Informal Remarks on the Orbit Structure of Discrete Approximations to Chaotic Maps

Oscar E. Lanford III

CONTENTS

We report the results of some computer experiments on the orbit structure of the discrete maps on a finite set which arise when an expanding map of the circle is iterated "naively" on the computer. We also comment on what mathematical questions ought to be answered in order to account for the reliability in practice of orbit following on the computer as an indicator of the ergodic properties of the underlying map.

1. INTRODUCTION

It is a fact of experience that computer simulations — of a relatively naive sort — are generally fairly reliable indicators of the properties of concrete dynamical systems. In the interest of brevity, let me explain what I mean by giving an explicit example, and leave it to the reader to think about generalizations. Consider the mapping

$$x \mapsto f(x) := 2x + 0.5x(1-x) \pmod 1, \quad \text{for } 0 \le x \le 1. \tag{1–1}$$

It is perhaps better to think of this mapping as acting on the unit interval with endpoints identified, i.e., on the circle. Note that $f'(x) \ge 1.5$ everywhere, so f is strictly expanding in a particularly clean and simple sense. As a consequence of expansivity, this mapping has about the best imaginable ergodic properties:

- It admits a unique invariant measure μ equivalent to Lebesgue measure.
- The abstract dynamical system (f, μ) is ergodic and in fact isomorphic to a Bernoulli shift.
- A central limit theorem holds, etc.

1058-6458/1998 $0.50 per page
Experimental Mathematics **7**:4, page 317

(One respect in which our example is less than optimal is that, when regarded as acting on the circle, it has a discontinuous first derivative at the origin. This could be fixed by studying instead, say, $x \mapsto 2x + \alpha \sin 2\pi x \pmod 1$ with $|\alpha| < 1/(2\pi)$. We chose the above example because it is cheap to compute and because the effects of round-off error are relatively easy to analyze.)

One consequence of the ergodicity of f relative to μ is that, for Lebesgue almost all initial points x in the unit interval, the corresponding orbit $f^n(x)$ is asymptotically distributed over the unit interval according to μ, i.e.,

$$\lim_{m\to\infty} \frac{1}{m} \sum_{n=0}^{m-1} \varphi(f^n(x)) = \int \varphi(t)\, \mu(dt)$$

for all well-behaved functions φ. (This assertion is simply a clever reading of the pointwise ergodic theorem applied to (f, μ) together with the fact that μ-almost all and Lebesgue-almost all are the same.) Intuitively: If we choose an initial point at random, then we should be very surprised if the corresponding orbit fails to distribute itself as indicated.

This mathematically rigorous result leads one to expect something quite similar to happen when the mapping is iterated *on the computer.* More formally:

Choose an initial point which is not too special (the points 0 and 1, for example, are obviously not typical); compute numerically a few tens of millions of points on the orbit of the point in question; divide the working interval into a few hundred equal subintervals; count the number of points of the computed orbit lying in each of these intervals; and plot the resulting histogram. It would be surprising if this process failed to produce a graph looking very much like that of the density of the absolutely continuous invariant measure.

I want to begin by making the very simple and general point that, *reasonable as this expectation is, it is not so obvious how to **prove** that it is correct.* The reason is that, because of the expansivity of the mapping, the growth of round-off error normally means that the computed orbit will remain near the true orbit with the chosen initial condition only for a relatively small number of steps — typically, of the order of the number of bits of precision with which the calculation is done. It is true that the above mapping — like many "chaotic" mappings — satisfies a *shadowing theorem* which ensures that the computed orbit stays near to some true orbit over arbitrarily large numbers of steps. The flaw in this idea as an explanation of the behavior of computed orbits is that the shadowing theorem says that the computed orbit approximates *some* true orbit, but not necessarily that it approximates *a typical* one. In fact, the simple example $x \mapsto 2x \pmod 1$ shows that computed orbits do not always approximate typical exact orbits (and also makes clear that the expectation expressed in the preceding paragraph is not always fulfilled).

