Beyond Belief:
Randomness, Prediction and Explanation in Science

Editors

John L. Casti

Institute of Econometrics, Operations Research, and System Theory
Technical University of Vienna
Vienna, Austria

Anders Karlqvist

Royal Swedish Academy of Sciences
Stockholm, Sweden

CRC Press
Boca Raton Ann Arbor Boston

CHAPTER 6

Dimensions of Atmospheric Variability

R. T. PIERREHUMBERT

Abstract

We consider a number of applications of the correlation dimension concept in the atmospheric sciences. Our emphasis is on the correlation dimension as a nonlinear signal-processing tool for characterizing the complexity of real and simulated atmospheric data, rather than as a means for justifying low-dimensional approximations to the underlying dynamics. Following an introductory exposition of the basic mathematics, we apply the analysis to a three-equation nonlinear system having some interesting points of affinity with the general subject of nonlinear energy transfers in two-dimensional fluids. Then we turn to the analysis of a 40-year data set of observed Northern Hemisphere flow patterns. Our approach deviates from most previous studies in that we employ time series of flow fields as our basic unit of analysis, rather than single-point time series. The main evidence for low-dimensional behavior is found in the patterns of interseasonal variability. However, even when this is removed, the streamfunction field shows clear indication of dimensionality in the range 20–200, rather than in the thousands. Due to problems connected with the nonequivalence of various norms in functions spaces, we do not claim that this represents the "true" dimensionality of the underlying system. Nevertheless, it does show the existence of a considerable amount of order in the system, a fact begging for an explanation.

From this, we turn to the matter of spatial complexity, examining the geometry of clouds of passive tracer mixed by spatially structured (atmospherically motivated) two-dimensional flow fields. Evidence is presented that the cloud ultimately mixes over a region characterized by dimension two. We also demonstrate that the correlation dimension of the tracer cloud is directly related to the algebraic power spectrum of the concentration distribution in Fourier space. From this we are led to some speculations on the role of chaotic mixing in the enstrophy cascade of two-dimensional turbulence.

Finally, we consider the spatial patterns of rate of predictability loss in the tracer problem. This information is obtained by computing finite-time estimates of the Lyapunov exponents for trajectories starting from various initial conditions. The rate of predictability loss is itself unpredictable, in the sense that it exhibits sensitive dependence on initial conditions. It is suggested that multifractal analysis could be used to characterize the spatial pattern of predictability.

1. *Introduction*

Of the various measures of atmospheric variability, the *correlation dimension,* introduced in dynamical systems theory by Grassberger and Procaccia (1983), is one that has recently received a lot of attention. This is a convenient and readily computable way to estimate the fractal dimension of an object traced out by a time series of points in some high-dimensional embedding space. Specifically, if we have a system whose time variation is described by state $\Psi(t)$ (where, for example, Ψ is the set of all wind and temperature values at every point in the atmosphere at time t), and a function $d(\Psi_1, \Psi_2)$ measuring the distance between states, then the correlation dimension is calculated as follows: Discretize the time series to form a sequence of states Ψ_j, and compute the distances d_{ij} between each pair of states. From this form the cumulative histogram $H(r)$ for the number of pairs of distance less than r. If $H(r)$ is well-approximated by the form $H = ar^d$, then the exponent d is said to be the correlation dimension. Grassberger and Procaccia give numerous examples showing why this quantity should be thought of as a dimension, as well as indicating how it relates to other measures of dimension.

The analysis has been applied to climatic variations on geological time scales (Grassberger 1986, Maasch 1989, *inter alia),* to daily weather fluctuations at a single station (Essex, et al. 1987), and to microscale fluctuations having time scales of seconds to minutes (Tsonis and Elsner, 1988). The analyses typically indicate low dimensionality and indeed often seem to be motivated by the (probably forlorn) hope of justifying the modelling of weather or climate in terms of a small set of ordinary differential equations. Whether or not this quest is tantamount to tilting at windmills, the correlation dimension analysis does provide useful information about atmospheric variability, information that goes beyond the bounds of traditional linear statistics. Such analyses reveal the extent to which the actual variations are concentrated on a limited subset of the space of all possible variations. It is in this spirit that we present the ideas of this chapter.

