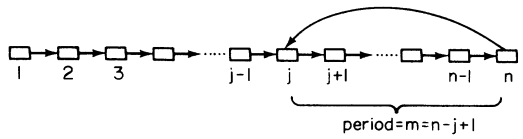

FIG. 3. Schematic illustrating the generation of a periodic orbit from a single initial condition.

histogram of $P(m)$ versus m using 30 equal size bins for m between 1 and 4.1 $\overline{m}$. The results are plotted in Fig. 5 where the dots correspond to the centers of the 30 bins. Also shown in Fig. 5 as a solid curve is the theoretical prediction, Eq. (2). Again the agreement is very good.

Finally we wish to check Eq. (3). For definiteness we choose to consider the case of ten initial conditions ($l = 10$), which, according to Eq. (3), should result in an average of 2.13... distinct periodic orbits. Again we choose a large number of ϵ values in a narrow interval about some central value. For each such ϵ value in the interval, we follow the orbits originating from ten randomly chosen initial conditions until they lock into periodic orbits. We then compare these ten periodic orbits to each other to determine if they are the same or not, and count the number N_p of distinct periodic orbits. Thus for each ϵ we determine N_p. Averaging N_p over the ϵ values in the small interval, we obtain $\overline{N}_p$. We have done this using 300 values of ϵ near $\epsilon = 10^{-4}$; we obtain $\overline{N}_p = 2.1 \pm 0.1$, in good agreement with the value predicted by Eq. (4), $\overline{N}_p = 2.13\ldots$. A somewhat surprising feature of Eq. (3) is its independence of ϵ. We therefore repeated the procedure at a different ϵ, namely, $\epsilon = 10^{-5}$ (100 values of ϵ were chosen in the interval) with the result $\overline{N}_p = 2.04 \pm 0.15$, again in agreement with Eq. (3).

We now turn to the derivation of Eqs. (1)–(3). Imagine that we construct a cubic grid of grid size ϵ in the phase space. Now suppose that we iterate the *exact* map (i.e., without roundoff), keeping track of which cubes in the grid have been visited by the orbit. We iterate until the orbit falls in a previously visited cube. Let $n+1$ denote the iterate at which this happens, and let j denote the iterate at which the cube was first visited. Let $q = n$

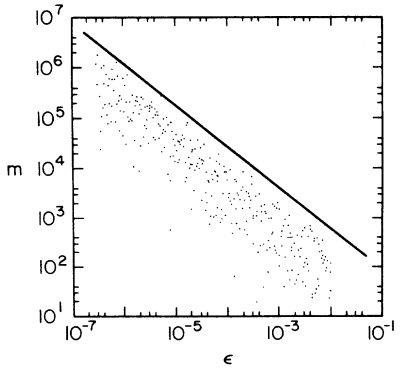

FIG. 2. Log-log plot of m vs ϵ for 300 values of ϵ evenly spaced along the logϵ axis between $\epsilon = 3 \times 10^{-7}$ and $\epsilon = 10^{-2}$.

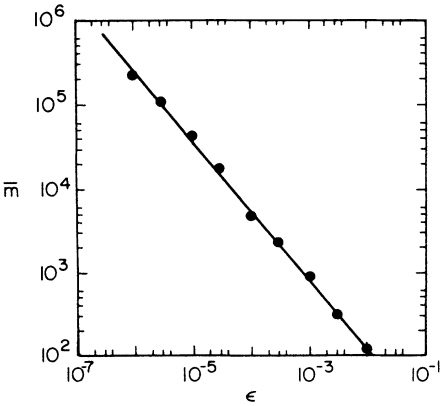

FIG. 4. Log-log plot of $\overline{m}$ vs ϵ.

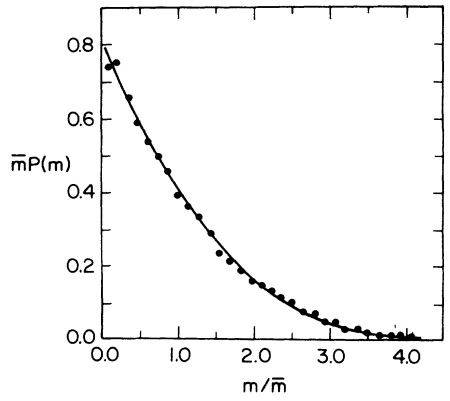

FIG. 5. $\bar{m}P(m)$ vs $m/\bar{m}$. Theory, Eq. (2), is the solid curve. Dots are data from the numerical experiments.