It appears to me that the precise formulation of these questions will require the setting-up of a limiting regime in which precision of calculation and number of steps of iteration both go to infinity, with relations between the rates. In fact: In a very general way, this problem reminds me quite a lot of the notoriously difficult one of deriving non-equilibrium statistical mechanics from atomic physics. As with statistical mechanics, there are two very different length scales, a macroscopic one on which the state space looks like a continuum and the mapping smooth, and a microscopic one on which the state space looks like a discrete set of points and the mapping has a certain amount of jaggedness. This suggests the discouraging possibility that this problem may be as hard of that of non-equilibrium statistical mechanics. As with statistical mechanics, the problem can probably be made *much* easier by the judicious introduction of a stochastic element in the microscopic evolution. As in statistical mechanics, I think this is cheating: For me, a satisfactory solution will have to take seriously the fact that computer iteration is perfectly deterministic.

2. EXPERIMENTAL RESULTS

But these are issues of general philosophy, and I don't want to discuss them in detail here. I mention them only as a preface to saying that what I do want to talk about is an experimental study of an *inappropriate* limiting regime — one in which the number of iterations is *too large* relative to the precision of calculation. The question addressed is the following: The exact mathematical problem concerns iterating a smooth mapping on an interval. The computer, working with fixed finite precision, is able to represent finitely many points in the interval in question. It is probably good, for purposes of orientation, to think of the case where the representable points are uniformly spaced in the interval. The true smooth mapping is then *approximated* by a discretized mapping, sending the finite set of representable points in the interval to itself. Describing the discretized mapping exactly is usually complicated, but it is *roughly* the mapping obtained by applying the exact smooth mapping to each of the discrete representable points and "rounding" the result to the nearest representable point. (The reason why this simple description is not quite realistic is that, in practice, intermediate quantities, and not just the final result, undergo rounding.) However the discretization is done, the upshot is that what is really iterated on the computer is *a mapping of a finite — albeit large — set to itself.* Every orbit of such a mapping is, trivially, eventually periodic, i.e., eventually lands exactly on a periodic cycle. The question addressed here is the orbit structure of such a discretized map:

- How many periodic cycles are there, and what are their periods?
- How large are their respective basins of attraction, i.e., for each periodic cycle, how many initial points give orbits which eventually land on the cycle in question?

I have done two kinds of experiments to explore this question:

- For relatively coarse discretizations — say about 10^7 representable points — determine the orbit structure completely, i.e., find all the periodic cycles and the exact sizes of their basins of attraction.
- For iteration using ordinary (IEEE-754) double-precision arithmetic — so that the working interval contains of the order of 10^{16} representable points — sample the orbit structure by choosing some number — 1000 in the case reported here — initial points at random and determining the cycles to which they converge.

For purposes of logical simplicity, it seemed mildly advantageous to look at evenly spaced discretizations. For the experiments with double precision, this was accomplished by shifting the working interval from $[0,1]$ to $[1,2]$, i.e., the mapping actually iterated was

$$x \mapsto 2x + 0.5(x-1)(2-x) \pmod 1 \qquad \text{from } [1,2] \text{ to itself.}$$

Some representative results results are given on the next page (Table 1).

Many more examples could be given, but those given may serve to illustrate the intriguing character of the results: The outcome proves to be extremely sensitive to the details of the experiment, but the results all have a similar flavor: A relatively small number of cycles attract nearly all orbits, and the lengths of these significant cycles are much larger than one but much smaller than the number of representable points.

