

Symbolic dynamics and periodic orbits of the Lorenz attractor*

Divakar Viswanath

Department of Mathematics, University of Michigan, 525 East University Avenue, Ann Arbor, MI 48109-1109, USA

E-mail: divakar@umich.edu

Received 4 November 2002, in final form 17 March 2003

Published 4 April 2003

Online at stacks.iop.org/Non/16/1035

Recommended by J Keating

Abstract

The butterfly-like Lorenz attractor is one of the best known images of chaos. The computations in this paper exploit symbolic dynamics and other basic notions of hyperbolicity theory to take apart the Lorenz attractor using periodic orbits. We compute all 111011 periodic orbits corresponding to symbol sequences of length 20 or less, periodic orbits whose symbol sequences have hundreds of symbols, the Cantor leaves of the Lorenz attractor, and periodic orbits close to the saddle at the origin. We derive a method for computing periodic orbits as close as machine precision allows to a given point on the Lorenz attractor. This method gives an algorithmic realization of a basic hypothesis of hyperbolicity theory—namely, the density of periodic orbits in hyperbolic invariant sets. All periodic orbits are computed with 14 accurate digits.

Mathematics Subject Classification: 34C35, 37D25, 37D45

1. Introduction

The differential equations

$$\dot{x} = \sigma(-x + y), \quad \dot{y} = -xz + rx - y, \quad \dot{z} = xy - bz,$$

with $\sigma = 10$, $r = 28$, and $b = \frac{8}{3}$, were studied by Edward N Lorenz [10] in an article of lasting elegance and clarity. Irregular, turbulent flow patterns that evolve in an irregular, nonperiodic manner are an obstacle to short-term weather forecasting. Lorenz showed that stochastic forcing is not the only way that such irregularity can arise and that the long-time behaviour of deterministic flows can be unstable, bounded, and nonperiodic at the same time. Lorenz based his argument partly on the equations above, now named after him, that are a finite mode truncation of the basic equations of fluid mechanics. It is possible to interpret solutions of the Lorenz equations as convection rolls, even though such an interpretation is not

* This work was supported by the Rackham Grant and Fellowship program, University of Michigan.

physically valid for all choices of the parameters [10, 14]. Lorenz understood the existence of the Lorenz attractor, the possibility of approximating the dynamics on the attractor by a one-dimensional map, the importance of the unstable fixed point at the origin (see figure 3 of [10]), the fractal structure of the attractor, the density of periodic orbits in the Lorenz attractor, and that nonperiodic points in a section of the Lorenz attractor were dense and uncountable. In the concluding section of his paper, Lorenz declared: ‘The foundation of our principal result is the eventual necessity for any bounded system of finite dimensionality to come arbitrarily close to acquiring a state which it has previously assumed.’ Recurrence is the foundation of all the computations given in this paper.

Figure 1 shows a periodic orbit of the Lorenz equations, not any trajectory like the usual plots. Symbol sequences are assigned to trajectories of the Lorenz attractor by following them forwards and backwards in time. Every time a trajectory passes through A (see figure 1), it is assigned an A and it is assigned a B every time it passes through B . In this way, every trajectory is assigned a bi-infinite label. Periodic orbits have symbol sequences with the same finite block of symbols repeating indefinitely, like decimal expansions of rational numbers. This block can be used to label the periodic orbit. Even though the periodic orbit in figure 1 has 240 symbols and a characteristic multiplier of magnitude $3.06e+59$, its Fourier series is computed with a relative error of $1.1e-14$. If an exact formula were known for that periodic orbit, simply evaluating it in double precision arithmetic will cause at least as much error. We compute numerous periodic orbits with equal accuracy and take apart the Lorenz attractor.

The computations in this paper are not exploratory in nature. Considered together, they form a Galilean experiment that puts a mathematical hypothesis to a severe test. The mathematical hypothesis that motivates this paper comes from hyperbolicity theory [8, 12]. Hyperbolicity theory tries to think of invariant sets as if they were hyperbolic fixed points with no critical eigenvalues. This analogy is developed by splitting the tangent space at every point




Figure 1. This plot shows the periodic orbit $ABA^2B^2 \dots A^{15}B^{15}$ of the Lorenz equations. This periodic orbit is a part of the Lorenz attractor. Its period is 171.863 729 1397 and its leading characteristic multiplier has magnitude $3.06e+59$. The sections marked A and B extend into the y -axis and are used to assign symbol sequences to trajectories of the Lorenz attractor.

of the invariant set into a stable and an unstable subspace. The splitting must be preserved by the flow or the map. A perturbation in a stable direction decreases exponentially with the linearized forward flow. Similarly, a perturbation in the unstable direction decreases exponentially with the linearized backward flow. There are periodic orbits arbitrarily close to any point on a hyperbolic invariant set, just as there are rational numbers arbitrarily close to any real number. The density of periodic orbits in hyperbolic invariant sets is the mathematical fact at the centre of this paper.

Density of periodic orbits can be a consequence of a definition or a part of a definition depending upon how hyperbolic invariant sets are defined. A consideration of recurrence of bounded trajectories, identified by Lorenz as the foundation of his work, makes the density of periodic orbits plausible. Let P be a point on an invariant set like the Lorenz attractor. As Lorenz explained, the flow must eventually bring back the point P (or a point in a small neighbourhood of P) to a point Q close to it. The flow is in addition unstable and continuity implies that splitting into stable and unstable manifolds at Q nearly coincides with the splitting at P , since Q is close to P . Therefore, it is possible to perturb P to P' along the unstable manifold of P and Q to Q' along the stable manifold of Q , in such a way that Q' lies on the trajectory starting at P' and Q' is much closer to P' than Q is to P . If the errors in this linearized argument are controlled, it is possible to make Q' coincide with P' and deduce the existence of a periodic orbit close to P . Recurrences of randomly chosen points are not easy to control, however, because very close recurrences take very long times. We will use symbolic dynamics to make the clever perturbations that are necessary to find periodic orbits that nearly go through a given point on the Lorenz attractor.

Symbolic dynamics and some theoretical deductions essential to our computations are based on a geometric model of the Lorenz attractor [7, 17]. Owing to its use in some of our computations, we give an abridged account of this model. The geometric model is based on a map of $(u, v) \in [-1, 1] \times [-1, 1]$ back to the same region. To state the correspondence of this region to the phase space of the Lorenz equations, it is necessary to take a Poincaré section at $z = 27$. The u -axis sort of runs along the surface of the wings (see figure 1) and extends over the region where the wings overlap (and therefore does not lie on the sections A and B in figure 1). The v -axis is transverse to the u -axis and is in the plane of the Poincaré section like the u -axis. The map of the (u, v) plane is written as $G(u, v) = (f(u), h(u, v))$. The form of the one-dimensional map is shown in figure 2. The slope of this map diverges to $\mp\infty$ as $u \rightarrow \pm 0$. The precise form of h is explained in [8], but below is a simplification that is also




Figure 2. A discontinuous one-dimensional map like the one above is a part of a geometric model of the Lorenz attractor. For convenience, the map $f(u)$ can be taken to map $u \in [-1, 1]$ to the same interval, with a discontinuity at $u = 0$. The interval $[-1, 0)$, which is the left half of the u -axis, corresponds to the symbol A . The right half of the u -axis corresponds to the symbol B .

explained in [8]:

$$\begin{aligned} h(u, v) &= av - b, & u > 0, \\ h(u, v) &= av + b, & u < 0. \end{aligned}$$

For the Lorenz attractor, a is something like 10^{-5} (see section 4) and b is chosen to make $G(u, v)$ a one-to-one map. If $a = 10^{-5}$, $b = \frac{1}{2}$ is an acceptable choice. This map has a Cantor set of lines for its attractor and its hyperbolic structure is easily deduced. In the complete geometric model, the attractor of the Lorenz flow is described as a Cantor set of leaves tied to the unstable manifold of the saddle at the origin [8, 17].