$-j+1$. (Note the analogy with Fig. 2.) Now imagine that we choose many random initial conditions according to the natural measure on the attractor, and that we calculate q for each one. The two fundamental hypotheses leading to our results in Eqs. (1)–(3) are the following.

(i) The statistics for q (generated using the exact map and randomly chosen initial conditions) is the same as the statistics for m (using the map with roundoff ϵ). In particular, if we write the probability distribution of q as

$$\hat{P}(q)=(1/\bar{q})\hat{F}[\sqrt{\pi/8}(q/\bar{q})] ,$$

where $\bar{q}$ denotes the mean value of q, then F [defined in Eq. (2)] and $\hat{F}$ are the same function.

(ii) The scaling of $\bar{q}$ and $\bar{m}$ with ϵ are the same (e.g., for the data in Fig. 2 we find that $\bar{m}\simeq 4\bar{q}$).

We have no rigorous argument supporting these hypotheses (although they seem reasonable on intuitive grounds). However, the excellent agreement between the resulting analytical predictions and the numerical experiments strongly suggests their validity.

We wish to obtain the distribution function P and the scaling of $\bar{m}$ using the hypotheses stated above. Thus we consider iterates of the *exact* map and where they fall in a cubic grid of edge length ϵ. Suppose we choose an initial condition at random with respect to the natural measure on the attractor. Suppose, in addition, that we iterate this initial condition for n iterates and find that the orbit has not visited any box of the grid twice. We then pick one of these n previously visited boxes. Let j denote the iterate at which this box was visited. What is the probability that iterate $n+1$ will fall in box j? To answer this question we assume that the dynamics on the attractor is mixing[6] and that the time between the first and second visits of the box is long. By long we mean that this time, $n-j$, is large enough so that, on the scale ϵ of the grid, the attractor measure in the box visited at iterate j has essentially spread over the attractor. This mixing process is particularly effective because of the subsequent exponential divergence (chaos) of points which were in the chosen box at iterate j. For the purposes of our subsequent calculation, a typical value of $n-j$ is $\bar{q}$, the average q. Thus we require $\bar{q}$ to be larger than the effective mixing time for the ϵ grid. (Note that for our numerical experiments $\bar{m}\simeq 4\bar{q}$ and $\bar{m}\gg 1$; e.g., $\bar{m}\sim 5000$ for $\epsilon\sim 10^{-4}$.) With these assumptions, the answer to our question is that the probability of repeating box j is the natural measure of the attractor in box j. This measure, however, is itself a random variable. This is because the initial condition was picked randomly. Thus its jth iterate can be anywhere on the attractor. Hence the probability of a specific box being visited on the jth iterate is also the attractor measure in the box. Therefore, the probability of the box visited on iterate j being revisited on iterate $n+1$ is

$$\langle p\rangle=\sum_i p_i^2 . \tag{5}$$

Thus the probability that an orbit, which does not repeat after n iterates, repeats at iterate $n+1$ by falling into *any one* of the previously visited n boxes is approximately

$$p_a(n)=n\langle p\rangle . \tag{6}$$

The probability that iterate $n+1$ does not repeat is $[1-p_a(n)]$. The probability that an orbit of length n has no repeats is thus

$$p_b(n)=[1-p_a(1)][1-p_a(2)]\cdots[1-p_a(n-1)] ,$$

or

$$p_b(n)=[1-\langle p\rangle][1-2\langle p\rangle]\cdots(1-(n-1)\langle p\rangle] .$$

This expression for $p_b(n)$ can be approximated for $n\langle p\rangle\ll 1$ as follows:

$$\ln p_b(n)=\sum_{i=1}^{n-1}\ln(1-i\langle p\rangle)\simeq-\sum_{i=1}^{n-1} i\langle p\rangle\simeq-n^2\langle p\rangle/2 ,$$

and thus

$$p_b(n)\simeq\exp[-(n^2\langle p\rangle/2)] . \tag{7}$$

The approximation $n\langle p\rangle\ll 1$ is justified, if $\bar{q}\langle p\rangle\ll 1$. Since we shall obtain $\bar{q}\sim\langle p\rangle^{-1/2}$, we require $\bar{q}\gg 1$ for the validity of Eq. (7). (This requirement is clearly satisfied for the data plotted in Fig. 2.) The probability that an orbit goes n steps without repeat and then repeats on step $(n+1)$ is

$$p(n)=p_a(n)p_b(n) . \tag{8}$$

It is equally likely that the orbit point $n+1$ map to each of the previous n orbit points (see Fig. 3). Thus, given that repeat occurs for the first time at $(n+1)$, the probability that $n-j+1=q$ is

$$p_c(n,q)=\begin{cases} n^{-1} & \text{for } q\le n , \\ 0 & \text{for } q>n . \end{cases} \tag{9}$$