As an example of the geometric information contained in the $H(r)$ curve, consider Figure 1. The upper example shows a "thin wire grid" embedded in R^2. The object is one-dimensional viewed at small scales, but of dimension two when viewed at large scales. This would be reflected in a change in the slope of the $\log H(r)$ vs. $\log r$ curve at a distance r_0 corresponding to the average spacing of the grids. In the bottom example we consider the $H(r)$ graph for a "relaxed earthworm,"

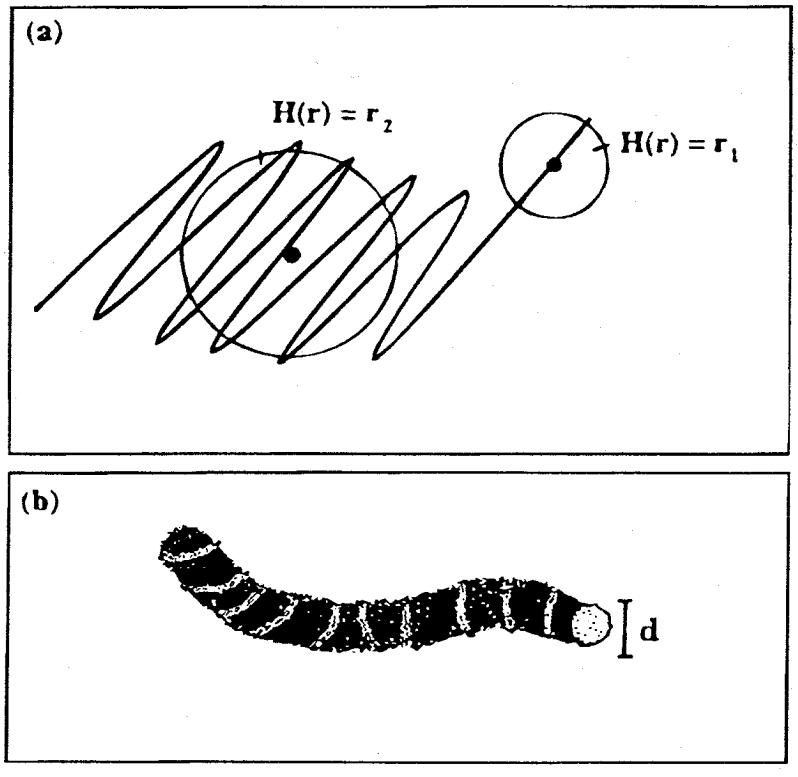

Figure 1. Geometric interpretation of the correlation dimension.

which displays a one-dimensional structure at large scales but is of dimension three at scales smaller than the thickness of the earthworm. If the earthworm should become agitated and decide to curl up into a spiral on the ground, the 1-D large-scale regime would be replaced by a 2-D regime. All objects which can be embedded in a compact subset of their phase space have $H(r)$ becoming constant for values of r sufficiently large to contain the entire object. Thus, all localized objects look zero-dimensional (pointlike) if you stand back far enough.

It is instructive to compare the correlation dimension analysis with the "empirical orthogonal function" (EOF) technique, which currently enjoys much wider application in the atmospheric sciences. This technique is a systematic way of finding a *linear subspace* of a high-dimensional embedding space, to which the trajectory of the system is approximately confined. Given a vector time series $(x_1(t), \ldots, x_n(t))$ in which each component has zero mean and unit variance, the EOF's are simply the n eigenfunctions of the $n \times n$ covariance matrix $< x_i x_j >$, where the brackets denote time averaging. The eigenvalues give the proportion of the variance of the signal explained by projection onto

the corresponding eigenvector. The use of EOF's in characterizing atmospheric variability was pioneered by Lorenz, though the concept first emerged in the statistics literature through the work of Karhunen and Loéve. The minimum number of EOF's needed to account for, say, 90% of the total variance is sometimes referred to as the "Karhunen-Loéve dimension."