3. MODELLING

It seems clear that there are regularities here which ought to be understood. *I know of* ***no*** *ideas which contribute, in my judgment, to a* ***fundamental*** *understanding of these regularities.* There is, on the other hand, a very persuasive idea about how one might *model* them. The idea, which I first heard from D. Ruelle (see Section 7 and the note following it), runs as follows: Since the mapping

$N = 2^{22} = 4{,}194{,}304$ 13 cycles		
period	basin size	
3,864	2,523,929	60%
1,337	538,712	13%
718	513,839	12%
295	238,486	6%
130	203,587	5%
1,338	152,942	4%
297	12,359	0.3%
169	5,056	0.1%
97	3,012	$< 0.1\%$
17	2,346	$< 0.1\%$
6	21	$< 0.1\%$
1	8	$< 0.1\%$
7	7	$< 0.1\%$

$N = 2^{24} = 16{,}777{,}216$ 2 cycles		
period	basin size	
5,300	16,777,214	100%
1	2	$< 0.1\%$

$N = 2^{23} = 8{,}388{,}608$ 7 cycles		
period	basin size	
4,898	5,441,432	65%
1,746	2,946,734	35%
13	205	$< 0.1\%$
6	132	$< 0.1\%$
30	96	$< 0.1\%$
4	8	$< 0.1\%$
1	1	$< 0.1\%$

$N = 2^{24} - 1 = 16{,}777{,}215$ 10 cycles		
period	basin size	
3,081	7,502,907	45%
699	3,047,369	18%
3,469	2,905,844	17%
1,012	2,774,926	17%
563	290,733	2%
2,159	221,294	1%
138	21,610	0.1%
421	12,477	$< 0.1\%$
9	54	$< 0.1\%$
1	1	$< 0.1\%$

$N = 2^{25} = 33{,}554{,}432$ 8 cycles		
period	basin size	
4,094	32,114,650	96%
621	918,519	3%
283	516,985	2%
126	2,937	$< 0.1\%$
6	887	$< 0.1\%$
55	433	$< 0.1\%$
4	20	$< 0.1\%$
1	1	$< 0.1\%$

double precision (sampling) 1000 initial points 7 cycles found		
period	"basin size"	
27,627,856	517	52%
88,201,822	320	32%
4,206,988	147	15%
4,837,566	17	2%
802,279	8	1%
6,945,337	6	1%
2,808,977	1	0.1%

TABLE 1. Census of cycles for representative discretizations. N is the order of the discretization; thus the representable points are the numbers j/N, with $0 \le j < N$.

is "chaotic," it is reasonable to think of modeling computed orbits by simply choosing the successive points at random. This model only makes sense, however, until some point has been chosen twice; thereafter, the fundamentally deterministic character of the mapping can no longer be neglected and the future of the orbit is unambiguously determined. An elementary calculation shows that the probability that there is *no* repetition in a set of n points chosen independently from a population of N (with equal weights) is about

$$e^{-n^2/(2N)},$$

provided that N is large and n not too much larger than $\sqrt{N}$. Loosely: The number of steps before the first repetition is typically of the order of $\sqrt{N}$. This rule of thumb is roughly consistent with the experimental results cited above. (A moment's thought shows that the above computation of distribution of "first-repeat times" is equivalent to computing the first-repeat time for a randomly chosen map. Hence, one speaks of the "random-map model" for computing periods of cycles, etc.)

These remarks also permit me to be a little more specific about what I think the right limiting regime for theorem proving about the reliability of numerical experiments of the sort described earlier should be: One should look at a discretization to N points, and at time averages over m steps, with m and N both large but with

$$\log N \ll m \ll \sqrt{N}.$$

The first $\ll$ allows the computed orbit to deviate macroscopically from the true one over most of its length; the second is in any case usually satisfied in practice and ought to mean that the times considered are short enough so that the effects of the strict finiteness of the space of states are not important. In fact: it might be prudent to replace the second $\ll$ by the stronger condition

$$\log m \ll \log N,$$

but it isn't so clear that this condition is satisfied in practice.