Hence, for large q, the probability distribution of q is

$$\hat{P}(q)=\int_0^\infty p(n)p_c(n,q)dn . \tag{10}$$

Combining Eqs. (6)–(10), we have

$$\hat{P}(q)=\sqrt{8/\pi}\langle p\rangle^{1/2}\hat{P}(\langle p\rangle^{1/2}q) , \tag{11}$$

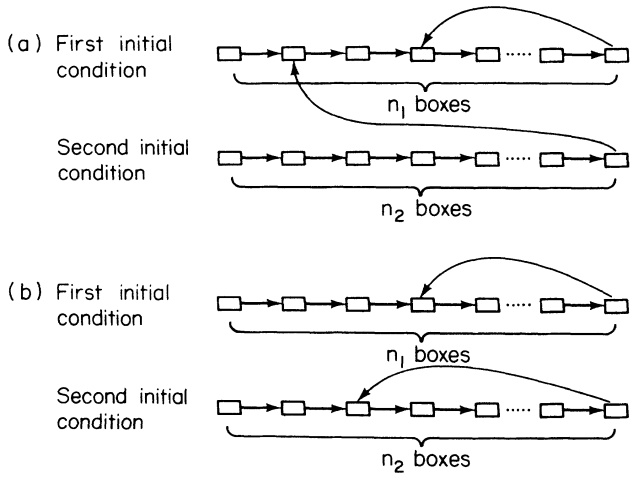

FIG. 6. Schematic illustrations of what can happen when two initial conditions are put down. (a) The second initial condition does not generate a new periodic orbit. (b) It does.

where $\hat{F}$ is the function defined following Eq. (2). Evaluating the integral $\bar{q}=\int_0^\infty q\hat{P}(q)dq$, we obtain

$$\bar{q}=\sqrt{\pi/8}\langle p\rangle^{-1/2} . \tag{12}$$

From Eq. (1b) and our hypothesis (ii), Eq. (12) yields $\bar{m}\sim\epsilon^{-d/2}$, which is Eq. (1a). Also, using Eq. (12) to eliminate $\langle p\rangle$ from Eq. (11), hypothesis (i) yields the desired expression for $P(m)$, Eq. (2). [In addition, from (9), the average time to lock in to a periodic orbit is $2m$.]

We note that Eq. (2) [and also Eq. (3)] do not explicitly involve $\langle p\rangle$. Thus they would also apply if each box had equal probability, in which case we would have $\langle p\rangle=1/N(\epsilon)$, where $N(\epsilon)$ is the number of boxes in the grid needed to cover the attractor. Hence with the replacement $\langle p\rangle\to 1/N$, the result (2) is the same as that for a random map of a set of N points into itself (i.e., a map where every point has equal probability of mapping to any point in the set). There is an extensive literature on random maps, and Eq. (2) has been previously derived.[2] We have included the derivation here for completeness and also because the intermediate steps are necessary for the derivation of Eq. (3) (which we have not found in previous literature).

We now derive Eq. (3). For illustrative purposes, first consider the case of two initial conditions on the attractor, $l=2$. We iterate the first initial condition until it repeats (necessarily generating a periodic orbit). Let n_1 denote the number of iterates of the first initial condition before repeat. Now we iterate the second initial condition until it repeats after, say, n_2 iterates by falling in one of the n_1 boxes visited by the orbit generated by the first initial condition or by falling in one of the n_2 boxes it visited by the orbit generated by the second initial condition. These two possibilities are illustrated in Figs. 6(a) and 6(b). In the former case [Fig. 6(a)], the second initial condition fails to generate a new periodic orbit, while, in the latter case [Fig. 6(b)], it does generate a new periodic orbit. Let P_2 denote the probability that the second initial condition will generate a new periodic orbit, i.e., that Fig. 6(b) will apply. Similarly, we let P_j be the probability that the jth initial condition will generate a new periodic orbit. In terms of the P_j, the average number of periodic orbits produced by l initial conditions is

$$\bar{N}_p(l)=\sum_{j=1}^{l} P_j . \tag{13}$$

(Clearly $P_1=1$.) To derive Eq. (3) we need to evaluate P_j. [We will obtain $P_j=1/(2j-1)$.] The probability that the kth initial condition on the attractor, will go n_k iterates without repeat *given* that initial conditions $1,2,\ldots,k-1$ went $n_1,n_2,\ldots,n_{k-1}$ iterates before repeat is [cf. Eq. (7)]

$$p_k=\frac{\exp[-(n_1+\cdots+n_k)^2\langle p\rangle/2]}{\exp[-(n_1+\cdots+n_{k-1})^2\langle p\rangle/2]} .$$

