Available online at www.sciencedirect.com
SCIENCE DIRECT®
Physica D 190 (2004) 115–128
PHYSICA D
www.elsevier.com/locate/physd

The fractal property of the Lorenz attractor

Divakar Viswanath

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

Received 30 September 2002; received in revised form 5 May 2003; accepted 14 October 2003
Communicated by C.K.R.T. Jones

Abstract

In a 1963 paper, Lorenz inferred that the Lorenz attractor must be an infinite complex of surfaces. We investigate this fractal property of the Lorenz attractor in two ways. Firstly, we obtain explicit plots of the fractal structure of the Lorenz attractor using symbolic dynamics and multiple precision computations of periodic orbits. The method we derive for multiple precision computation is based on iterative refinement and can compute even highly unstable periodic orbits with long symbol sequences with as many as 100 digits of accuracy. Ordinary numerical integrations are much too crude to show even the coarsest splitting of surfaces, and there appear to be no other explicit computations of the fractal structure in the extensive literature about the Lorenz attractor. Secondly, we apply a well known formula that gives the Hausdorff dimension of the Lorenz attractor in terms of the characteristic multipliers of its unstable periodic orbits. The formula converges impressively and the Hausdorff dimension of the Lorenz attractor appears to be 2.0627160. We use comparison with explicit computations of the fractal structure and discuss the accuracy of this formula and its applicability to the Lorenz equations. Additionally, we apply periodic orbit theory to the Lorenz attractor and exhibit its spectral determinant and compute its Lyapunov exponent.

Keywords: Symbolic dynamics; Hausdorff dimension; Lorenz equations

1. Introduction

The systematic study of the differential equations:

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

with $\sigma = 10$, $r = 28$, and $b = 8/3$, by Lorenz [10] led to the discovery of the butterfly-like Lorenz attractor, an image that has become commonplace in textbooks on chaos theory. The path that led Lorenz to these equations began with an effort to find a simple model problem for which the methods used for statistical weather forecasting would fail. The Lorenz equations are also connected to other physical phenomena [15].

In Fig. 1, we have plotted a single periodic orbit of the Lorenz equations. By an ingenious argument, Lorenz inferred that although the Lorenz attractor appears to be a single surface, it must really be an infinite complex of surfaces; in other words, the Lorenz butterfly must be a fractal. Lorenz's paper has been as influential as it is insightful, yet there has not been a single computation of the splitting of the Lorenz attractor into a complex of

E-mail address: divakar@umich.edu (D. Viswanath).

0167-2789/$ – see front matter © 2003 Elsevier B.V. All rights reserved.
doi:10.1016/j.physd.2003.10.006

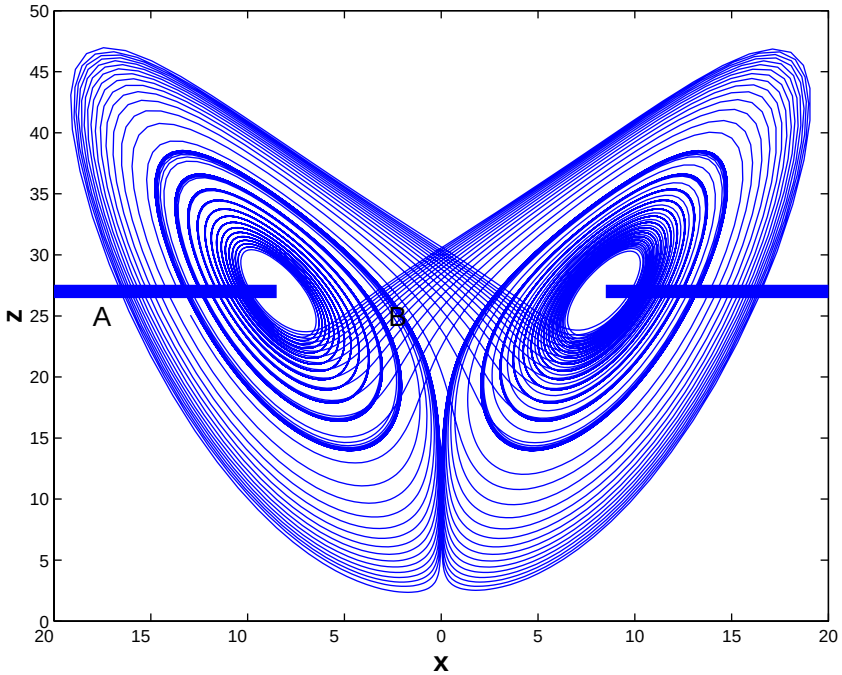

Fig. 1. This plot shows a single periodic orbit of the Lorenz equations. The crossing of trajectories to different wings of the butterfly and the apparent merging of the two surfaces are both evident in this plot. The symbols A and B are assigned to every intersection of a trajectory with the two sections shown in the figure. The symbol sequence that corresponds to the periodic orbit shown above is $\mathrm{ABA}^2\mathrm{B}^2\cdots\mathrm{A}^{15}\mathrm{B}^{15}$. Even though this periodic orbit is highly unstable with a characteristic multiplier of 3.06e+59, the method described in Section 2 can compute starting data for this periodic orbit with 100 digits of accuracy.

surfaces. This is probably because the splitting of the surface of the Lorenz attractor cannot be computed by direct numerical integration. If a trajectory is integrated for even a long period of time and its intersection with the Poincaré section $z = 27$ is plotted, the intersection appears to be a one-dimensional object—a simple curved line and its symmetric image. Lorenz [11, p. 140] has said that he initially concluded that the two wings do indeed merge and that the curved line is what it appears to be—a one-dimensional object. This initial conclusion, however, could not be reconciled with the mathematical fact that distinct trajectories of the Lorenz equations cannot merge. For this conflict to go away, the Lorenz attractor has to be an infinite complex of surfaces.

A description of the first split in the surface can be used to pinpoint the reason why plots of the fractal structure of the Lorenz attractor cannot be obtained by direct numerical integration. It is obvious from Fig. 1 that trajectories that start off on section A can have their next landings on either section A or B. In fact, the points where these landings occur appear to span the intersection of the entire Lorenz attractor with both sections A and B. The same is true for trajectories which start off from section B. If we consider the intersection of the Lorenz attractor with the section A, the points that emanate from A must be on a line that is distinct from the points that emanate from B. However, a simple estimate, such as that given by Lorenz [10], shows that the distance between these two lines will about 3×10^{-5}. This distance is about five orders of magnitude smaller than the extent of the Lorenz attractor. Therefore, for the separation of surfaces to be visible, a plot of the intersection of the Lorenz attractor with the Poincaré section $z = 27$ has to zoom in on a very tiny region of the section. Further, the direct numerical integration has to be carried on for a very long time if we are to find sufficiently many points in that tiny region, and the integration has to be accurate for the plot to be meaningful. The surface of the Lorenz attractor splits first into two

surfaces, as we just explained, and those two surfaces split again into four, and those four split into eight, and so on. The plot will have to zoom in on ever tinier regions to make these further splits visible. Computing even this first split is probably outside the reach of direct numerical integration. For plots of further splits, not only do the direct numerical integrations have to be much longer, they also have to be much more accurate. Plots of the fractal structure of the Lorenz attractor shown in Fig. 2 are impossible to obtain using direct numerical integration.