By way of example, consider the motion $x(t) = \cos\omega t$, $y(t) = \sin\omega t$. The eigenvectors of the correlation matrix are $(1,0)$ and $(0,1)$, and each explains 50% of the variance; the Karhunen-Loéve dimension is 2. However the correlation dimension is 1, as the motion traces out a circle in the x-y plane. The correlation dimension analysis provides a much better description of the confinement of the trajectory to a limited portion of the embedding space, because it does not need to confine the trajectory within a *linear* subspace. On the other hand, a strength of the empirical orthogonal function analysis is that it provides some information on the *patterns* of variation via the eigenvectors, which have no counterpart in the correlation dimension analysis.

An extension of this example demonstrates that the correlation dimension analysis must be interpreted with great care if one's objective is to determine the dimensionality of the underlying dynamics. Suppose that the frequency ω is time dependent (say, $\omega = \omega_0 + \alpha \sin \Omega t$, for the sake of concreteness), rather than being constant. Then a correlation dimension estimate based on the trajectory in x-y space is bound to yield a dimension of unity or less, since the trajectory is restricted to the unit circle. However, knowlege of the initial point on the circle does not determine the future course of the system, as one must also know the phase Ωt; that is, the system is actually two-dimensional. Thus, while the naively determined correlation dimension gives us a useful description of the kind of set the trajectory lives on in x-y space, it does not tell us how many equations are needed to predict the future.

The problem stems from the fact that the x-y plane in which the trajectory is being embedded is not a *phase space* for the system. The rectification of the problem is, in principal, equally simple. One forms the family of $2(N+1)$-dimensional synthetic phase spaces consisting of the x-y point and its time derivatives up to Nth order, and performs the dimension analysis for ever larger N until the answer converges. In the example above, the x-y point and all its time derivatives can be expressed in terms of the two phases $\theta_1 = \Omega t$ and $\theta_2 = \omega(\theta_1)t$, which define a two-dimensional manifold embedded in the $2(N+1)$-dimensional space. In practice, when applying the method to data one forms the synthetic phase spaces by employing time-lagged se-

quences $(x_n, y_n, x_{n-1}, y_{n-1}, \ldots)$, which is equivalent to computing the time derivatives by finite differencing.

In this chapter we examine the application of the correlation dimension analysis as a tool for characterizing the variability of a number of actual and simulated atmospheric time series. We begin in §2 with an analysis of a simple system of ordinary differential equations whose behavior can also be directly visualized geometrically. In §3 we apply the methods to a 40-year data set of actual atmospheric flow patterns. In contrast to most previous studies, this one deals with time series of *fields* covering most of the Northern Hemisphere, rather than single-point time series of weather or climate data. We then take up the characterization of the complexity of individual two-dimensional patterns in §4, showing how geometrically simple flow fields can nevertheless generate intricate patterns in tracers and vorticity. In this section we also discuss the relation between fractal dimensions and the spatial spectrum of concentration variance. In §5 we come up against yet another form of spatial complexity. The spatial complexity in tracer patterns arises from loss of predictability of the individual particle trajectories. It turns out that the spatial pattern characterizing the *rate* of predictability loss is itself fractal in nature, so that even the degree of unpredictability is in some sense unpredictable. Some general conclusions and speculations are given in §6.