4. SPATIAL DISTRIBUTION OF CYCLES

As noted at the beginning, almost all the orbits of the mathematically exact mapping distribute themselves asymptotically over the working interval according to the unique invariant probability measure μ which is absolutely continuous with respect to Lebesgue measure. On the other hand, a very long computed orbit simply runs many times around whatever periodic cycle it lands on, i.e., has the same asymptotic distribution as that cycle. It is therefore interesting to know whether the periodic cycles are distributed according to the absolutely continuous invariant measure. This can of course only hold in an approximate sense, since a periodic cycle is only a finite set of points. Somewhat surprisingly, it does appear that, at least in the example we are considering, the periodic cycles do approximate the absolutely continuous invariant measure quite well. To be somewhat more concrete: I looked, from this point of view, only at the double-precision discretization. The working interval was partitioned into 200 equal subintervals and, for each of the seven cycles listed in Table 1, the points of the cycle in each subinterval were counted. The resulting histogram for the first cycle — the one which seems to attract a majority of the orbits — is shown in Figure 1. The same figure shows the density of the absolutely continuous invariant measure, but the agreement is so good that it is nearly impossible to distinguish the two plots (except when magnified, as in the inset box). The results for the other cycles (not shown) are similar, with the following systematic variation: Very simple ideas lead to the surmise that the occupation numbers n_i of the various intervals should show fluctuations of the order of $\sqrt{n_i}$. The n_i's are roughly proportional to the period of the cycle, and the periods of the cycles vary by about a factor of 100; thus, it is to be expected that the histograms for cycles with relatively small periods look significantly noisier than those for cycles with relatively large periods. This is indeed what happens.

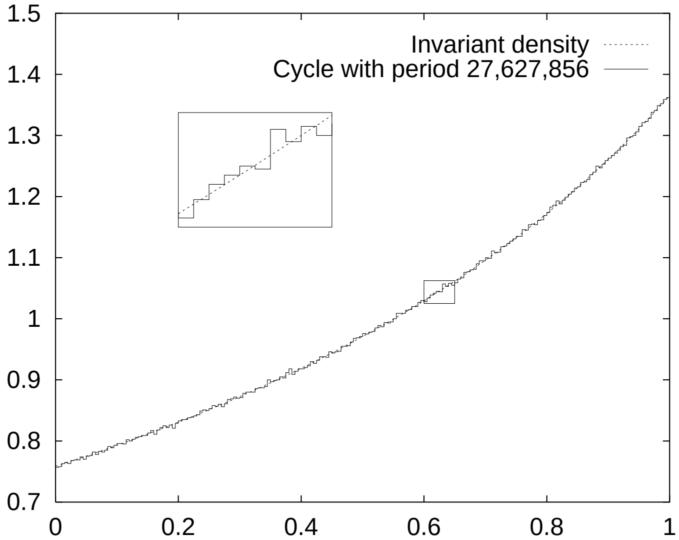

FIGURE 1. Histogram of the main cycle and density of the absolutely continuous invariant measure.

To get a more quantitative measure of the agreement between the distribution of the cycles and the invariant measure, we compute for each cycle the χ^2-statistic,

$$\chi^2 = \sum_{i=1}^{200} \frac{(n_i - \bar{n}_i)^2}{\bar{n}_i}, \tag{4–1}$$

where n_i is the number of points on the cycle lying in the ith subinterval, and $\bar{n}_i$ the "expected" number, i.e., the period of the orbit multiplied by the probability assigned to the ith subinterval by the

period	χ^2
27,627,856	170.480
88,201,822	229.728
4,206,988	122.258
4,837,566	264.920
802,279	184.270
6,945,337	215.295
2,808,977	197.632

TABLE 2. The χ^2-statistic as defined in (4–1) for each of the seven cycles found for the double-precision discretization.

absolutely continuous invariant probability measure. The results are as shown in Table 2. For comparison: For sample of p points chosen independently according to the invariant probability measure, χ^2 has probability about 0.05 of being smaller than 167 and also probability about 0.05 of being larger than 233.[1] The third and fourth orbits look a little suspect at first glance; the probability that χ^2 is < 123 (respectively > 264) is only about 5.4×10^{-6} (respectively 1.4×10^{-3}). On reflection, however, it is clear that these probabilities should be taken with a grain of salt; it would be more convincing to compare the fluctuations of the distributions of periodic cycles about the invariant measure with the typical sizes of fluctuations of long segments of true orbits around this same measure. The latter fluctuations are indeed Gaussian – i.e., the system satisfies a central limit theorem — but the covariance (presumably) does not have the simple form corresponding to an independent choice of p points (and furthermore does not seem to be easy to compute).