It is necessary then to devise a method for finding points of the Lorenz attractor in even very tiny regions of either section A or section B shown in Fig. 1. We describe how to find such points in Section 3. That method relies on the ability to compute periodic orbits of the Lorenz equations with specified symbol sequences. It is possible to associate a symbol sequence with any periodic orbit of the Lorenz equations in the following way. Assume that the starting point of the periodic orbit is on either section A or section B. The first symbol is either A or B depending upon which section the periodic orbit starts from. As the periodic orbit is traversed, it lands on the sections A or B in some order. A symbol is appended for every landing till the periodic orbit returns to its starting point. The symbol sequence assigned to a periodic orbit depends upon its starting point, and we will always assume that the starting point is on either section A or section B.

Fig. 1 shows a periodic orbit of the Lorenz attractor that was computed to have a specified symbol sequence of length 240. We used methods derived in [17,18] to compute the periodic orbit shown in Fig. 1. Those methods compute periodic orbits with accuracy close to the best possible in IEEE double precision arithmetic. Thus the starting point of a periodic orbit of specified symbol sequence can be located with about 14 or 15 digits of accuracy. Unfortunately, even this level of accuracy is insufficient to obtain detailed plots of the fractal structure of the Lorenz attractor. In Section 2, we describe a numerical method for refining a computation of a periodic orbit to increase its accuracy. This numerical method uses multiple precision arithmetic, as it must, but it limits its usage of multiple precision arithmetic operations to only a fraction of the total number of arithmetic operations necessary at each step of refinement. Thus even though the cost of multiple precision arithmetic operations is great, the computations remain manageable. In Section 2, we report a few periodic orbits computed with 100 digits of accuracy. With this numerical method, it is possible to compute the fractal structure of the Lorenz attractor in detail as shown in Section 3. For fractals like the Mandelbrot set, it is easy to zoom in on successively smaller regions and exhibit the self similarity property. The Lorenz attractor too has the same self similarity property as Lorenz inferred, and our Fig. 2 shows plots that convincingly exhibit the fractal property of the Lorenz attractor.

Although numerical computations are widely used in investigations of dynamics, they are not always reliable. As discussed above, typical numerical computations are incapable of revealing the fractal structure of the Lorenz attractor. More generally, it can be difficult to assert the existence of chaos with confidence based on numerical computations. Motivated partly by a wish to remedy this deficiency, methods to perform rigorous computations of dynamics have been developed [12,16]. Indeed the Lorenz equations have been a favored target for these methods [13,16]. While the computation of mathematical objects such as the Conley index or the invariant cone field by these methods uses machine arithmetic, the errors are rigorously bounded using interval arithmetic and other techniques, so that, in the end, it is possible to assert theorems based on the computation. Methods based on the Conley index, in particular, have already been applied to a variety of problems [12]. Although the methods developed here and in [17,18] lead to highly accurate computations, they do not lead to a proof that the computed objects, such as the one displayed in Fig. 1, actually exist. The type of rigor that these computations aim for is not that of the mathematical proof, but the equally exacting notion of rigor found in the experimental sciences. When the velocity of light is experimentally found to be a certain number, the validity of that number relies on two principles. Firstly, it must be possible to reproduce the experiment. In practice, however, reproduction of experiments can require skill, and this is certainly true not only for the determination of the velocity of light but also for the computations reported here. The computations reported here rely on the methods developed in [17,18] in addition to the methods developed in this paper, and altogether, the code is nearly several thousand lines of C++. The second—and far more

important—principle of experimentation is that the experiments can be either verified or refuted using experiments based on independent ideas. Ultimately the validity of an experiment to determine the velocity of light depends entirely on the digits in the answer that is found, because each of these digits can be picked apart by other experiments and is open to refutation. With this principle in mind, we have reported computations in great detail, and with many digits of precision, in this paper and in [17,18]. Considerations that lead us to believe that the numbers in this paper are as accurate as we report them to be are given in Section 2. A convergence theorem proved in Section 2 is a part of those considerations.

There are formulas available for the Hausdorff dimension and other fractal dimensions of chaotic sets like the Lorenz attractor [2,6]. These formulas can be applied to derive estimates of the Hausdorff dimension of the Lorenz attractor when periodic orbits corresponding to all symbol sequences of length about 10 or less are known [4,5]. It is necessary to compute periodic orbits with symbol sequences of length 16 or more to obtain a satisfactory plot of even the first split in the surface of the Lorenz attractor. Therefore, it may be puzzling why the formulas can be applied using only periodic orbits with symbol sequences of length around 10. Unlike a direct computation of the fractal structure, the formulas use information from the local linearization around the periodic orbits. Thus estimates of the Hausdorff dimension derived from the formulas may be expected to be better than estimates derived from a direct computation of the fractal structure of the Lorenz attractor. In Section 4, we use data computed in [18] and apply the formula to estimate the Hausdorff dimension of the Lorenz attractor using all 111 011 periodic orbits of the Lorenz equations with symbol sequences of length 20 or less. We find that the impressive convergence reported in [5] continues. We compare this computation with the explicit computation of the fractal structure in Section 3 and discuss the accuracy of the formula and its applicability to the Lorenz attractor. Section 5 applies periodic orbit theory to the Lorenz attractor.

2. Numerical method

The methods described in [17,18] can compute the starting point of a periodic orbit of the Lorenz equations of specified symbol sequence with 14 or 15 digits of accuracy. The entire periodic orbit is represented as a Fourier series. The numerical method described here uses multiple precision arithmetic and is based on iterative refinement [8,14]. Its input is a periodic orbit represented as a Fourier series with 14 or 15 accurate digits. It refines the input and outputs Fourier series for the same periodic orbit with as many as 100 accurate digits. The output Fourier series may need to be shifted in the time domain so that the starting point of the given symbol sequence occurs at the first point on the Fourier grid. The plots of the fractal structure of the Lorenz attractor given in Section 3 use starting points of periodic orbits computed with 40 accurate digits.