We make two minor changes to the geometric model as described in [8]. The assumption $f'(u) > \sqrt{2}$ is used in [8] to prove that upon repeated applications of f , any open interval covers all of $(-1, 1) - \{0\}$. It is better to assume only that $f'(u) > 0$ and to assume, rather than prove, that upon repeated iterations any open interval covers all of $(-1, 1) - \{0\}$. That $f'(u) > \sqrt{2}$, or $f'(u) > 1$, must be false is obvious from even simple computations; a close examination of the tight winding in the innermost part of the wing in figure 9.3.2 of [14] or in our figure 6, which requires a more sophisticated computation, reveals this (also, see [15]). The other change is that we use the sections A and B shown in figure 1 instead of taking a section where the two wings overlap. With this choice, the interval $[-1, 0)$ of the u -axis corresponds to the intersection of the Lorenz attractor with the section A and the interval $(0, 1]$ corresponds to the intersection of the Lorenz attractor with B ; see figures 1 and 2. Neither change affects the topology of the attractor of the geometric model. However, in section 5 we suggest that the topology of the Lorenz attractor may differ from the topology of the attractor of the geometric model in a subtle way.

The assumption $f'(u) > \sqrt{2}$ is used to deduce that a symbol sequence can correspond to at most one periodic orbit of the one-dimensional map f in the geometric model. Section 2 discusses why we expect the Lorenz attractor to have at most one periodic orbit for a given symbol sequence. This can be deduced for the slightly modified geometric model that we use with the additional assumption $(f^n)'(u) > 1$, for $u \in [-1, 1]$ and some $n > 1$.

2. Symbolic dynamics

The symbol sequence of any point on the Lorenz attractor is a bi-infinite sequence of A s and B s. If the symbol sequence of a point P is $\dots x_{-2}x_{-1}.x_1x_2\dots$, where each x_i is A or B , it means that the forward trajectory from the point P goes through the section x_1 , and then the section x_2 , and so on. Similarly, it means that the backward in time trajectory goes through the section x_{-1} , and then the section x_{-2} , and so on. If the point is already on one of the two sections, that section becomes the x_1 symbol. The sections A and B are defined in figure 1.

The idea behind symbolic dynamics is extremely simple and broadly applicable. One example is to label people in the United States by the sequence of states they visit. For this to be a good labelling scheme, people have to move between states often and move independently of one another. Symbolic dynamics labels trajectories and points in phase space in this way. Symbolic dynamics is useful if, and only if, trajectories diverge rapidly; it will buy very little if used to describe quasi-periodic dynamics, for example. Rapid divergence of trajectories is of course true for the Lorenz attractor.

As mentioned earlier, symbol sequences of periodic orbits look like $\dots ss.ss\dots$, where s is a repeating block of A s and B s. If s is something like ABB and it is rotated to something like BAB , we get another sequence which represents the very same periodic orbit. The rotated symbol sequence only starts at a different point on the same periodic orbit. The symbol sequence ss represents two circuits around the periodic orbit represented by s . A few simple

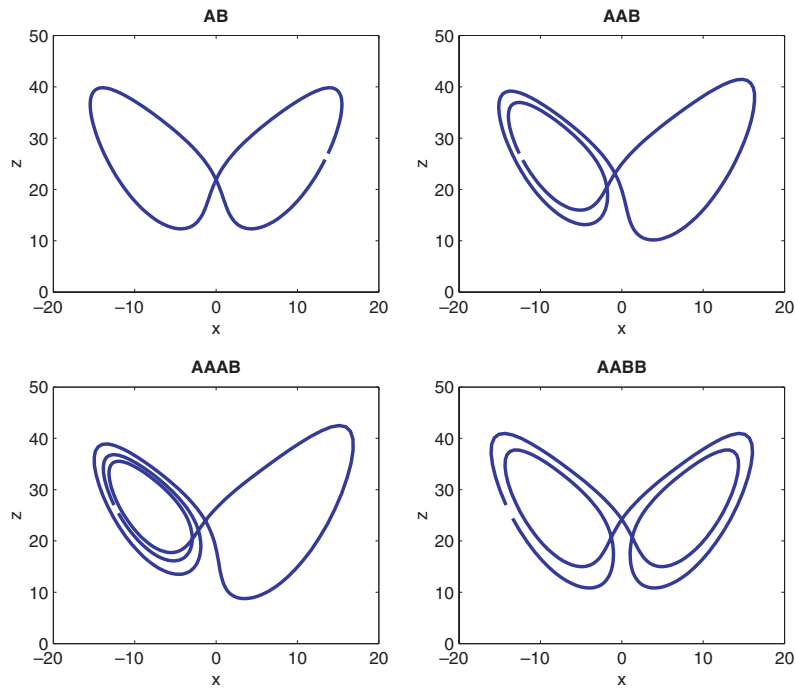


Figure 3. A few simple periodic orbits of the Lorenz equations. Data for these and other periodic orbits are given in table 1.

Table 1. Data for ten periodic orbits. The s column lists the symbol sequence; x and y are the x and y coordinates of a point on that periodic orbit with z coordinate equal to 27. A trajectory starting at this point, which is always on A to begin with, visits the sections A and B (see figure 1) in the order given by s . The Lyapunov exponent λ and period T are also reported.

s	x	y	λ	T
AB	-1.376 361 068 213 4e+01	-1.957 875 194 245 2e+01	9.946e-01	1.558 652 210 716 2e+00
AAB	-1.259 511 539 768 9e+01	-1.697 052 530 708 4e+01	9.607e-01	2.305 907 263 939 8e+00
ABB	-1.442 640 802 503 5e+01	-2.111 123 005 699 4e+01	9.607e-01	2.305 907 263 939 8e+00
AAAB	-1.199 852 328 006 2e+01	-1.568 425 409 688 3e+01	9.192e-01	3.023 583 703 433 9e+00
AABB	-1.291 513 797 031 1e+01	-1.767 310 017 264 6e+01	9.685e-01	3.084 276 775 822 1e+00
ABBB	-1.483 980 067 171 7e+01	-2.208 680 716 008 9e+01	9.192e-01	3.023 583 703 433 9e+00
AAAAB	-1.158 695 166 372 2e+01	-1.481 461 528 812 2e+01	8.838e-01	3.725 641 771 555 8e+00
AAABB	-1.223 140 659 627 3e+01	-1.618 273 059 889 4e+01	9.460e-01	3.820 254 163 436 8e+00
AABAB	-1.269 894 134 991 5e+01	-1.719 749 724 771 3e+01	9.776e-01	3.869 539 112 564 6e+00
AABBB	-1.305 693 014 634 5e+01	-1.798 721 404 928 1e+01	9.460e-01	3.820 254 163 436 8e+00

periodic orbits of the Lorenz equations are shown in figure 3. Table 1 gives data for the first ten periodic orbits. The Lorenz equations are invariant under the change of variables $(x, y, z) \rightarrow (-x, -y, z)$. If s represents a periodic orbit, and s' is obtained by flipping A s to B s and B s to A s in s , the periodic orbit corresponding to s' must exist. It is easily obtained by changing the sign of x and y in the periodic orbit s . From here on, we sometimes identify periodic orbits with their symbol sequences.

The Lorenz equations do not have periodic orbits with the symbol sequence A or with the symbol sequence B . Excluding these, there are in all 111011 symbol sequences of length 20 or less that can represent distinct periodic orbits. We have computed all these periodic orbits

Table 2. This table reports symbol sequences of certain lengths, shown in parentheses, which have the minimum or maximum Lyapunov exponent among all symbol sequences of that length. Of the 111011 periodic orbits with symbol sequences of length 20 or less, the periodic orbit AB has the largest Lyapunov exponent and the periodic orbit $A^{14}B$ has the smallest Lyapunov exponent.

s	λ	s	λ
AB(2)	0.994 645	AB(2)	0.994 645
AAB(3)	0.960 742	AAB(3)	0.960 742
AAAB(4)	0.919 239	AABB(4)	0.968 55
AAAAB(5)	0.883 752	AABAB(5)	0.977 592
AAAAAB(6)	0.854 468	AABABB(6)	0.974 272
AAAAAAB(7)	0.830 498	AABABAB(7)	0.982 872
AAAAAAAB(8)	0.810 897	AABABABB(8)	0.978 859
AAAAAAAAB(9)	0.795 095	AABABABAB(9)	0.985 597
AAAAAAAAB(10)	0.782 318	AABABABABB(10)	0.982 01
AAAAAAAAB(11)	0.772 28	AABABABABAB(11)	0.987 248
AAAAAAAAB(12)	0.764 715	AABABABABABB(12)	0.984 095
AAAAAAAAB(13)	0.759 454	AABABABABABAB(13)	0.988 371
AAAAAAAAB(14)	0.756 413	AABABABABABABB(14)	0.985 579
AAAAAAAAB(15)	0.755 588	AABABABABABABAB(15)	0.989 177
AAAAAAAAB(16)	0.757 063	AABABABABABABABB(16)	0.986 675
AAAAAAAAB(17)	0.761 264	AABABABABABABABAB(17)	0.989 873
AAAAAAAAB(18)	0.768 147	AABABABABABABABABB(18)	0.987 631
AAAAAAAAB(19)	0.778 429	AABABABABABABABABAB(19)	0.990 372
AAAAAAAABBBBBBBBBB(20)	0.781 745	AABABABABABABABABABB(20)	0.988 329