5. ANOTHER MAPPING

To show that the above is not the whole story, we present the results of one other experiment — a sampling study in double precision of a discretization of the mapping

$$x \mapsto 1 - 2x^2 \quad \text{on } [-1, 1]. \tag{5–1}$$

This mapping also has excellent ergodic properties but in a more subtle and unstable way than the previous example. (The precise discretization studied is obtained by first exploiting evenness to fold the interval $[-1, 0]$ onto $[0, 1]$, i.e. replacing (5–1) by

$$x \mapsto |1 - 2x^2| \quad \text{on } [0, 1]. \tag{5–2}$$

It is not difficult to see that the folded mapping has the same set of periods as the original one. The working interval is then shifted from $[0, 1]$ to $[1, 2]$ by translation, and the iteration of the translated folded mapping is programmed in a straightforward way.) Out of 1000 randomly chosen initial points,

- 890 — the overwhelming majority — converged to the fixed point corresponding to the fixed point -1 in the original representation (5–1);
- 108 converged to a cycle of period 3,490,273;
- the remaining 2 converged to a cycle of period 1,107,319.

Thus, in this case at least, the very long-term behavior of numerical orbits is, for a substantial fraction of initial points, in flagrant disagreement with the true behavior of typical orbits of the original smooth mapping.

6. SOME DETAILS

As should now be apparent, the orbit structure of a discretized map tends to depend sensitively on the details of the discretization. One consequence is that an attempt to reproduce the reported orbit structures is not likely to give the same results unless care is taken to use *exactly* the same discretization. For the relatively low-precision "artificial" discretization, it is not difficult to describe the discretization precisely. We denote by N the

[1] As is customary, the computation of these probabilities is done under the assumption that the central limit theorem gives a "sufficiently accurate" description of the fluctuations of the n_i about their mean values $\mu_i p$. This assumption should certainly hold very well in this situation, but it seems to be hard to estimate the error reliably.

number of points in the discretized working interval, so that the points themselves are the

$$x_j = \frac{j}{N}, \quad \text{with } 0 \leq j \leq N-1.$$

The discretized map we study can be described as the result of the following three-step procedure:

1. Apply the exact mapping (1–1) to x_j. The result lies between 0 and $2 - 1/N$.
2. If the result of step 1 is $\geq 1 - 1/2N$, subtract 1; otherwise, leave it unchanged. In either case, the result in now in $[-1/2N,\ 1-1/2N)$.
3. Round the result of step 2 to the nearest number of the form k/N. If it lies exactly halfway between two of these lattice points, choose the one with k even.

(We emphasize that this is an "implementation-independent"characterization of the the discretized map, not an algorithm adapted to computing it.) The prescription in step 2 to reduce to the interval $[-1/2N,\ 1-1/2N)$, rather than the more natural-seeming $[0, 1]$, serves to ensure that we round to a number of the form k/N with $k \leq N-1$. Note that the plausible alternative of rounding *before* reducing to the interval $[0, 1)$ gives the same results for N even, but not necessarily for N odd, because of rounding-to-even in the case of a tie.

For the sampling experiments in double precision, it is not so easy to give a short precise specification of how the mapping is discretized. The way the experiment was actually done was to write reasonably straightforward C-language code for the mapping and pass it through a compiler. The C code was:

```
double f (double x) {
    double w;
    w = 2*(x-1)+0.5*(x-1.0)*(2.0-x);
    if (w < 1.0) w += 1.0;
    return(w);
}
```

The compiler used for the experiment reported here was the GNU C compiler, version 2.7.0, running on an Intel Pentium under the Linux operating system. "Unoptimized" compilation was requested. Because the Pentium's internal floating point registers provide 64 bit precision rather than the 52 bits of double precision numbers, the exact discretization algorithm depends on which intermediate quantities are kept in registers and which are rounded to double precision in order to be stored.[2] I have not attempted to sort out in all detail what was really happening in the experiment performed; the objective was simply to imitate how such a computation would be done in practice. It may be worth noting that the above code produces *different* results when compiled and run on a Sparc (again using the GNU C compiler).