The Newton iteration described in [17] starts with a guess $x_0(\tau) \in R^d$, which is 2π periodic, and a frequency ω_0. It tries to improve this guess, and obtain a more accurate periodic solution of the differential equation $\mathrm{d}x/\mathrm{d}t = f(x)$, with $x,\ f(x) \in R^d$, by solving the linear system:

$$\omega_0 \dot{y}(\tau) = A(\tau)y + r(\tau) - \delta\omega \dot{x}_0(\tau) \tag{2.1}$$

for y and $\delta\omega$. The residual $r(\tau)$ is given by $r(\tau) = f(x_0(\tau)) - \omega_0 \dot{x}_0(\tau)$ and $A(\tau) = (\partial f/\partial x)|_{x=x_0(\tau)}$. The use of the Lindstedt–Poincaré technique to solve (2.1) is described in [17]. Relative residual error is obtained by dividing the usual Hilbert space norm of the 2π periodic function $r(\tau)$ by the norm of $f(x_0(\tau))$.

Iterative refinement is combined with the Newton iteration as explained below. For concreteness, we will assume that the final output must be a periodic orbit accurate to 100 digits:

1. Compute a periodic orbit represented as a 2π periodic function $x_0(\tau)$ and a frequency ω_0 to about 14 digits of accuracy. Compute $A(\tau)$ and $\dot{x}_0(\tau)$ for repeated use in step 5. Even though this guess has only about 14 digits of accuracy, the number of Fourier points used to represent it must be sufficient for 100 digits of accuracy.

2. Represent $x_0(\tau)$ and ω_0 using numbers with 100 digit mantissas.
3. Compute the residual $r(\tau)$ using 100 digit arithmetic.
4. Convert the residual $r(\tau)$ to IEEE double precision.
5. Solve (2.1) in IEEE double precision for y and $\delta\omega$.
6. Perform the updates $x_1 = x_0 + y$ and $\omega_1 = \omega_0 + \delta\omega$ using 100 digit arithmetic.
7. Return to step 3, replace x_0 and ω_0 by x_1 and ω_1, and repeat till the relative residual error is smaller that 10^{-100}.

The iterative refinement described above fails unless the guess $x_0(\tau)$ is smoothed at the beginning or at the end of every iteration. We used the software package MPFUN++ for multiple precision arithmetic. The FFT supplied with that package was used to compute the residual and to smooth the guess.

The effect of freezing $A(\tau)$ and $\dot{x}_0(\tau)$ at the beginning is that Eq. (2.1) cannot be solved exactly; it is not even formed exactly. This is not a problem because in any event (2.1) is solved only with a relative error of about 10^{-5}. We prove the convergence of the iterative refinement procedure in two steps. Theorem 2.1 is a modified version of a theorem given in [17]. The modification supplies different estimates for the constants C_1, C_2, and C_3. The earlier estimates are wildly unrealistic for periodic orbits with characteristic multipliers greater than 10^{10} such as the periodic orbit displayed in Fig. 1. Theorem 2.1 assumes that the correction equation (2.1) is solved exactly, which is not the case in the iterative refinement procedure described above. Theorem 2.2 only assumes that the correction equation (2.1) is solved with a certain relative accuracy as in the iterative refinement procedure described above.

Theorem 2.1. *Assume that $\bar{\omega}$ and $\bar{x}(\tau)$ give a hyperbolic 2π periodic orbit of $\omega\dot{x}(\tau) = f(x)$. They also give a $2\pi/\bar{\omega}$ periodic orbit of $\dot{x}(t) = f(x)$. Assume $x \in R^d$ and that $f(x)$ is twice continuously differentiable. Let ω_0 and $x_0(\tau)$, which must be 2π periodic, be approximations to $\bar{\omega}$ and $\bar{x}(\tau)$. Let $\delta\bar{\omega} = \bar{\omega} - \omega_0$ and $\bar{y} = \bar{x} - x_0$. Assume*

$$|\delta\bar{\omega}| < \epsilon, \qquad \|\bar{y}(\tau)\| < \epsilon, \quad 0 \le \tau \le 2\pi, \qquad \|\dot{\bar{y}}(\tau)\| < \epsilon, \quad 0 \le \tau \le 2\pi.$$

Let $\delta\omega$ and $y(\tau)$ be the corrections to ω_0 and $x_0(\tau)$ obtained by solving the correction equation (2.1) along with the condition (2.3) exactly. Let $\omega_1 = \omega_0 + \delta\omega$ and $x_1 = x_0 + y$. Then

$$|\omega_1 - \bar{\omega}| < C_1\epsilon^2, \qquad \|x_1(\tau) - \bar{x}(\tau)\| < C_2\epsilon^2, \quad 0 \le \tau \le 2\pi, \qquad \|\dot{x}_1(\tau) - \dot{\bar{x}}(\tau)\| < C_3\epsilon^2, \quad 0 \le \tau \le 2\pi$$

for $\epsilon < \epsilon_0$ for some $\epsilon_0 > 0$. The constants C_1, C_2, and C_3 are given by $C_1 = c_1\sqrt{m+1}(\sigma_1/\sigma_2)$, $C_2 = c_2(\sqrt{m+1}(\sigma_1/\sigma_2) + \sigma_1)$, $C_3 = c_3(\sqrt{m+1}(\sigma_1/\sigma_2) + \sigma_1)$. The constants c_1, c_2, and c_3 depend only upon the magnitudes of $f(x)$ and its first and second derivatives in a neighborhood of the periodic orbit. The constant m is an arbitrary positive integer. The constant σ_1 is given by

$$\sigma_1 = \sup_{0 \le s \le t \le 2\pi, |s-t| \le 2\pi/m} \|Y(t)Y^{-1}(s)\|,$$

where $Y(t)$ is the fundamental solution matrix of the linear system $\dot{y} = A(\tau)y$ with $Y(0)$ being the identity matrix. The constant σ_2 is the minimum singular value of a matrix that is described in the proof. The dimension of that matrix and its entries depend upon the choice of the constant m.

Proof. As shown in the proof of Theorem 3.1 in [17], $\delta\bar{\omega}$ and $\bar{y}$ satisfy the equation:

$$\omega_0\dot{\bar{y}} = A(\tau)\bar{y} + r(\tau) - \delta\bar{\omega}\dot{x}_0(\tau) + p(\tau), \tag{2.2}$$

where $p(\tau)$ is a 2π periodic function corresponding to the term ignored when the correction equation (2.1) is formed by linearizing. If the correction equation (2.1) is solved exactly, all the error in x_1 and ω_1 comes from neglecting the term $p(\tau)$. The proof of Theorem 3.1 in [17] used a single shooting procedure to represent the solution of

(2.2). To obtain the expressions for the constants C_1, C_2, and C_3 given above it is necessary to use a multiple shooting procedure to represent the solution of (2.2). To this end, let $h = 2\pi/m$ for some positive integer m. Let $M_i = Y(ih)Y^{-1}((i-1)h)$ for $1 \le i \le m$. Let $f_i = z(ih)$ for the solution of $\dot{z} = A(\tau)z + r(\tau)$, $z((i-1)h) = 0$ for $1 \le i \le m$. Similarly, let $g_i = z(ih)$ for the solution of $\dot{z} = A(\tau)z + \dot{x}_0(\tau)$, $z((i-1)h) = 0$ for $1 \le i \le m$; and $P_i = z(ih)$ for the solution of $\dot{z} = A(\tau)z + p(\tau)$, $z((i-1)h) = 0$ for $1 \le i \le m$. Let $y_i = \bar{y}(ih)$ for $0 \le i < m$. Then

$$M_1 y_0 + f_1 + P_1 = y_1 + \delta\bar{\omega} g_1, \qquad M_2 y_1 + f_2 + P_2 = y_2 + \delta\bar{\omega} g_2, \ldots,$$
$$M_m y_{m-1} + f_m + P_m = y_0 + \delta\bar{\omega} g_m.$$