Table 3. This table reports symbol sequences of certain lengths, shown in parentheses, which have the minimum or maximum period among all symbol sequences of the same length.

s	T	s	T
AB(2)	1.558 65	AB(2)	1.558 65
AAB(3)	2.305 91	AAB(3)	2.305 91
AAAB(4)	3.023 58	AABB(4)	3.084 28
AAAAB(5)	3.725 64	AABAB(5)	3.869 54
AAAAAB(6)	4.417 77	AABABB(6)	4.637 14
AAAAAAB(7)	5.103 04	AABABAB(7)	5.429 13
AAAAAAAB(8)	5.783 41	AABABABB(8)	6.194 6
AAAAAAAAB(9)	6.460 26	AABABABAB(9)	6.987 97
AAAAAAAAB(10)	7.134 63	AABABABABB(10)	7.753
AAAAAAAAB(11)	7.807 39	AABABABABAB(11)	8.546 66
AAAAAAAAB(12)	8.479 27	AABABABABABB(12)	9.311 6
AAAAAAAAB(13)	9.150 97	AABABABABABAB(13)	10.105 3
AAAAAAAAB(14)	9.823 17	AABABABABABABB(14)	10.870 2
AAAAAAAAB(15)	10.496 6	AABABABABABABAB(15)	11.664
AAAAAAAAB(16)	11.172	AABABABABABABABB(16)	12.428 9
AAAAAAAAB(17)	11.850 5	AABABABABABABABAB(17)	13.222 6
AAAAAAAAB(18)	12.533 1	AABABABABABABABABB(18)	13.987 5
AAAAAAAAB(19)	13.221 6	AABABABABABABABABAB(19)	14.781 3
AAAAAAAAB(20)	13.918 5	AABABABABABABABABABB(20)	15.546 2

and some data is summarized in tables 2 and 3. The Lyapunov exponent of a periodic orbit is defined as $\log m/T$, where m is the magnitude of its leading characteristic multiplier and T is its period. The periodic orbit $A^{14}B$ has the smallest Lyapunov exponent among all periodic orbits of the Lorenz equations we have computed.

The maximum relative error in the computation of 111011 periodic orbits with symbol sequences of length 20 or less was $3.2e-14$. The average relative error was $1.3e-14$; the mere application of the fast Fourier transform followed by Fourier inversion, the most stable of numerical computations, creates a relative error that is only slightly smaller. All computations were in double precision. The periodic orbits were represented as Fourier series in time, using 64 Fourier points per symbol. If the number of symbols was not a power of 2, the number of Fourier points used was 64 times the first power of 2 that exceeds the number of symbols. For example, the periodic orbit shown in figure 1 was represented using 64×256 Fourier points since it has 240 symbols.

The computation of periodic orbits used the Lindstedt–Poincaré algorithm derived by Viswanath [16] and an extension that is described in the appendix. The extension to the Lindstedt–Poincaré algorithm described in the appendix makes it possible to compute periodic orbits accurately even from inaccurate starting guesses. Use of this extension is essential to carry out the computations described in this paper.

To describe the method used to generate starting guesses for the Lindstedt–Poincaré algorithm, we start by assuming the availability of a library holding all 1375 periodic orbits with symbol sequences of length 13 or less computed with 14 accurate digits. The computation of this library is explained later. Let us say we wish to obtain a starting guess for a periodic orbit with the symbol sequence $s_x = x_1 \dots x_n$, where each x_i is either *A* or *B*. The starting guess is put together by splicing together n slices of trajectories taken from the library. The n slices correspond to the n intervals $(x_1, x_2), \dots, (x_n, x_1)$ in the symbol sequence.

To obtain a slice for an interval (x_i, x_{i+1}) , we make one pass through the library and compute the best fit and the length of the best fit as follows. The length of the forward fit of an interval (y_j, y_{j+1}) of the symbol sequence $s_y = y_1 \dots y_m$ in the library to the interval (x_i, x_{i+1}) of the symbol sequence s_x is equal to k if

$$y_j = x_i, \quad y_{j+1} = x_{i+1}, \dots, \quad y_{j+k-1} = x_{i+k-1}.$$

Above, the subscripts of x are added modulo n and the subscripts of y are added modulo m . The length of the backward fit is equal to $k - 1$ if

$$y_j = x_i, \quad y_{j-1} = x_{i-1}, \dots, \quad y_{j-k+1} = x_{i-k+1}.$$

Again, the arithmetic in the subscripts is modular. The length of the fit is defined as the minimum of the forward fit and a certain number times the backward fit. A good choice for this number would be 14; it is more important to match symbols forwards than backwards (see section 4 for an explanation). In fact, for most of the computations we simply take the length of the forward fit as the length of the fit. A long fit implies that even though the symbols x_i and y_j may correspond to different points on the Poincaré section, the points are likely to be close. The best fit is the interval among all symbol sequences in the library that gives the longest fit to (x_i, x_{i+1}) . Since the arithmetic is modular, the length of the best fit might be infinite. This means that the periodic orbit s_x is identical to some periodic orbit s_y in the library.

The initial guess for the Lindstedt–Poincaré algorithm is put together using the best fits for every interval in s_x . If the best fit of (x_i, x_{i+1}) is to (y_i, y_{i+1}) , the segment of the trajectory of s_y between the symbols y_i and y_{i+1} is taken as the guess for the segment of s_x from x_i to x_{i+1} . The guess is made up of n segments corresponding to the n intervals. If the first segment is between the symbols *A* and *B*, it will take off from section *A* and land on section *B*. The second segment will take off from *B*, but the point it takes off from will not coincide with the point where the first segment lands. Therefore, it is necessary to splice all the segments. We used splines, with appropriate boundary conditions, to splice the segments together. The guess for the period is the sum of the traversal times for each of the n segments.

The library of 1375 periodic orbits was built up as follows. We started with a small library which had all periodic orbits with symbol sequences of length 5 or less. The starting guesses extracted from a Runge–Kutta integration are good enough to build a small library with symbol sequences of length 5 or less. This small library was used to generate starting guesses to build a larger library, and that library was used to generate starting guesses for an even larger library, and we finally ended up with a library of all periodic orbits whose symbol sequences were of length 13 or less. The starting guesses generated using this method may have relative errors as high as 5%. Therefore, the use of the Lindstedt–Poincaré algorithm, with the extension described in the appendix, is essential.

For a concrete example of the use of symbolic dynamics to construct initial guesses for the Lindstedt–Poincaré algorithm, assume that the periodic orbits AB , AAB , and ABB are available, and that the periodic orbit $AABB$ needs to be computed. The match between the first A in AAB and the first A in $AABB$ is valid for three symbols in the forward direction. This is the best match for the first A in $AABB$. Similarly, the best match for the third symbol, which is a B , in $AABB$ is the second symbol in ABB . The method described above uses a library of 1375 periodic orbits to construct initial guesses. This is not the only way to set up the computation. If a periodic orbit with a specified symbol sequence needs to be computed, it is possible to recursively compute only those periodic orbits that are necessary to generate a good guess for the specified symbol sequence.

Why is there a unique periodic orbit corresponding to a symbol sequence such as AB ? We give a heuristic argument using figure 1. The starting point of AB must surely be on the section marked A in figure 1. It can be seen in that figure that the trajectories starting on A that land on B must start from the part of A closer to the origin. Thus the starting point of AB must be on this part of A . After it lands on B , we want the trajectory to return to A . This means only a part of the trajectories that start from A and land on B are valid candidates to be the periodic orbit AB . Continuing this argument further leaves us with a unique periodic orbit AB . An expansion assumption, such as $f'(u) > \sqrt{2}$ or $(f^n)'(u) > 1$ for some $n > 1$, is necessary for a mathematically valid argument based on the geometric model. In most cases, the point of difficulty in constructing adequate symbolic dynamics is in verifying the expansion and contraction requirements. The topological part of the construction, which might involve finding a topological horseshoe for example, is often comparatively easy.