7. A FEW BIBLIOGRAPHIC REMARKS

As noted earlier, I learned from D. Ruelle the idea that the "random iteration" described earlier might be a sensible way to model the structure of periodic cycles of chaotic maps (and hence that typical periods should be of the order of the square root of the number of accessible states). Ruelle proposed this idea to account for some anomalies in numerical experiments performed by Y. Levy, and the idea appears in [Levy 1982]. These ideas were developed more generally — and a technical emendation proposed to Levy's Ansatz — in [Grebogi et al. 1988]. Recently, a group at Deakin University and the University of Queensland has been studying discretized mappings systematically. One direction they have pursued is the quantitative study of how well the random-maps model predicts the orbit structure of discretizations. They have also studied ways of improving on the random-maps model to take into account relevant properties of the map being discretized, notably the presence of

[2] There is in fact yet another complication. In its default mode, the Intel floating point hardware performs floating point operations by first rounding to 64 bits, then rounding *that* to 52 when it is stored. Because of the round-to-even tie-breaking rule, this is not the same as rounding directly to 52 bits. Thus, even if all intermediate quantities are stored, the results will not always be the same as with pure 52-bit arithmetic.

critical points. See, for example, [Diamond et al. 1996]; the extensive bibliography of this article also provides a more thorough overview of prior work than can be given here.

NOTE ADDED

I am indebted to the anonymous referee for calling to my attention the very relevant work of T. Erber and his collaborators, in which, among many other things, the idea of modelling the orbit structure of a chaotic map by that of a random map appears prior to the aforementioned work of Levy. In [Erber et al. 1979], orbit structure of discretizations of $x \mapsto 2 - x^2$ and its dependence on the precision of the discretization are studied in very much the same spirit as the experiments reported here. (Curiously, the phenomenon of collapse of a significant fraction of the orbits onto the fixed point — at $x = -2$ in this way of writing the mapping — which was so prominent in our experiments turned up only rarely in those reported in this paper.) This paper and the related one [Erber et al. 1983] present a wealth of intriguing heuristic ideas bearing on the modelling of the orbit structure of these discretizations.

ACKNOWLEDGEMENTS

The programming on which this report is based was mostly done during an extended stay at the Institut des Hautes Études Scientifiques in Buressur-Yvette, France, in the Fall and Winter of 1993–94. I thank the IHES and especially Professor Marcel Berger, its Director at that time, for their hospitality. I am also grateful to Henri Epstein for many stimulating discussions and for encouragement.

REFERENCES

[Diamond et al. 1996] P. Diamond, A. Klemm, P. Kloeden, and A. Pokrovskii, "Basin of attraction of cycles of discretizations of dynamical systems with SRB invariant measures", *J. Statist. Phys.* **84** (1996), 713–733.

[Erber et al. 1979] T. Erber, P. Everett, and P. W. Johnson, "The simulation of random processes on digital computers with Čebyšev mixing transformations", *J. Comput. Phys.* **32** (1979), 168–211.

[Erber et al. 1983] T. Erber, T. M. Rynne, W. F. Darsow, and M. J. Frank, "The simulation of random processes on digital computers: unavoidable order", *J. Comput. Phys.* **49** (1983), 394–419.

[Grebogi et al. 1988] C. Grebogi, E. Ott, and J. A. Yorke, "Roundoff-induced periodicity and the correlation dimension of chaotic attractors", *Phys. Rev. A* **38** (1988), 3688–3692.

[Levy 1982] Y. E. Levy, "Some remarks about computer studies of dynamical systems", *Phys. Lett. A* **88** (1982), 1–3.

Oscar E. Lanford III, Mathematics Department, ETH-Zürich, 8092 Zürich, Switzerland (oscar.lanford@math.ethz.ch, http://www.math.ethz.ch/~lanford)

Received December 3, 1997; accepted in revised form June 22, 1998