Since a periodic solution of the autonomous system $\dot{x} = f(x)$ can be shifted in time, $\bar{y}$ and $\delta\bar{\omega}$ are not determined uniquely by the starting guesses x_0 and ω_0 and the exact solution. We eliminate this freedom by imposing the condition:

$$\sum_{i=0}^{m-1} f(x_0(ih))^T y_i = 0, \tag{2.3}$$

where $f(\cdot)$ is the vector field. The equations above form a linear system of dimension $md + 1$ for the variables $\delta\bar{\omega}$ and y_i, $0 \le i < m$. The constant σ_2 is the minimum singular value of this linear system. When the correction equation (2.1) is solved exactly to compute the corrections y and $\delta\omega$, all the error is due to neglecting the terms P_i in the equations above.

From this point, the proof can be completed by following the proof of Theorem 3.1 in [17]. However, a difference is in the appearance of the $\sqrt{m+1}$ factor in the bounds for the constants C_1, C_2, and C_3. We now explain this factor. As noted in [17], $\|p(\tau)\| < c\epsilon^2$ for $0 \le \tau < 2\pi$ and a constant c that depends only the second derivative of f in a neighborhood of the periodic orbit. Again as in [17], it follows that $\|P_i\| < c'\sigma_1\epsilon^2$. Since all the P_i are neglected in solving the linear system, the norm of the neglected vector is bounded above by $\sqrt{m+1}c'\sigma_1\epsilon^2$ and this is the source of the $\sqrt{m+1}$ factors in the theorem stated above. □

The development of the multiple shooting method is described in [9]. It appears possible to state Theorem 2.1 using the Banach space approach described in [9]. In the statement, we have adopted, the bounds are given in terms of quantities formed in the intermediate stages of the computation of a periodic orbit.

Theorem 2.2. *If the assumptions are as in* Theorem 2.1, *except that the corrections y, $\delta\omega$ and the derivative $\dot{y}$ are computed with relative errors less than ϵ' and not exactly, then*

$$|\omega_1 - \bar{\omega}| < C_1(\epsilon\epsilon' + \epsilon^2), \qquad \|x_1(\tau) - \bar{x}(\tau)\| < C_2(\epsilon\epsilon' + \epsilon^2), \quad 0 \le \tau \le 2\pi,$$
$$\|\dot{x}_1(\tau) - \dot{\bar{x}}(\tau)\| < C_3(\epsilon\epsilon' + \epsilon^2), \quad 0 \le \tau \le 2\pi$$

for $\epsilon < \epsilon_0$, $\epsilon' < \epsilon_0$, for some $\epsilon_0 > 0$. The constants C_1, C_2, and C_3 can be chosen as in Theorem 2.1.

Proof. We first argue that the corrections y_i and $\delta\omega$ obtained by solving the $md+1$-dimensional linear system in the proof of Theorem 2.1 *after* dropping the P_i terms have norms bounded above by a constant times ϵ. The right-hand side of that linear system is formed by the quantities f_i and these quantities are formed by propagating the residual error $r(\tau)$ of the starting guess $x_0(\tau)$, ω_0. Since the errors in the starting guess $|\omega_0 - \bar{\omega}|$, $\|x_0 - \bar{x}\|$, and $\|\dot{x}_0 - \dot{\bar{x}}\|$ are all bounded above by a constant times ϵ, the residual error $r(\tau) = f(x_0(\tau)) - \omega_0\dot{x}_0(\tau)$ is also bounded above by a constant times ϵ. This constant only depends upon f and its first two derivatives in a neighborhood of the periodic orbit. It follows that the quantities $\|f_i\|$ are bounded above by $c\sigma_1\epsilon$, for a constant c that depends on f and its two

derivatives in a neighborhood of the periodic orbit. This implies that $\|y_i\|$, $0 \le i < m$, and $\|\delta\omega\|$ are bounded above by $c(\sigma_1/\sigma_2)\epsilon$, with a constant c of the same type.

Now assume that y_i and $\delta\omega$ are computed with a relative accuracy of ϵ' and not by solving the linear system exactly as in the previous paragraph. Denote y_i and $\delta\omega$ computed as in the previous paragraph by y_i^e and $\delta\omega^e$. Then $\|y_i - y_i^e\| < c(\sigma_1/\sigma_2)\epsilon\epsilon'$ and $\|\delta\omega - \delta\omega^e\| < c(\sigma_1/\sigma_2)\epsilon\epsilon'$. Adding these bounds to those given in Theorem 2.1 completes the proof of the bounds for $|\omega_1 - \bar{\omega}|$ and $\|x_1(\tau) - \bar{x}(\tau)\|$.

To prove the bound for $\|\dot{x}_1(\tau) - \dot{\bar{x}}(\tau)\|$, it is necessary to obtain a bound for $\|\dot{y}(\tau)\|$, where $y(\tau)$ is the exact solution of the correction equation (2.1) with the condition (2.3). Such a bound is readily obtained since $\dot{y} = A(\tau)y + r(\tau) - \delta\omega\dot{x}_0(\tau)$ and y, r, and $\dot{x}_0$ have already been bounded. □

Below we give data for three periodic orbits of the Lorenz equations, all of them unstable. We give the x and y coordinates of a point on the periodic orbit with $z = 27$, and the period T of the periodic orbit. The periodic orbit AB alternates between the negative and positive wings of the Lorenz attractor. Its data is as follows:

$$x = -13.763610682134200525014401054361653864100864854092368453537864292120282774726811585294023934639503828,$$

$$y = -19.578751942451795538838041446009558866114240053427643864979133429542635474614752641597316550670467617,$$

$$T = 1.558652210716174727567870209212696070528480548997243935889521578319019875625888085435585108266014237 4.$$

Data for the periodic orbit $A^{25}B^{25}$ is as follows:

$$x = -9.166753145420366808882343568854760778922580033495914790846306848960292336794904346919383558764501032 4,$$

$$y = -9.974395112827182734516188428170964301381594712594830103556702553883086852222537322343622021232368895 6,$$

$$T = 36.043958707914849533151673848862974502009243730020190798586123062370108760553858014094134100572851733.$$

Finally, the data for the periodic orbit $ABA^2B^2 \cdots A^{15}B^{15}$ is as follows:

$$x = -13.568317317591138693791116532738086146665425413802770267307341928920639925115986035379124247913350182,$$

$$y = -19.134575113926898610482106145675590555238063694018831440659257986585209042730623744601562225619287641,$$

$$T = 171.863729139731746760144812713699868043538365275465728148429841690032093001635615971235963768748054 44.$$

All 100 digits in the numbers reported above are accurate after rounding.