There have been many other computations of periodic orbits of the Lorenz attractor. The Lorenz equations undergo many bifurcations as the parameter r is varied and there are computations of periodic orbits near known bifurcation points [9, 13]. With the same choice of parameters as in this paper, computation of all periodic orbits of the Lorenz attractor with symbol sequences of length 11 or less is reported in [6]. The level of accuracy obtained is 7 digits. In [5], the damped Newton method is used to compute all periodic orbits corresponding to symbol sequences of length 9 or less. The level of accuracy obtained is fewer than 7 digits. The shadowing principle is used to accurately shadow periodic orbits of the Lorenz attractor in [2]. For a result about the Lorenz attractor see [11].

Periodic orbit theory described in [3] is an elaborate framework for computing dynamical averages over chaotic sets using periodic orbits. One might expect that dynamical averages like the Lyapunov exponent can be obtained by taking weighted averages over periodic orbits. This is indeed true; however, it must be noted that in periodic orbit theory the averages are taken over individual periodic orbits as well as groups of periodic orbits, and crucially for fast convergence, some of the weights are negative. This theory treats the Lorenz attractor as a chaotic set without assuming it to be an attractor. The theory gives a method for approximating the escape rate using periodic orbits, as the logarithm of the leading eigenvalue of a certain Fredholm determinant, and this escape rate must be zero for an attractor. The computation of

the escape rate is a useful check of the accuracy of the formulae of periodic orbit theory. For the Lorenz attractor, an estimate of 5.0×10^{-6} for the escape rate may be deduced using all periodic orbits with symbol sequences of length 20 or less.

3. Periodic orbits with long symbol sequences

It is easy to compute periodic orbits with long symbol sequences. For example, the symbol sequence

```

AAABAABABBAABABBABAABBBBAAAAAABABBAAAAABBAABAAAAABBBBBBAABAABBABBBBA
ABBAABBBBABBBAABBBBBAABABBBBBAABABABABAAABBAABBBBABBBAABAAAABBBBABB
AAABBBBBBBBBABAABBBBAAABBBAAABAAAAABBBBAABBBBABBABABBBBBBBABABAAABBA
ABABAABABABABBABAABBBAAAAABABBBAABAAAAABBAABBAABAAAAABBAABAAAAABBBABBBBA
BABBABABBABAABAABAABBBBAABBAABAAAAABBBBABAABBAABAAAAABBBBAABAABAABA
    
```

is of length 347. The periodic orbit with that sequence has period $2.630\,875\,008\,674\,2e+02$ and was computed with a relative error of $1.5e-14$. Its leading characteristic multiplier is of magnitude $9.56e+105$. A periodic orbit with a symbol sequence that is ten times as long can be computed just as easily.

It is clear from table 4 that if a single symbol is removed from the sequence of length 347, the periodic orbit corresponding to the shorter sequence will be so close that the double precision representation of two points on the supposedly distinct periodic orbits will be identical. The question arises in what sense the periodic orbits are being computed. The answer is that each point in the Fourier series is accurate to machine precision and so is the period.

Consider the map $x_{n+1} = 2x_n \text{ mod } 1$ over $[0, 1)$. This map shifts the leading bit with every application. If the interval $[0, \frac{1}{2})$ is labelled 0 and $[\frac{1}{2}, 1)$ is labelled 1, the symbol sequence of every point is identical to its binary expansion. Suppose a machine word can represent only three bits and we wish to compute a periodic orbit with symbol sequence $b_1 \dots b_{10}$. In a computation such as ours, the first point in the orbit will be $.b_1b_2b_3$, the second point will be $.b_2b_3b_4$, and the tenth point will be $.b_{10}b_1b_2$.

We now go back to table 4 and explain a numerical issue that arises in computing distances between periodic orbits. In the Lindstedt–Poincaré algorithm [16], nonlinearities in the vector field $F(x)$ are handled by interpolating at Fourier grid points. This can lead to aliasing errors. Nevertheless, the Lindstedt–Poincaré algorithm computes periodic orbits to machine precision without requiring any more Fourier points than necessary to represent those periodic orbits to machine precision. This is possible because, even though $F(x)$ can be nonlinear, the left-hand side of the differential equation $\dot{x}(t) = F(x)$ is very much linear. A price has to be

Table 4. In this table, s is a randomly generated symbol sequence, whose length is in parentheses. The symbol sequence s' is identical to s but without the last symbol. The second column is the distance between the periodic orbits represented by s and s' at closest approach.

s	Dist(s, s')
ABABB(5)	0.120 595
ABABBAAABA(10)	0.003 582 89
BBAABAABBABABBA(15)	0.000 250 035
ABAAAABBBBBBABBAAAAB(20)	3.497 62e-05
BAABBAAABAABBBBBBABAABAAAB(25)	2.361 38e-07
BBAABAAAABBBAAAAABBBBABAABAABA(30)	1.927 9e-08
AABAAAABABBABBBBBBAAAABABBAABBABBA(35)	3.299 1e-10
ABAABBBBAABBABAABAAAABAAABBBABBABBABABAAAB(40)	2.015 42e-11

paid to eliminate aliasing errors, however, in computing the distance between two periodic orbits.

First consider the problem of finding the distance from a point (x_0, y_0, z_0) to a periodic orbit. A simple strategy, which fails to avoid aliasing errors, is to start by finding the distance d_i from (x_0, y_0, z_0) to the i th point of the periodic orbit on the Fourier grid. The distance of the point to the periodic orbit must be a periodic function itself; we can treat d_i as Fourier interpolants of that periodic function. To find the minimum distance, we simply need to minimize this Fourier series. This method is plagued by aliasing errors, however. If the distance of the point from the periodic orbit is of the order of machine precision, as it is both in table 4 and section 7, the minimum can in fact be negative. The distance between two points is the square root of a sum of squares and because of this square root the number of Fourier points needed to represent the distance of a point from a periodic orbit as a Fourier series can be easily 50 times as many. The way to eliminate aliasing without paying too steep a price is to minimize the square of the distance and not the distance itself, and then take the square root of this minimum. The Fourier series to represent the square of the distance from (x_0, y_0, z_0) to a periodic orbit to machine precision may require at most twice as many points.

Next there is the problem of finding the distance between two nearby periodic orbits accurately. The distance of points of the first periodic orbit from the second periodic orbit can be represented as a Fourier series. In theory, it is enough to minimize this Fourier series. Again there is the problem of aliasing, and this time, there is no easy solution to it. We are saved, however, because we do not even need to do a continuous minimization. Both periodic orbits are solutions of the same differential equation and the divergence factor of trajectories over small intervals of time will be close to 1, even for the Lorenz equations (for a proof use the Lipschitz constant). So it is enough to find the distance of points of the first periodic orbit on the Fourier grid from the second periodic orbit and take the minimum of these numbers. This distance will not be accurate to machine precision, but it will be off by a factor that is quite close to 1.

4. Cantor leaves of the Lorenz attractor

The passage below from Lorenz's paper [10] has been quoted by Strogatz [14]. It is worth quoting once again.

It would seem, then, that the two surfaces merely appear to merge, and remain distinct surfaces. Following these surfaces along a path parallel to a trajectory, and circling around C and C' , we see that each surface is really a pair of surfaces, so that, where they appear to merge, there are really four surfaces. Continuing this process for another circuit, we see there are really eight surfaces, etc, and finally we conclude that there is an infinite complex of surfaces each extremely close to one or the other of the two merging surfaces.

To explain this passage, we return to figure 1. The points C and C' are the fixed points located inside the holes in the two wings. The z coordinate is 27 for both the fixed points. However, C has positive x and y coordinates, while C' has negative x and y coordinates. Trajectories of the Lorenz equations are directed upwards, with increasing z , when they intersect the sections A and B . They are directed downwards where the wings overlap. Lorenz based his argument on the apparent merging of the two wings.

Figure 4 shows a section of the Lorenz attractor. There is no evidence of an infinite complex of surfaces here. When trajectories starting at the intersection of the Lorenz attractor with the section A are followed, they intersect the Poincaré section once again at l_a (see figure 4). The




Figure 4. This plot shows a section of the Lorenz attractor at $z = 27$. The holes in the two wings can be seen in this plot. Of the four lines in the plot, the ones that point to the corners of the plot are on sections A and B defined in figure 1. When these two lines are integrated forwards in time, they trace out the top part of the two wings by first moving up, with z increasing, and then moving down. They end up on the Poincaré section $z = 27$ once again, but face each other this time as shown in the plot. We will denote the two lines facing each other by l_a and l_b .

observation behind Lorenz’s argument is that upon further integration, the trajectories spread out, and at their next intersection with the Poincaré section $z = 27$ seem to nearly completely cover the intersection of the attractor with *both* the sections A and B (see figure 1). The same thing is true if we follow trajectories from B . Therefore, the lines in figure 4 must be really two lines close to each other.