[We assume $(n_1+\cdots+n_k)\langle p\rangle\ll 1$.] Given that initial condition k goes n_k iterates without repeat, the probability that it will repeat at iterate n_{k+1} is

$$p_k'=(n_1+\cdots+n_k)\langle p\rangle ,$$

and the probability that its repeat will generate a new periodic orbit is

$$p_k''=n_k\langle p\rangle .$$

Hence the probability that the kth initial condition will generate an orbit that repeats for the first time at step n_k+1 is p_kp_k', and the probability that it will repeat for the first time at step n_k+1 and will generate a new periodic orbit is p_kp_k''. Thus

$$P_j=\int_0^\infty dn_j\cdots\int_0^\infty dn_2\int_0^\infty dn_1(p_1p_1'p_2p_2'\cdots p_{j-1}p_{j-1}'p_jp_j'') .$$

Letting $y_k=(n_1+n_2+\cdots+n_k)\langle p\rangle^{1/2}$, the above integral becomes

$$P_j=\int_0^\infty dy_j\cdots\int_0^{y_3}dy_2\int_0^{y_2}dy_1y_1y_2\cdots y_{j-1}(y_j-y_{j-1})e^{-y_j^2/2} .$$

Performing the integrations we obtain

$$P_j=1/(2j-1) , \tag{14}$$

from which Eq. (3a) follows. To obtain the approximation, Eq. (3b), we note the definition of Euler's constant

$$\gamma\equiv\lim_{m\to\infty}[1+(\tfrac{1}{2})+(\tfrac{1}{3})+\cdots+(1/m)-\ln(m)]$$

in terms of which we obtain for the sum in Eq. (3a),

$\bar{N}_p = (\frac{1}{2}) \ln l + (\ln 2 + \gamma/2) + o(1)$, for large l.

For the validity of Eq. (7) and hence the derivative of Eq. (14) we required $(n_1 + n_2 + \cdots + n_l)\langle p \rangle \ll 1$. The expected value of $n_1 + \cdots + n_l$ is

$$\int_0^\infty \cdots \int_0^\infty (n_1 + \cdots + n_l) p_1' p_1 \cdots p_l' p_l dn_1 \cdots dn_l ,$$

which is approximately $\sqrt{2l/\langle p \rangle}$. Thus Eq. (3) requires $l\langle p \rangle \ll 1$ or $l \ll \bar{m}^2 \sim \epsilon^{-d}$ for its validity. This is well satisfied for our numerical experiments on Eq. (3) and in typical practical situations. Another reason why we must have $l\langle p \rangle \ll 1$ for Eq. (3) to be valid is that we take the probability of falling into a previously visited box as constant at $\langle p \rangle$. If we have many initial conditions, the high probability boxes will tend to get used up as more and more boxes are visited. Hence the probability of return to a box visited by the orbit generated by initial condition l will decrease with l if l becomes so large that $l\langle p \rangle \ll 1$ is violated. A problem somewhat related to Eqs. (3) has been considered by Kruskal[7] who calculates the expected number of *all* periodic orbits of a random map of a set of N points into itself. He obtains for this quantity $(\frac{1}{2})\ln N + \gamma + o(1)$, for large N. This equation would be difficult to realize in numerical experiments because of the large number of initial conditions ($l \sim N$) needed to obtain all periodic orbits. Also to treat chaotic attractors such as Fig. 1 one would have to generalize Kruskal's result to the case of multifractal chaotic attractors (i.e., to highly nonuniform p_i).

As an example of an application of Eq. (1), we discuss the case of one-dimensional maps of the form $x_{n+1} = 1 - 2|x_n|^z$, $z > 0$, which was previously considered in Ref. 4. In this case almost every initial condition in $-1 \le x \le 1$ generates an orbit yielding a smooth density $\rho(x)$ in $-1 < x < 1$ with singular behavior, $\rho(x) \sim (1 - |x|)^{-(z-1)/z}$ near $x = \pm 1$. In this case the contributions to $\langle p \rangle = \sum_i p_i^2$ are of two types:[8] that arising from the singularities at $x = \pm 1$ and that arising from the bulk of the ϵ intervals in $-1 < x < 1$,

$$\langle p \rangle = \sum_i p_i^2 \sim \kappa_1 \epsilon^{2/z} + \kappa_2 \epsilon \sim \begin{cases} \epsilon & \text{for } z \le 2 , \\ \epsilon^{2/z} & \text{for } z \ge 2 . \end{cases}$$