The periodic orbit AB is the simplest periodic orbit of the Lorenz equations. Its leading characteristic multiplier is 4.71. The periodic orbit $A^{25}B^{25}$ passes very close to the origin, which is a part of the Lorenz attractor, as shown in [18]. Its leading characteristic multiplier is 4.8×10^{17}. The periodic orbit $AB \cdots A^{15}B^{15}$ has leading characteristic multiplier equal to 3.06×10^{59}. It is possible to verify the first 15 or 16 digits of the data given for AB using a high order numerical integrator and short time steps. If an eighth order Runge–Kutta method is used, for example, the time steps have to be about 0.01 or 0.001. However, it is not possible to verify all 100 digits of the data even for the simple periodic orbit AB using a multiple precision numerical integrator. For the local truncation error to be of the order of 10^{-100}, the time steps have to be so small that the computation becomes prohibitively expensive. In comparison, the number of Fourier grid points needed to represent the periodic orbit AB with 14 or 15 accurate digits is 128, and if the level of accuracy desired is 100 digits, the number of Fourier grid points needed is 1024. The high precision computations given in this paper take full advantage of the exponential decrease of the Fourier coefficients of analytic functions, and indeed those computations are possible only because of the use of Fourier series.

Methods of quantifying the error can be based on the residual $r(\tau)$. One such quantity, the relative residual error, was defined earlier in this section. Since we represent approximations to solutions as Fourier series, the computation of the residual is particularly direct. All that is needed is the direct substitution of the Fourier series into the differential equation $dx/dt = f(x)$ followed by an examination of how nearly the two sides of the equality cancel. If the

arithmetic uses 100 digits, the method described in this section can find an *approximation* $x_0(t)$ to the true solution whose relative residual error is as small as 10^{-100}. This does not imply that the relative error in $x_0(t)$ itself is as small as 10^{-100}. It can in fact be greater and for the sake of discussion, we will suppose that the relative error in $x_0(t)$ is as large as 10^{-50}. Even though the loss of 50 digits may seem disappointing, such a loss of accuracy is inevitable in any computation. Suppose there is another guess $x_1(t)$ whose relative accuracy is 10^{-100}. When the relative residual error is computed using $x_1(t)$, it will again be around 10^{-100}, just as for $x_0(t)$, since the arithmetic does not permit more that 100 digits in computing the residual. Thus within the confines of 100 digit arithmetic, the guesses $x_0(t)$ and $x_1(t)$ are indistinguishable and there can be no computational method for improving $x_0(t)$ to get something closer to $x_1(t)$.

The data for periodic orbits given above is mindful of the distinction between the relative residual error and the relative error of the guess itself. The connection between these two types of errors is made by the correction equation (2.1) which attempts to compute the corrections $y(\tau)$ and $\delta\omega$ from the residual $r(\tau)$. The conditioning of the connection between $r(\tau)$ and the actual errors in the guess becomes evident during the intermediate stages of the computation, and numbers that govern this conditioning have been explicitly exhibited in the bounds of Theorems 2.1 and 2.2. We have taken those numbers into account in reporting data here and in [18]. A more heuristic approach is to look at the digits of the computed approximation and observe which ones do not change with repeated iterations. It may be verified, for example, that all data for the periodic orbits AB, $A^{25}B^{25}$, and $AB\cdots A^{15}B^{15}$ is in agreement with the less accurate computations of [18]. As pointed out in Section 1, these considerations by no means imply the existence of a periodic orbit such as $A^{25}B^{25}$; they only suggest that the chances of nonexistence are extraordinarily small.

Above, data for some periodic orbits is reported at particular points along the orbits. Therefore, the distinction between normwise errors and pointwise errors must be considered as well. It is impossible in general to make the pointwise errors small in a relative sense. One of the components of a solution can be tiny, even zero, and there is no way to prevent remote parts of the periodic orbit from contributing to the error in that component. The number of digits in pointwise data that can be deemed correct is determined by absolute errors. The Fourier series used to represent periodic orbits use only thousands of terms. Thus even in the worst case where all the normwise error is concentrated in a single component, the pointwise data will be accurate in an absolute sense. Typically, we may expect the normwise error to be spread around the periodic orbit and that the relative error at a typical point along the orbit is about the same as that of the relative error of the approximation to the periodic orbit considered as a whole.

3. Fractal structure of the Lorenz attractor

As mentioned in Section 1, it is necessary to compute points located in tiny regions of the Poincaré section $z = 27$ to obtain plots of the fractal structure of the Lorenz attractor. The method we use for locating points of the Lorenz attractor in tiny regions (like the region plotted in the plot numbered 6 in Fig. 2, whose extent is 35 orders of magnitude smaller than the extent of the Lorenz attractor) is related to the inverse limit construction [7,10,19]. Locally the intersection of the Lorenz attractor with the Poincaré section looks like the product of a closed interval with a Cantor set. This Cantor set has the structure of a middle-α Cantor set, where the αth portion of a closed interval is removed from its middle and the same operation is recursively applied to the two remaining closed intervals at the ends. The fraction α is less than 1 by about 10^{-5} and is hence very close to 1. The resulting Cantor set is an uncountable union of points, which are like the leaves of an infinite binary tree, and has fractional Hausdorff dimension. If the root of the binary tree is assigned a depth 0, there are 2^d nodes at depth d, and each of these nodes corresponds to a closed interval. The lengths of the intervals corresponding to the nodes decreases geometrically with the depth of the nodes. Since the attractor looks like the product of a line segment with a middle-α Cantor set locally, it can be thought of as a union of Cantor lines, where each Cantor line is the product of the line segment