Figure 5 shows how the infinite complex of surfaces that Lorenz talked about is formed. The two plots were obtained using symbolic dynamics. The basic idea for obtaining the plots is contained in the passage quoted before, although perhaps it is more explicit in [8]. To see the line in figure 4 split into two, we need to separate the points that come from A from the points that come from B . If we consider periodic orbits, the starting points of the periodic orbits (these points will be located on either section A or section B) whose symbol sequences end with A , come from A . Those whose symbol sequences end with B , come from B . To zoom in on a little segment of the line, we need to fix the first several symbols. For the first of the two plots in figure 5, we fixed this at $ABABABABAB$, a sequence of length 10. The upper of the two lines in the first plots shows the starting points on section A of periodic orbits $ABABABABABxA$, where x varies over all the 128 sequences of length 7. The lower line is obtained from periodic orbits $ABABABABABxB$.

All the points in the upper line in the first plot in figure 5 land on A if the Lorenz equations are solved backwards in time. However, if we continue to integrate backwards some of them land on B and some of them land on A . These points must be separated to see that line split into two. We used periodic orbits of symbol sequences $ABABABABABABABABABABABABABABxA$ and $ABABABABABABABABABABABABABABxB$ to obtain the two lines in the second plot in figure 5. The leading symbol sequence used to zoom in on a segment of the upper line in the first plot must be an extension of $ABABABABAB$ with some more symbols. The extension we used has a total length of 24.




Figure 5. The two plots show the intersection of the Cantor leaves of the Lorenz attractor with the plane $z = 27$. The first plot is centred at $(a, b) = (-13.76451256298730, -19.57987258133632)$. It zooms in on a small region of the plot in figure 4, so that what appears to be a line there splits into two. The second plot is centred at $(a, b) = (-13.76431891885149, -19.57849004063768)$. It zooms on a small region on the upper line in the first plot. What appears to be a single line in the first plot is now seen to be really two lines. The text explains how the two plots were obtained.

Returning to section 2, the reason that a long forward fit is more important than a long backward fit in finding a match for an interval in a given symbol sequence is now clear. It takes 10 forward symbols to define the same scale as a single backward symbol; and 24 forward symbols to define the same scale as two backward symbols. It will take 38 forward symbols to define the same scale as three backward symbols.

It is clear from figures 4 and 5 that the intersection of the Lorenz attractor with a Poincaré section at $z = 27$ is a Cantor set of lines. The attractor of the flow must then be a Cantor set of leaves. In the geometric model, these leaves are pinched together at the unstable manifold of the origin. We will have more to say about this in the next section.

It is possible to estimate the factor by which the distance between the Cantor lines decreases every time the lines split. Lorenz [10] gave the estimate 7×10^{-5} . This estimate can be improved slightly. We will estimate the factor by which the distance between the two lines in the second plot of figure 5 is smaller than the distance between the two lines in the first plot. The second plot can be obtained by pushing the first plot forwards in time until it intersects the section A once again. The approximate time for a point on section A to return to section A is 0.72. The rate of volume contraction is $\frac{41}{3}$ and the rate of expansion on the unstable manifold is approximately 0.9. This implies that the rate of contraction along the stable manifold is approximately -14.6 . Therefore, the distances between the two lines must decrease approximately by the factor $e^{-14.6 \times 0.72} \approx 3 \times 10^{-5}$. Figure 5 shows that the actual contraction is in-between Lorenz's estimate and this estimate.

5. Periodic orbits close to the origin

The origin is a distinguished part of the Lorenz attractor. In this section, we compute periodic orbits close to the origin. We also begin to make more careful use of the geometric model to guide our computations. The propositions in this section and another proposition in section 7 assume the geometric model, partly explained in section 1. All the proofs follow standard symbolic dynamics arguments.

Proposition 5.1. *The geometric model possesses a periodic orbit with a certain symbol sequence if and only if the one-dimensional map f admits a periodic orbit with the same symbol sequence.*

Proof. Assume that the one-dimensional map admits a periodic orbit with a certain symbol sequence. This symbol sequence fixes the u coordinate of the first point of this periodic orbit uniquely. To find a periodic orbit of the geometric model, we need to determine a v coordinate for the given symbol sequence. This can always be done as the geometric model allows a complete, infinite binary tree of Cantor lines (transverse to lines with a fixed u) in the (u, v) plane.

In the geometric model, the u coordinate is mapped independently of the v coordinate. Therefore, the converse follows easily. \square

Proposition 5.1 allows us to look for periodic orbits by considering only the one-dimensional map f . Periodic orbits of the one-dimensional map that pass close to $u = \pm 1$ will pass close to the origin in the geometric model. The points $u = \pm 1$ correspond to the intersections of the two branches of the unstable manifold of the origin with the sections A and B (see figures 1 and 2). The point $u = 0$ corresponds to the stable manifold of the origin. The following proposition helps locate periodic orbits close to the origin.

Proposition 5.2. *Assume that -1 has the symbol sequence $A^m B^n A^o \dots$ under the one-dimensional map f . Then*

- *The periodic orbits $A^k B^k$ exist for $k < m$. If $m \neq n$, the periodic orbit $A^m B^m$ also exists.*
- *The periodic orbits get closer to ± 1 as k increases.*

Proof. Let $-1 < u_1 < 0$ be such that $f(u_1) = 0$. Define u_{i+1} , for $1 \leq i$, by $-1 < u_{i+1} < u_i$ and $f(u_{i+1}) = u_i$, as long as such a u_{i+1} can be found. If $-1 < u_i < 0$ for $1 \leq i < M$, but u_M cannot be found, define $u_M = -1$. Define $u_0 = 0$. Denote the half-open interval $[u_k, u_{k-1})$, $1 \leq k \leq M$, by I_k . Define the interval J_k to be $-I_k$ for $1 \leq k \leq M$.

A point in I_k has a symbol sequence beginning with k A s and then one or more B s. Since -1 has the symbol sequence $A^m B^n A^o \dots$, we can say that $M = m$.

Let $k < m$. Then $f^k(I_k)$ is the whole of $[0, 1]$ and monotone. There must be an interval $I'_k \subset I_k$ such that $f^k(I'_k) = J_k$. Since $f^k(J_k)$ is the whole of $[-1, 0]$, there must be $I''_k \subset I'_k$ such that $f^k(I''_k) \subset J_k$ and $f^{2k}(I''_k) = I_k$. In addition, f^{2k} is monotone over I''_k . Therefore, f^{2k} has a fixed point in I''_k . This fixed point is the first point of a periodic orbit of f with the symbol sequence $A^k B^k$.

The argument above breaks down for $k = m$ because we can no longer assert that $f^m(I_m)$ includes all of $[0, 1]$. In particular, J_m may not be fully contained in $f^m(I_m)$. However, if $n < m$, it is easy to argue that $J_m - \{1\} \subset f^m(I_m)$.

All the points of the periodic orbit $A^k B^k$ are contained in the interval $[u_k, -u_k]$ and the starting point is in I_k . Therefore the periodic orbits $A^k B^k$ get closer to -1 for increasing k . \square

Figure 6 illustrates proposition 5.2. The periodic orbit $A^{25} B^{25}$ was the last periodic orbit we were able to compute. It is necessary to add the periodic orbits $A^k B^k$ to the library described in section 2 as they are computed for increasing k . While almost all periodic orbits out of the 111011 periodic orbits with symbol sequences of length 20 or less took less than 15 iterations of the Lindstedt–Poincaré algorithm to compute, the periodic orbit $A^{23} B^{23}$ took a few hundred iterations, the periodic orbit $A^{24} B^{24}$ took nearly a thousand iterations, and the periodic orbit $A^{25} B^{25}$ took tens of thousands of iterations. The periodic orbit $A^{26} B^{26}$ may not exist. Our computation went into a local minimum of the residual error and was stuck with a relative residual error around $6.0e-4$. It is stated in [13] that the unstable manifold of the origin passes




Figure 6. The left-hand plot shows the periodic orbit $A^{25}B^{25}$. It is symmetric like the Lorenz attractor and gives a good idea of the full extent of the entire Lorenz attractor including the origin. However, the separation of trajectories that emanate from the section A into those that return to A and those that land on B is not as clear as in figure 1. The right-hand plot illustrates proposition 5.2.