(The term $\kappa_2 \epsilon$ arises because there are of order $1/\epsilon$ intervals, and each interval, except for those near $x = \pm 1$, has $p_i \sim \epsilon$.) Thus, from Eq. (1b), the exponent in Eq. (1a) is

$$d/2 = \begin{cases} \frac{1}{2} & \text{for } z \le 2 , \\ 1/z & \text{for } z \ge 2 . \end{cases} \tag{15}$$

The equation for d versus z is nonanalytic in z at $z = 2$ (the derivative is discontinuous). This type of behavior was pointed out in Ref. 8 and its analogy to a phase transition noted in Ref. 9. A main result of Ref. 4 is $\bar{m} \sim \epsilon^{-1/z}$, which agrees with Eq. (15) for $z \ge 2$. Reference 4, however, makes no distinction between the $z > 2$ and the $z < 2$ regimes, and applies the formula $\bar{m} \sim \epsilon^{-1/z}$ to $z = 1$ for which it yields[10] the incorrect result $\bar{m} \sim \epsilon^{-1}$ (rather than $\bar{m} \sim \epsilon^{-1/2}$ [cf. Eq. (15)]).

This work was supported by the U. S. Defense Advanced Research Projects Agency under the Applied and Computational Mathematics Program and by the Naval Surface Weapon Center, Navy Dynamics Institute Program. We would like to thank Eric Kostelich and Frank Varosi for making their correlation dimension data available to us, and to Ittai Kan for useful conversations.

---

*Permanent address: Laboratory for Plasma Research, University of Maryland, College Park, MD 20742.

†Permanent address: Institute for Physical Science and Technology, University of Maryland, College Park, MD 20742.

[1]D. E. Knuth, *The Art of Computer Programming* (Addison-Wesley, Reading, MA 1969), Vol. II, p. 8 and pp. 519 and 520. Knuth discusses the problem for random-number generators which are essentially very chaotic one-dimensional maps.

[2]B. Harris, Ann. Math. Stat. **31**, 1045 (1960). This paper obtains results on random maps.

[3]Y. E. Levy, Phys. Lett. **88A**, 1 (1982). This paper also considers fractal strange attractors and does numerical experiments on the Henon map. Formula (1a) of our paper is conjectured by Levy but with d incorrectly identified as the Hausdorf dimension of the attractor rather than the correlation dimension.

[4]C. Beck and G. Roepstorff, Physica **25D**, 173 (1987). This paper investigates one-dimensional maps and also gives numerical results for the Hénon map.

[5]P. Grassberger and I. Procaccia, Phys. Rev. Lett. **50**, 346 (1983).

[6]Single-component chaotic attractors are typically mixing. If we had a chaotic attractor consisting of M disjoint components through which the orbit cycled, we could consider the M-times-iterated map; then each of the individual M components would be a single mixing chaotic component.

[7]M. D. Kruskal, Ann. Math. Monthly **61**, 392 (1954).

[8]E. Ott, W. Withers, and J. A. Yorke, J. Stat. Phys. **36**, 687 (1984).

[9]D. Katzen and I. Procaccia, Phys. Rev. Lett. **58**, 1169 (1987).

[10]Their numerical results for $z = 1$ are also apparently erroneous.

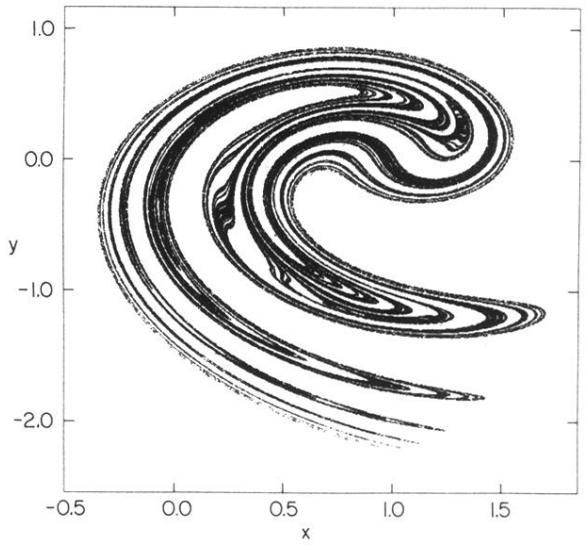

FIG. 1. Chaotic attractor for the Ikeda map.