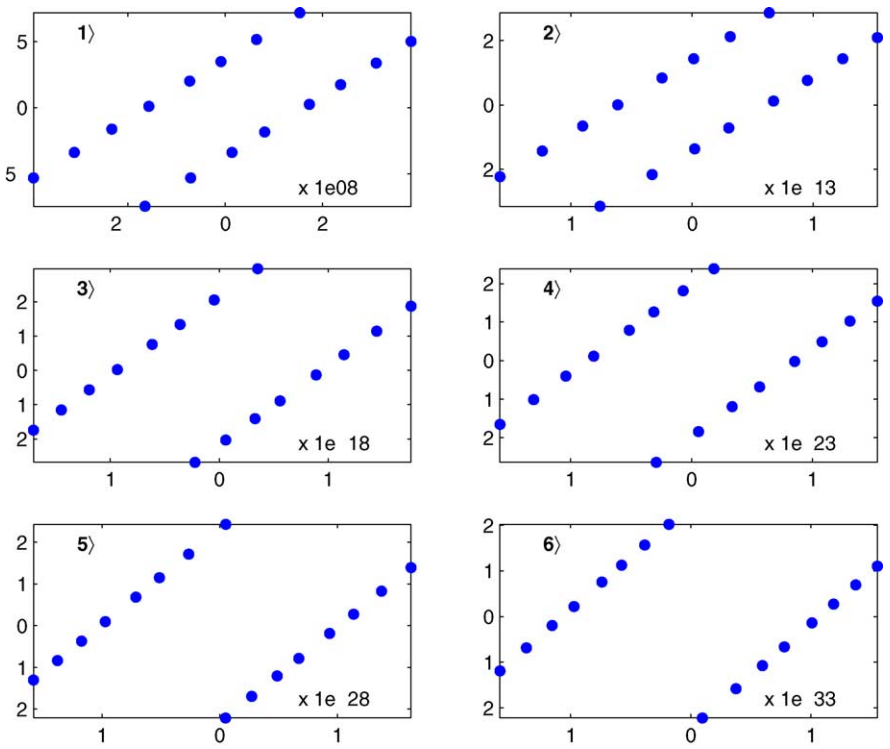

Fig. 2. The plots above are numbered from 1 to 6. Each is a plot of $y - b$ against $x - a$. Table 2 lists the points (a, b). Each plot shows two Cantor bundles of the Lorenz attractor at the Poincaré section $z = 27$. The connection of these plots to Table 1 is explained in the text. The number in the lower right-hand corner is the scale factor for the x and y axes.

with a point in the Cantor set. The product of the line segment with an interval corresponding to a node at depth d is therefore a bundle of Cantor lines. The lines in the plots in Fig. 2 correspond to these Cantor bundles. It must also be pointed out that the fraction that is removed from closed intervals is not a fixed number α for the Lorenz attractor, but varies from node to node.

Let P be the starting point of a periodic orbit whose symbol sequence is s. The symbol sequence s is related to the location of P in the Lorenz attractor. The location of P is determined by its position on the one-dimensional skeleton of the intersection of the Lorenz attractor with the Poincaré section and by the Cantor line it lies upon. Both these are related to the symbol sequence, but in different ways. The symbol sequence s determines the position of P on the one-dimensional skeleton in the same way the binary expansion of a fraction determines its position in the interval [0, 1]. Let Q be the starting point of another periodic orbit whose symbol sequence is s′. Assume that the first i symbols in s and s′ coincide but that the $(i + 1)$st symbols are different. Then the distance between P and Q along the one-dimensional skeleton will be smaller for larger i. For the Lorenz attractor, the distance roughly halves for every additional matching symbol.

The Cantor line that P lies upon is also determined by the symbol sequence s. The symbol sequence now has to be read backwards beginning with the last symbol. The symbol sequence read backwards will correspond to a node in the infinite binary tree, and if only the last few symbols of s are considered, they will correspond to a node that is the parent of the node that corresponds to s. These nodes in turn correspond to Cantor bundles, and P will lie in those Cantor bundles. It is possible to take the symbol sequence corresponding to P to be ss, sss, and so on, and find thinner and thinner Cantor bundles, and eventually a Cantor line, that must contain P. For every additional

Table 1
Plots in Fig. 2 correspond to this table[a]

n	α	β	Distance
1	ABABABABABABABABABABAB	A	2.96044e−0.8
2	ABABABABABABABABABABABABABABABABABAB	BA	1.12766e−13
3	ABABABABABABABABABABABABABABABABABABABABABABABABABABABABA	ABA	1.73144e−1.8
4	ABABABABABABABABABABABABABABABABABABABABABABABABABABABAB-ABABABABABABAB	BABA	1.57786e−23
5	ABABABABABABABABABABABABABABABABABABABABABABABABABABABAB-ABABABABABABABABABABABABABA	ABABA	1.86012e−28
6	ABABABABABABABABABABABABABABABABABABABABABABABABABABABAB-ABABABABABABABABABABABABABABAB	BABABA	1.95584e−33

[a] The columns α and β give the starting and ending symbol sequences used to determine the Cantor bundles plotted in Fig. 2. The last column gives the distance between two nearly parallel straight lines that fit the lines in the plot numbered n in Fig. 2. Another way to measure the distance between the lines is along the stable foliation.

backward symbol, the corresponding node in the infinite binary tree is at a depth greater by 1, and the girth of the corresponding Cantor bundle is smaller by a factor that is approximately 10^{-5}.

Assume that P and Q are the starting points of periodic orbits with symbol sequences $\alpha x\beta$ and $\alpha y\beta$, where α, β, x, and y are strings of As and Bs. If $x \neq y$, the distance between P and Q along the one-dimensional skeleton of the intersection of the Lorenz attractor with the Poincaré section is determined by the number of symbols in α. The distance between the Cantor lines that contain P and Q (this distance is along a line transverse to the one-dimensional skeleton) is determined by the number of symbols in β. Since every additional symbol in α causes the distance to decrease by a factor that is about 1/2, while every additional symbol in β causes the distance to decrease by a factor that is about 10^{-5}, it is necessary to take many more symbols in α than in β if these distances are to be balanced. The lengths of the αs and the βs listed in Table 1 are chosen so that the two distances are of the same order for points P and Q. Each plot in Fig. 2 is obtained by plotting the starting points of all periodic orbits with symbol sequence of the form $\alpha x\beta$, where x varies over the 16 possible symbol sequences of length 4. The 16 points line up along two lines in each case. These lines are the Cantor bundles for the ending sequences Aβ and Bβ, respectively. It is also evident from Table 1 that the starting sequence α for a higher numbered plot is obtained by appending symbols at the end of the α for a lower numbered plot. Similarly, the ending sequence β for a plot is obtained by appending a single symbol at the front of the β for the previous plot. This ensures that each plot zooms in on a tiny region in the previous plot, too tiny in fact to be shown in Fig. 2. The containment relationship between a plot and its previous plot can be inferred from Table 2, however. Further, because the β's increase one symbol at a time, the Cantor bundles that can be seen in a plot are obtained by removing the middle portions of one of the two Cantor bundles in the previous plot.