Figure 7. Two plots of the periodic orbit $A^{23}B$. The right-hand plot is the only one in this paper to plot the z coordinate against the y coordinate.

through the section A 24 times before crossing over to the other side. Our computations show that this number is 25 or greater.

A symbolic dynamics argument assuming the geometric model implies that if $A^{25}B^{25}$ exists, so must the periodic orbit $A^{24}B$. We tried to test this implication by computing periodic orbits of the form $A^k B$. We were able to compute these periodic orbits up to $k = 23$ (see figure 7), but not with $k = 24$. Therefore, either the geometric model does not give a completely correct description of the Lorenz attractor or the Lindstedt–Poincaré algorithm fails to compute some periodic orbits that actually exist. We discuss both these possibilities after the next paragraph.

Figure 8 shows the periodic orbit $A^{25}B^7$ whose distance from the origin is 0.169. We were unable to compute $A^{25}B^6$. If that periodic orbit indeed fails to exist, then the symbol sequence of the negative branch of the origin could begin as $A^{25}B^6 A \dots$. However, it is not possible to draw this conclusion with certainty. Data for some of the periodic orbits that came up in this section is given in table 5.

The geometric model of the Lorenz attractor implies that the Cantor leaves of the Lorenz attractor, computed in section 4, are pinched in at the unstable manifold of the origin. Thus the Cantor lines on a Poincaré section at $z = 27$, such as the ones shown in figure 5, are




Figure 8. Among the periodic orbits we have computed, the periodic orbit $A^{25}B^7$ passes closest to the origin. The left-hand plot shows this periodic orbit. It is possible to extend proposition 5.2 and deduce the upward trend of the distances to the origin in the right-hand plot.

Table 5. Data for a few of the periodic orbits discussed in section 5. Comparison with table 2 shows that periodic orbits that pass close to the origin tend to have larger Lyapunov exponents.

s	λ	T
$A^{20}B$	0.813	1.462 753 180 496 2e+01
$A^{21}B$	0.842	1.535 545 749 548 1e+01
$A^{22}B$	0.886	1.611 597 203 895 0e+01
$A^{23}B$	0.958	1.694 510 320 552 6e+01
$A^{22}B^{22}$	0.842	3.071 091 470 802 8e+01
$A^{23}B^{23}$	0.886	3.223 194 408 118 1e+01
$A^{24}B^{24}$	0.958	3.389 020 642 303 8e+01
$A^{25}B^{25}$	1.13	3.604 395 870 791 7e+01
$A^{25}B^{10}$	1.04	2.516 150 035 094 0e+01
$A^{25}B^9$	1.05	2.449 012 425 167 2e+01
$A^{25}B^8$	1.06	2.381 812 609 353 6e+01
$A^{25}B^7$	1.08	2.314 577 553 153 6e+01

... tied in at a point that lies on the unstable manifold of the origin. It seems possible that the Cantor lines (or leaves) merge and disappear in a cascade, and not merge simultaneously as the model implies. Our computations provide no direct support to this possibility, however. This alternative picture has implications for symbolic dynamics (e.g. proposition 5.1 will fail) as well as the topology of the Lorenz attractor. These implications apply only in a very tiny region around the first intersections of the unstable manifold of the origin with the sections A and B .

The Lindstedt–Poincaré algorithm too becomes unreliable in this region. It appears possible, however, to combine the Lindstedt–Poincaré algorithm with a normal form computation near the origin, such as the one used in [15]. Such a method might be capable of determining the topology of the Lorenz attractor in a small neighbourhood of the origin. To focus this discussion, we ask the following question:

What are the first 250 symbols of the negative branch of the unstable manifold of the origin?

The extended precision computation necessary to answer this question will also shed light on two related matters. The first is the topology of the Lorenz attractor near the unstable manifold of the origin and the second is the question of exactly what symbol sequences are admitted by the Lorenz attractor.

6. Density of periodic orbits in the Lorenz attractor I

Table 6 shows 11 points close to the Lorenz attractor. They were generated using the 7th-order Fehlberg Runge–Kutta method with a constant time step. A random point was integrated for $t > 10\,000$ to generate the points shown in table 6. The time step used was $\frac{1}{128}$ for the first two points, $\frac{1}{512}$ for the next two, $\frac{1}{4192}$ for the ones numbered 7 and 8, and $\frac{1}{2096}$ for the others. All these points are a little below the section A.

We wish to find periodic orbits that pass as close to these points as possible. The method used in this section makes no attempt to be clever. A better method, derived using the geometric model in the next section, shows that all these points are as close to the Lorenz attractor as machine precision permits. The simpler method of this section is not so successful, however. We present it nevertheless because it may apply in situations where a detailed geometric model is not available.

Table 7 shows the periodic orbits computed using this simpler method and their distances from the 11 points. Although they do pass quite close to the points, the distances are not of the order of machine precision.

The method works as follows. Given a point, it prepares a list of symbol sequences of length from 2 to 32 such that each symbol sequence is the closest to the given point among

Table 6. This table shows points close to the Lorenz attractor. They are numbered from $n = 1$ till $n = 11$. The text explains the method used for generating these points.

n	x	y	z
1	-1.479 351 263 303 5e+01	-2.265 460 975 160 8e+01	2.601 751 639 387 0e+01
2	-1.179 336 677 270 8e+01	-1.567 815 010 226 7e+01	2.634 734 133 982 2e+01
3	-1.336 188 799 485 0e+01	-1.878 134 776 736 4e+01	2.683 377 564 595 5e+01
4	-1.030 743 945 264 6e+01	-1.218 022 416 084 0e+01	2.690 220 173 720 9e+01
5	-1.321 633 041 880 2e+01	-1.838 169 888 555 2e+01	2.693 709 938 314 3e+01
6	-1.356 230 039 378 3e+01	-1.913 466 467 849 7e+01	2.697 655 815 165 9e+01
7	-1.396 763 919 801 8e+01	-2.007 068 580 792 8e+01	2.696 531 491 804 2e+01
8	-1.211 990 132 847 3e+01	-1.594 589 290 517 6e+01	2.699 658 557 704 7e+01
9	-1.231 125 028 898 3e+01	-1.637 063 704 277 9e+01	2.697 740 231 982 3e+01
10	-1.324 777 611 718 5e+01	-1.843 802 870 444 0e+01	2.696 397 233 116 9e+01
11	-1.426 039 058 036 1e+01	-2.079 380 159 649 7e+01	2.689 994 450 658 4e+01

Table 7. The first column refers to the numbering of the points in table 6, and s is a symbol sequence of length 32. The last column gives the distances of the point numbered n from the periodic orbit with symbol sequence s .

n	s	Distance
1	ABBBBABBBBBBBAABBBBAAAAAAAAAAAA	7.04e-07
2	AAABAAABAAAABABBAAAAAAAAABAAAAAAAA	5.37e-07
3	ABAAABABAAAABBABAABBABBAAAAAAAAAAAA	1.85e-07
4	AAAAAAAAAAABAAAABABAABBAAAAAAAAAAAA	2.01e-05
5	AABBBBBBBBBBAAAABBABBABBAAAAAAAAAAAA	1.59e-04
6	ABAABBAABBAAAAAAAAAAAAAAAAAAAAAAAA	2.01e-06
7	ABABBBBABBABABAABAABAABBBBBBBBBB	7.62e-06
8	AAABABABAAAABBABBABBABBABBABBABB	1.14e-06
9	AAABBBAAAABAABAABAABAABBBBBBBBBB	1.10e-04
10	ABAAAAAAAAABBBBAAAABAABAABBBBBBBBBB	8.85e-05
11	ABBABAAAABAAAAAAAAABBBAAABBBBBBBBB	1.64e-06

a large number of symbol sequences of the same length. To begin with, it makes one pass through the library described in section 2 and finds symbol sequences of length 2–13 that are the closest to the given point among symbol sequences of the same length. These 12 sequences are entered into the list. The 12 symbol sequences for the point numbered 7 in table 6 and their distances from that point are shown below.

```

AB(2) dist:0.533595
ABB(3) dist:1.12959
ABBA(4) dist:0.614564
ABABB(5) dist:0.248908
ABABBB(6) dist:0.0982862
ABABBBB(7) dist:0.0119864
ABABBBBA(8) dist:0.0355487
ABABBBBAB(9) dist:0.015166
ABABBBBAB(10) dist:0.00513484
ABABBBBAB(11) dist:0.000335084
ABABBBBAB(12) dist:0.00122588
ABABBBBAB(13) dist:0.000131703
    
```

Above and in table 7, the given point is closest to the segment of the periodic orbit after the last symbol, but before the first symbol. To find a symbol sequence of length n , $n > 13$, the method starts with the symbol sequences of shorter length entered into the list. It inserts a single symbol in one of the last six positions of the previous symbol sequence of length $n - 1$. It then adds two symbols together in one of the last six positions of the one before the previous symbol sequence. Finally, it adds three symbols together to the symbol sequence of length $n - 3$ in the list. Of all the possible strings obtained in this way, the symbol sequence closest to the given point is entered into the list as the sequence of length n . Table 7 shows that this simple method does reasonably well. The list can be continued to sequences longer than 32, but these longer sequences do not get appreciably closer to the points in table 6 than the symbol sequences of length 32 shown in table 7.