Fig. 2 may give the impression that Cantor bundles of the Lorenz attractor are straight lines. This is far from being true. For example, it is clear from Table 3 that the distance between the same Cantor bundles can depend a

Table 2
The points of the Lorenz attractor plotted in Fig. 2 are centered at $x = a$ and $y = b$ in the plane $z = 27$

n	a	b
1	−13.764318917690414925022830902023965876189	−19.578490037982618846341494795388351482041
2	−13.764318894388820868962416471783533350650	−19.578490021642010277850042121844164884100
3	−13.764318894388828755035202962819358269485	−19.578490021641887589995227600060138207047
4	−13.764318894388828754638773391293198413650	−19.578490021641887591249375829935211870775
5	−13.764318894388828754638777146411062863846	−19.578490021641887591249364726660511246751
6	−13.764318894388828754638777146356549931052	−19.578490021641887591249364726767981118157

Table 3
The first and second set of rows in the table below use the same ending symbol sequences β but different starting sequences α[a]

α	β	Distance
BAAAAABABBAAABABBABABBAB	A	1.79459e−10
BAAAAABABBAAABABBABABBABBAAABABBABBBAAAB	AA	8.70484e−15
BAAAAABABBAAABABBABABBABBAAABABBABBBAAABAABABBAAABBBAAA	BAA	2.37141e−20
BBBAAABAAABBBBAAABBBBAAB	A	1.47793e−08
BBBAAABAAABBBBAAABBBBAABABABBABBBAAABABB	AA	4.22756e−13
BBBAAABAAABBBBAAABBBBAABABABBABBBAAABABBAABABABBBBABABB	BAA	2.35906e−18

[a] The distance is between the Cantor bundles corresponding to the ending sequences Aβ and Bβ in the region of the Lorenz attractor determined by the starting sequence α.

lot on the point along their lengths where the distance is computed. The Cantor bundles appear to come close and move apart in a complicated way. It is difficult to predict precisely where the Cantor bundles come close and where they move apart. But it is possible that the Cantor bundles corresponding to the first row of the first table in Table 3 come close because of the string of 5 As following a single B in the starting sequence. From the plots in Fig. 2 and Tables 1 and 3, it seems that the Lorenz attractor is a very complicated type of fractal.

In one respect, the investigation of the fractal structure of the Lorenz attractor given in this paper is incomplete. For a mathematically tractable model of the Lorenz equations, it is shown in [19] that the Cantor leaves are tied to the unstable manifold in the same way that the pages of a book are attached to its spine. In the Lorenz attractor too something unusual has to happen to Cantor bundles, such as the ones shown in Fig. 2, as they approach the part of the unstable manifold near the boundary of the attractor. A computation of the fractal structure of the Lorenz equations in this region has not been carried out. We have discussed this issue at greater length in Section 5 of [18], where we also discuss a possible method for investigating the Lorenz attractor in this tiny region.

4. Hausdorff dimension of the Lorenz attractor

The Hausdorff dimension is a number that can be associated with any fractal, however complicated. The number of ϵ-balls needed to cover a fractal increases as ϵ decreases. The Hausdorff dimension may be recovered from the power law scaling relationship between this number and ϵ. Unless the construction of the fractal set is especially simple, the Hausdorff dimension is a difficult number to get hold of. The computations we have reported in the previous section are sufficient only for a crude estimate of the Hausdorff dimension.

A known formula connects characteristic multipliers of periodic orbits to the Hausdorff dimension of a chaotic attractor. Let F be an invertible map of the plane with a chaotic attractor. Define D_m by

$$\sum_i \frac{1}{\lambda_i} \lambda_i^{D_m} = 1, \tag{4.1}$$

where the summation is over all points x_i in the attractor with $F^m(x_i) = x_i$ and λ_i the stable characteristic multiplier of F^m at x_i. It is argued in [2,6] that D_m converges to the Hausdorff dimension as $m \to \infty$.

In Table 4, we have used computations reported in [18] and the formula for D_m above. However, each D_m listed in Table 4 gives an estimate for the entire Lorenz attractor and is therefore greater than the D_m obtained from (4.1) by 1. Computations using the formula are reported in [4,5]. Our computation agrees with [5] except for the case $m = 12$. The disagreement for $m = 12$ is because the computation reported in [5] misses some periodic orbits with

Table 4
Estimates D_m for the Hausdorff dimension of the Lorenz attractor obtained using the characteristic multipliers of 111 011 periodic orbits with symbol sequences of length 20 or less and (4.1)

m	D_m
4	2.05915144e+00
5	2.06124450e+00
6	2.06207150e+00
7	2.06244322e+00
8	2.06257995e+00
9	2.06265149e+00
10	2.06268522e+00
11	2.06270128e+00
12	2.06270901e+00
13	2.06271278e+00
14	2.06271464e+00
15	2.06271557e+00
16	2.06271607e+00
17	2.06271589e+00
18	2.06271598e+00
19	2.06271603e+00
20	2.06271606e+00

symbol sequences of length 12 that actually exist. The computation in [5] goes up to $m = 13$. The computation in [4] goes up to $m = 9$, but the numbers in Table 1 of [4] do not agree with the numbers in our Table 4.

The numbers in Table 4 are converging to 2.0627160... But is that number really the Hausdorff dimension of the Lorenz attractor? To answer this question, we attempt to connect the computations in the previous section to (4.1). Arguments that lead to (4.1) use periodic orbits like our computations in the previous section, but the geometry of the Lorenz attractor is inferred from the periodic orbits in a way that is entirely different. To explain the way periodic orbits are connected to the geometry of the Lorenz attractor to obtain (4.1), we assume $m = 3$. With $m = 3$, the Lorenz attractor has two distinct periodic orbits with the symbol sequences AAB and BBA, respectively. These give six periodic points of the map of the Poincaré plane induced by the Lorenz flow. We have seen in the previous section that Cantor bundles can be obtained by fixing the ending sequence and varying the starting sequence. If we consider the six Cantor bundles corresponding to the ending sequences AAB, ABA, BAA, BBA, BAB, and ABB, each of the six periodic point will lie on a different one of these Cantor bundles. There are Cantor bundles corresponding to the ending sequences AAA and BBB as well, but there are no periodic points in the sum with $m = 3$ that correspond to these. To derive formula (4.1), it is necessary to think of each of these Cantor bundles as having O(1) length and constant girth. The formula uses the stable characteristic multiplier of the periodic point on the Cantor bundle as an estimate for its girth which, it is hoped, is off by atmost a constant factor. With $m = 20$, the formula pictures the section of the Lorenz attractor as being made up of 2^{20}—2 Cantor bundles of uniform girth and O(1) length, and it infers the scale of the girth of a Cantor bundle using the stable characteristic multiplier of a single periodic point on that Cantor bundle. We can use the computations reported in the previous section to check if the Cantor bundles are of uniform girth and if their girth is correctly estimated by characteristic multipliers.

The distances listed in Tables 1 and 3 are the distances between Cantor bundles with ending sequences Aβ and Bβ and give a good approximation to the girth of the Cantor bundle corresponding to the ending sequence β. It follows from Table 3 that the girth can vary along the length of the section of the Lorenz attractor. There is still the possibility however that the girth might be more and more uniform over the length of Cantor bundles corresponding to longer ending sequences. For the Cantor bundles with the ending sequences BA, BABA, and BABABA the girths

are about a thousandth of the estimate obtained from characteristic multipliers. This is not so bad, because it is more important to get a scaling factor that is independent of the Cantor bundle than to get a scaling factor that is close to 1. Thus the computations of the previous section raise questions about the assumptions behind (4.1), but are unable to conclusively verify or deny their validity. More extensive computations needed to examine those assumptions are not possible with available computing resources.