7. Density of periodic orbits in the Lorenz attractor II

We will devise another method to find periodic orbits close to a given point on the Lorenz attractor in this section. Table 8 shows that periodic orbits computed by this method come within a distance of about $1e-12$ in all 11 cases. Since points near the section A have two digits

Table 8. The last column gives the distance from the periodic orbit with the symbol sequence s to the point numbered n . Data for the points is given in tables 6 and 9.

n	s	Distance
1	ABBBBABABBBBBBAABABBABBBBBBABAABBBBBBBAABAA	3.34e-12
2	AAABAAAABAAAABABBAAAAAAABBBAAABBBABBBABBBBBBAAA	8.71e-13
3	ABAAABABAAAABBABAABBABBAABBABABABBABAABAA	4.95e-12
4	AAAAAAAAAABAAAABABAABBAABBAABAAABABBBBAAAABBBABBBABA	5.62e-13
5	AABBBBBBBBBBAABBBABAABBBABBABBBBABBAABBAABAABAAA	8.97e-13
6	ABAABBAABBBABBBBBBBBBBBBABBBBBBABABAAAAAABBAAAAAABAA	2.06e-13
7	ABABBBBABBBABABAABAAAAAABBBAAABAABABB	1.30e-12
8	AAABABABAAAABBABBBABBBAAABBBABBBAAAAAABBBAAABBBBAB	5.74e-13
9	AAABBBAAAABAAAABBBBBBBBBBBABABAABBBBAABABABBAB	1.98e-12
10	ABAAAAAAAABBBBAAAABABBAABAAAAABAABBAABABABBAAB	9.26e-13
11	ABBABAAAABAAAAAABABBAABAABABBBAAABBBBBBBBABBB	3.08e-13

before the decimal point, 14 digits of the randomly generated points on the Lorenz attractor coincide with some point on the periodic orbits. In effect, we have a method here that merely changes the last digit or two of a given point on the Lorenz attractor and then finds a periodic orbit that passes right through the perturbed point.

The proposition below is easily deduced from the geometric model. It forms the basis for a method to find periodic orbits close to a given point on the Lorenz attractor.

Proposition 7.1. *Let $-1 < p < 1$. Consider the one-dimensional map f on $[-1, 1]$. Let $d(p, \alpha)$ denote the distance of that point from the first point of the periodic orbit of symbol sequence α . Assume that the symbol sequence of p begins with α_p and that α_p ends with a block of A s. Let*

$$d(k) = d(p, \alpha_p B^k A).$$

Let m be as in proposition 5.2. Then

- As k takes the values $1, 2, 3, \dots, m$ successively, $d(k)$ begins to decrease and then increases.
- If $d(k)$ is minimum for $k = k_0$, the block of B s that follows α_p in the symbol sequence of p is of length $k_0 - 1, k_0$, or $k_0 + 1$.

The statements above are still true if A and B are interchanged.

Proof. The proof uses notation defined in the proof of proposition 5.2. Let l be the number of symbols in α_p . By assumption, the symbol following α_p in the symbol sequence of p must be B . Denote the periodic point whose symbol sequence repeats $\alpha_p B^k A$ by p_k .

The point $f^l(p)$ must be positive. If the number of B s that follow α_p in the symbol sequence of p is r , $f^l(p)$ will lie in the interval J_r . Similarly, $f^l(p_k)$ lies in the interval J_k . Therefore, $d(f^l(p), f^l(p_k))$ is minimum for $k = r - 1$, or $k = r$, or $k = r + 1$. Furthermore, this quantity first decreases and then increases as k takes the values $1, 2, \dots, m$. The map f^l is monotone over the interval of points whose symbol sequences begin with α_p . Therefore, the statements made about $d(k)$ in the proposition are true.

To obtain a proof for the proposition with A and B interchanged, interchange A and B in the proof and replace J by I . □

We will assume the point on the Lorenz attractor to be on the section A . Table 9 shows the points obtained by advancing the points in table 8 till they lie on section A . The proposition

Table 9. This table gives the x and y coordinates of points shown in table 6 after advancing them to the section A . The last column gives the Lyapunov exponents of periodic orbits shown in table 8. Their distance from the corresponding points is of the order of machine precision in every case.

n	x	y	λ
1	-1.507 688 334 028 0e+01	-2.265 317 531 783 1e+01	0.906
2	-1.200 699 650 035 0e+01	-1.567 871 724 314 0e+01	0.888
3	-1.341 179 249 534 0e+01	-1.877 738 672 920 9e+01	0.957
4	-1.034 111 146 724 7e+01	-1.217 777 331 046 8e+01	0.899
5	-1.323 527 218 244 6e+01	-1.837 995 753 542 3e+01	0.912
6	-1.356 925 853 612 8e+01	-1.913 398 873 181 7e+01	0.867
7	-1.397 778 360 200 3e+01	-2.006 971 168 098 6e+01	0.926
8	-1.212 097 834 492 1e+01	-1.594 578 578 610 2e+01	0.931
9	-1.231 832 021 112 9e+01	-1.636 995 430 629 1e+01	0.905
10	-1.325 860 991 718 6e+01	-1.843 699 521 219 5e+01	0.917
11	-1.428 938 112 515 1e+01	-2.079 122 015 447 7e+01	0.902

above leads to a method that can find the u coordinate of a given point p in the geometric model. However, since we wish to carry our computations till the periodic orbit is as close to the given point as machine precision permits, it is also necessary to do something about the v coordinate. The v coordinate of the beginning point of a periodic orbit is determined by the last few symbols in its symbol sequence. Computations in section 4 imply that the last two or three symbols are sufficient to determine the v coordinate. Since the number of possibilities is not too many, we do not attempt to derive a systematic method to find the last few symbols. We simply use a guess for the last few symbols.

In the description of the method below, we assume that the last two or three symbols are known to be β . In practice, β is guessed by running the method two or three times. At an intermediate stage in the method, there will be up to five guesses, α_i , $1 \leq i \leq 5$, for the starting symbol sequence of the given point p . To pass to the next stage, the first step is to trifurcate, following proposition 7.1, from each of the α_i . The trifurcation proceeds as follows. Assume that α_1 ends with a B . Then we compute $d(k) = d(p, \alpha_1 A^k B \beta)$, where $d(p, s)$ is the distance from the point p to the periodic orbit of symbol sequence s , for $k = 1, 2, 3, \dots$. This first decreases and then increases. We stop the computation of distances as soon as the distances begin to increase, and record the value of k where $d(k)$ is minimum as k_0 . It is possible that k_0 is equal to 1. The symbol sequence α_1 trifurcates into the symbol sequences $\alpha_1 A^{k_0-1}$, $\alpha_1 A^{k_0}$, and $\alpha_1 A^{k_0+1}$, if $k_0 > 1$. If $k_0 = 1$, it bifurcates into the latter two sequences. After all the five sequences either trifurcate or bifurcate, we will be left with at most 15 sequences. Of these, we retain only the five sequences that are the closest to the given point p . The last few symbols must be fixed at β for computing these distances. The five remaining sequences are used to begin the next stage.

For the point numbered 7 in table 9, the beginning collection of α_i is:

```
A dist:0.249867
AA dist:2.98898
```

At the 7th stage, we are left with the following collection:

```
ABABBBBABBBBA dist:0.000220551
ABABBBBABBBAA dist:0.00120674
ABABBBBABBA dist:0.00535478
ABABBBBABBA dist:0.00790187
```

At the 17th and final stage, we have the following collection:

```
ABABBBBABBBABABABAAAAAABBBAABAA dist:1.30158e-12
ABABBBBABBBABABABAAAAAABBBAABAAA dist:7.89513e-10
ABABBBBABBBABABABAAAAAABBBAABA dist:1.49464e-09
```

The computation for this point was run with $\beta = ABB$.