It is not difficult to construct fractals in the plane whose local structure is given by a product of a line segment with a middle-α type Cantor set, but whose Cantor bundles are of nonuniform girth. This can be done for example by smoothly varying α along the length of the line segment. If such fractals can be chaotic attractors, the main assumption behind (4.1) will be false.

5. Application of periodic orbit theory to the Lorenz attractor

The starting point for the application of periodic orbit theory to the Lorenz attractor is a linear operator that propagates probability densities supported inside the Poincaré section according the return map of the Lorenz flow; see [1,3]. The traces of powers of this operator can be formally expanded as a sum over periodic orbits. The information contained in these traces can be converted into a spectral determinant which bears the same relationship to the return map as the characteristic polynomial does to a matrix. The spectral determinant, however, is not a polynomial although it is known to be an entire function for some flows [3]. The first 20 terms of the spectral determinant for the Lorenz attractor are given below:

$$\begin{aligned} 1 \; & -0.269331z^2 - 0.244951z^3 - 0.172774z^4 - 0.113168z^5 - 0.072512z^6 - 0.0462723z^7 \\ & - 0.0295876z^8 - 0.018885z^9 - 0.0120788z^{10} - 0.00772456z^{11} - 0.00492473z^{12} - 0.00312169z^{13} \\ & - 0.00196396z^{14} - 0.00122612z^{15} - 0.00076226z^{16} - 0.00038957z^{17} - 0.000197666z^{18} \\ & - 0.0000999565z^{19} - 0.0000493542z^{20}. \end{aligned}$$

The escape rate for the Lorenz attractor implied by the above polynomial is 5×10^{-6}. Since the terms above use only periodic orbits with symbol sequences of length equal to or less than 20, it is natural to expect that an occurrence of a repeated sequence of As or Bs of length 20 or greater will be interpreted by the above truncation as an escape from the attractor. If As and Bs are taken to be equally probable and independent, the probability of occurrence of 20 consecutive As or Bs is $1/2^{19}$, which leads to an estimate of the escape rate that is in the same ball park.

There is a formula to obtain an estimate of the Lyapunov exponent of the Lorenz flow using periodic orbits; see appendix of [3]. Use of this formula leads to the number 0.905630. The magnitude of the last term before truncation suggests that the error in this number is less than 0.00005.

6. Conclusions

The Lorenz equations are but one set of differential equations that exhibit chaos. They are peculiar in several ways, as so definite a set of equations must be, and also very simple algebraically. The greater complexity and differences in detail of physical models may lead one to conclude that the significance of the numerical investigations presented here and in [18] is limited to the Lorenz equations and is therefore narrow.

Lorenz explains in his book [11] that the Lorenz equations were derived in order to suggest a theory of turbulence. Although the Lorenz equations are derived by Galerkin truncation of the basic equations of fluid dynamics, the elimination of all but a few modes in the truncation and the somewhat arbitrary choice of parameters imply that

any direct connection to fluid flow is unlikely. However, since the work of Lorenz the presence of chaos in fluid flow has been discovered to be widespread and theories of types of turbulence based on strange attractors have been advanced. Thus in the work of Lorenz we have an example of the investigation of a narrow mathematical problem having broader consequences.

Numerical investigations, even though unreliable, can provide information where mathematics is intractable. In almost all of computational science, dynamical models are integrated forwards in time by Runge–Kutta-like methods that attempt to match terms of Taylor series expansions. As a result, a number of possibilities suggested by the mathematical theory of dynamical systems are completely out of reach of typical computations. Attempts to close this gap between conceptual possibilities and mathematical theorems on the one hand, and what can be computed on the other have to begin with particular examples.

Acknowledgements

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

References

[1] R. Artuso, E. Aurell, P. Cvitanović, Recycling of strange sets. I. Cycle expansions, Nonlinearity 3 (1990) 325–360.
[2] D. Auerbach, P. Cvitanović, J.-P. Eckmann, G.H. Gunaratne, I. Procaccia, Exploring chaotic motion through periodic orbits, Phys. Rev. Lett. 58 (1987) 2387.
[3] P. Cvitanović, R. Artuso, P. Dahlquist, R. Mainieri, G. Tanner, G. Vattay, N. Whelan, A. Wirzba, Classical and Quantum Chaos, 2001. http://www.nbi.dk/ChaosBook/.
[4] B. Eckhardt, G. Ott, Periodic orbit analysis of the Lorenz attractor, Zeit. Phys. B 93 (1994) 258–266.
[5] V. Franceschini, C. Giberti, Z. Zhang, Characterization of the Lorenz attractor by unstable periodic orbits, Nonlinearity 6 (1993) 251–258.
[6] C. Grebogi, E. Ott, J.A. Yorke, Unstable periodic orbits and the dimension of multifractal chaotic attractors, Phys. Rev. A 37 (1988) 1711–1724.
[7] J. Guckenheimer, A strange, strange attractor, in: J. Marsden, M. McCracken (Eds.), The Hopf Bifurcation and its Applications, Springer, New York, 1979.
[8] N.J. Higham, Accuracy and Stability of Numerical Algorithms, SIAM, Philadelphia, 1996.
[9] H.B. Keller, Numerical Methods for Two-point Boundary Value Problems, Dover, New York, 1992.
[10] E.N. Lorenz, Deterministic non-periodic flows, J. Atmos. Sci. 20 (1963) 130–141.
[11] E.N. Lorenz, The Essence of Chaos, University of Washington Press, Seattle, 1995.
[12] K. Mischaikow, Topological techniques for efficient rigorous computations in dynamics, Acta Numer. 11 (2002) 435–477.
[13] K. Mischaikow, M. Mrozek, A. Szymaczak, Chaos in the Lorenz equations: a computer assisted proof III, classical parameter values, J. Diff. Eqs. 169 (2001) 17–56.
[14] G.W. Stewart, Introduction to Matrix Computations, Academic Press, New York, 1973.
[15] S.H. Strogatz, Nonlinear Dynamics and Chaos, Addison-Wesley, Reading, MA, 1994.
[16] W. Tucker, A rigorous ODE solver and Smale's 14th problem, Found. Comput. Math. 2 (2002) 53–117.
[17] D. Viswanath, Lindstedt–Poincaré technique as an algorithm for computing periodic orbits, SIAM Rev. 43 (3) (2001) 478–496.
[18] D. Viswanath, Symbolic dynamics and periodic orbits of the Lorenz attractor, Nonlinearity 16 (2003) 1035–1056.
[19] R.F. Williams, Structure of Lorenz attractors, Publ. Math. IHES 50 (1980) 73–100.