Computation of each of the periodic orbits listed in table 8 using this method involved the computation of a few hundred periodic orbits in the intermediate stages. Table 9 shows the Lyapunov exponents of the periodic orbits that pass as closely as machine precision allows to 11 randomly generated points. The wide variation of Lyapunov exponents in table 9 indicates the difficulty faced by numerical methods that attempt to compute the Lyapunov exponent by direct numerical integration.

8. Conclusion

Lorenz [10] used numerical integrations to support his argument that bounded, nonperiodic flows are necessarily unstable, and to draw a series of conclusions about the geometry of the Lorenz attractor. His integration used a time step of 0.01 and a one step numerical method

(the use of 64 Fourier points per symbol implies that the time step in our computations can be longer). We have obtained plots of the fractal structure of the Lorenz attractor in section 4 and derived a method to compute periodic orbits arbitrarily close to a given point on the Lorenz attractor in sections 6 and 7. From a computational point of view, we have demonstrated the symbolic dynamics approach to the numerical investigation of chaos. Even for the widely studied Lorenz attractor, we have shown that this approach leads to far greater precision and level of descriptive detail than any other approach.

Acknowledgments

The author thanks Professor Nick Trefethen for encouragement and for his comments about the paper.

Appendix

This appendix discusses an extension of the Lindstedt–Poincaré algorithm for computing periodic orbits presented in [16]. The computational problem is to find periodic solutions of $dz/dt = f(z)$, $z \in R^d$, starting with an initial guess $z(t) = x(\omega t)$, where $x(\tau)$ is a 2π periodic function of τ and ω is a guess for the frequency. The Lindstedt–Poincaré algorithm proceeds by solving the linear equation

$$\omega \dot{y}(\tau) = A(\tau)y + r(\tau) - \delta\omega \dot{x}(\tau),$$

where $A(\tau) = \partial f/\partial x|_{x=x(\tau)}$ and $r(\tau) = f(x(\tau)) - \omega \dot{x}(\tau)$ is the residual error in the starting guess. The solution $y(\tau)$ is the correction to $x(\tau)$ and $\delta\omega$ is the correction to ω . It is different from other Newton methods mainly in the systematic use of Fourier series to represent $x(\tau)$ and also in the use of the Lindstedt–Poincaré technique to compute $\delta\omega$. We make some points about the method below.

- (1) The numerical computation of any mathematical object can be only as accurate as its representation permits. The representation of periodic orbits by Fourier series is both accurate and efficient because, for analytic vector fields, the Fourier coefficients decrease exponentially. Moreover the residual error $r(\tau)$ can always be computed accurately, even when the guess itself is inaccurate, as its computation requires only the fast Fourier transform and the evaluation of the vector field at points on the equispaced Fourier grid. It might appear, however, that the Fourier representation of $f(x(\tau))$ could require many more points than the Fourier representation of $x(\tau)$ because f is nonlinear; for example, the Fourier representation of $\sqrt{3 + \cos \tau}$ up to machine precision requires many more points than the Fourier representation of $\cos \tau$. Such a need for padding Fourier series to eliminate aliasing error never arises in the Lindstedt–Poincaré algorithm because $x(\tau)$ is not just any Fourier series but a representation of a Fourier series close to a solution of the system $\dot{x} = f(x)$. Even though $f(x)$ is nonlinear, it is equal to \dot{x} which is a linear function of x . Therefore, there must be a *conspiracy* between the nonlinear vector field f and the solution $x(\tau)$ that completely eliminates any need to pad Fourier series while computing the residual error $r(\tau)$.
- (2) Every step of the method attempts to eliminate the residual error $r(\tau)$ by computing corrections $y(\tau)$ and $\delta\omega$. This residual error cannot be completely eliminated in a single step because the correction equation is obtained by keeping only the linear part. Further, it is not necessary to solve the correction equation with any great accuracy

even when the final solution is desired to be as accurate as machine precision allows. Computing the correction with only three digits of accuracy, e.g., is enough to improve the guess by three digits of accuracy, and upon successive iterations, the guess will become more and more accurate. The key is to compute the guess $r(\tau)$ accurately and it is here that the Fourier representation proves essential.

- (3) There are some auxiliary linear ODEs that arise in the solution of the correction equation [16]. It is convenient to take the step size of the Runge–Kutta method used for solving those ODEs to be equal to the spacing between neighbouring points on the Fourier grid. Data needed for the intermediate stages of the Runge–Kutta methods can be easily obtained by shifting Fourier series. The ability of the methods derived in [1] to compute oscillations, which is due to the use of the matrix exponential by those methods, makes them a good choice in the Lindstedt–Poincaré algorithm.
- (4) The Lindstedt–Poincaré algorithm presented in [16] uses single shooting to solve the correction equation. A multiple shooting version requires only a slight modification. As even a coarse solution of the correction equation suffices as explained above, even highly unstable periodic orbits can be computed with only a few shooting intervals. For example, a periodic orbit with leading characteristic multiplier of 10^{60} can be computed using only five or six shooting intervals.

A particularly felicitous way to improve the convergence of Newton iterations is the dogleg trust region method discovered by M J D Powell (see [4]). The only ingredient necessary to incorporate the dogleg method into the Lindstedt–Poincaré iteration is a formula for the gradient of the residual. A simple variational calculation shows the gradient along $x(\tau)$ to be

$$2(A^T(\tau)r(\tau) + \omega\dot{r}(\tau)),$$

and the gradient along ω to be

$$-2 \int_0^{2\pi} \langle r(\tau), \dot{x}(\tau) \rangle d\tau.$$

For the computations presented in the preceding sections, the above formulae were used to modify the Lindstedt–Poincaré method.

References

- [1] Blanes S, Casas F and Ros J 2000 Improved high order integrators based on Magnus expansion *BIT* **40** 434–50
- [2] Coomes B A, Kocak H and Palmer K J 1997 Long period shadowing *Numer. Algorithms* **14** 55–78
- [3] Cvitanović P, Artuso R, Dahlquist P, Mainieri R, Tanner G, Vattay G, Wheelan N and Wirzba A 2003 *Classical and Quantum Chaos* www.nbi.dk/ChaosBook/
- [4] Dennis J E Jr and Schnabel R B 1996 *Numerical Methods for Unconstrained Optimization and Nonlinear Equations* (Philadelphia, PA: SIAM)
- [5] Eckhardt B and Ott G 1994 Periodic orbit analysis of the Lorenz attractor *Z. Phys. B* **93** 259–66
- [6] Franceschini V, Gilberti C and Zheng Z 1993 Characterization of the Lorenz attractor by unstable periodic orbits *Nonlinearity* **6** 251–8
- [7] Guckenheimer J 1979 A strange, strange attractor *The Hopf Bifurcation and its Applications* ed J Marsden and M McCracken (New York: Springer)
- [8] Guckenheimer J and Holmes P 1996 *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields* (New York: Springer)
- [9] Kaplan J L and Yorke J A 1979 Preturbulence, a regime observed in a fluid flow model of Lorenz *Commun. Math. Phys.* **67** 93–108
- [10] Lorenz E N 1963 Deterministic non-periodic flows *J. Atmos. Sci.* **20** 130–41
- [11] Mischaikow K, Mrozek M and Szymczak A 2001 Chaos in the Lorenz equations: a computer assisted proof III, classical parameter values *J. Diff. Eqns* **169** 17–56

-
- [12] Smale S 1967 Differentiable dynamical systems *Bull. Am. Math. Soc.* **73** 747–817
 - [13] Sparrow C 1982 *The Lorenz Equations: Bifurcations, Chaos, and Strange Attractors* (New York: Springer)
 - [14] Strogatz S H 1994 *Nonlinear Dynamics and Chaos* (Reading, MA: Addison-Wesley)
 - [15] Tucker W 2002 A rigorous ODE solver and Smale’s 14th problem *Found. Comput. Math.* **2** 53–117
 - [16] Viswanath D 2001 Lindstedt–Poincaré technique as an algorithm for computing periodic orbits *SIAM Rev.* **43** 478–96
 - [17] Williams R F 1980 Structure of Lorenz attractors *Publications Mathématiques, IHES* **50** 73–100