2. *A Toy Atmosphere*

To illustrate the key ideas, we first treat a simple model embodying some of the energy transfers and dissipative characteristics of atmospheric flow. Consider the equations

$$\frac{dA_1}{dt} = A_2 A_3 - d_1 A_1 \tag{2.1}$$

$$\frac{dA_2}{dt} = -2A_1 A_3 + gA_2 \tag{2.2}$$

$$\frac{dA_3}{dt} = A_1 A_2 - d_3 A_3 \tag{2.3}$$

In the absence of dissipation ($d_1 = d_3 = g = 0$) these equations reduce to the conventional triad equations, and conserve the energy and enstrophy related quantities

$$E = A_1^2 + A_2^2 + A_3^2, \quad F = A_1^2 - A_3^2 \tag{2.4}$$

Hence the solutions are closed orbits defined by the intersection of the sphere E = constant and the hyperboloid F = constant. These equations are identical to the gyroscopic equations describing the evolution of the angular momentum vector of a rigid body. The inviscid triad equations can be derived from a number of *ad hoc* and formal asymptotic approximations to 2-D fluid flow, and the amplitudes A_j may be thought of as the amplitudes of three scales of motion (e.g. the amplitudes of a triplet of resonantly interacting Rossby waves, as described in Pedlosky 1979). Motion concentrated initially in the middle scale ($j = 2$) is unstable, and transfers energy to the neighboring scales. This process is at the heart of two-dimensional turbulence.

But the real atmosphere is not a closed system; it is forced by energy input at intermediate scales, and nonlinear processes transfer energy to other scales where it is dissipated. In Eq. (2.2), g represents the generation process, while d_1 and d_3 in Eqs. (2.1) and (2.3) represent dissipation. This system has the interesting property that a trajectory located initially on one of the axes remains there indefinitely. Thus the three trajectories

$$(A_1, A_2, A_3) = (a\exp(-d_1 t), 0, 0) \tag{2.4a}$$

$$= (0, a\exp(gt), 0) \tag{2.4b}$$

$$= (0, 0, a\exp(-d_3 t)) \tag{2.4c}$$

are all exact nonlinear solutions of Eqs. (2.1)–(2.3). This means that the dissipation in itself cannot be relied on to prevent the growth of energy in the system; to keep the trajectory from running off to infinity along the 2-axis, energy must be transferred out of mode 2 and into modes 1 and 3 where it can be dissipated.

This property is not an artifact of the *ad hoc* truncation; it is also encountered in partial differential equations describing fluid systems that are more like the real atmospheric flow. For example, consider the damped barotropic vorticity equation, augmented to make a single mode unstable:

$$\partial_t \nabla^2\psi + \mathcal{J}(\psi, \nabla^2\psi) = -d\nabla^2\psi + g(\nabla^2\psi)_{nm} \tag{2.5}$$

where ψ is the 2-D streamfunction, $\nabla^2\psi$ the vorticity, $(\nabla^2\psi)_{nm}$ the projection of the vorticity on the Fourier mode (n, m), and $\mathcal{J}$ is the Jacobian matrix defined to be $\mathcal{J}(A, B) = \{\partial_x A\partial_y B - \partial_x B\partial_y A\}$. One can think of the artificial instability as a surrogate for various instability processes (such as baroclinic instability) that cannot be captured in a

single-layer model. If the system is initialized with a pure Fourier mode (n, m), the Jacobian vanishes identically, and the amplitude of the mode grows without bound. However, the mode is subject to a sideband instability rather similar to that occuring in the triad model, and so nonlinear transfer out of the mode can cause enough dissipation to limit growth. We are not aware of any integrations of Eq. (2.5) addressing the boundedness of the flow, but a similar question occurs naturally in connection with baroclinic instability in an infinite x-y plane, and there the runaway mode appears to have bounded amplitude. (I. Held, personal communication).

Now let's get a bit more mathematical. The origin, (0, 0, 0) is a fixed point of Eqs. (2.1)–(2.3). The directions (1, 0, 0) and (0, 0, 1) are stable, while the direction (0, 1, 0) is unstable. In fact, given Eq. (2.4), the *unstable manifold* of the system (defined crudely as the set of trajectories coming out of the fixed point at directions tangent to the local unstable direction) is simply the 2-axis. The *stable manifold,* which is the set of points which flow into the fixed point with time, is two-dimensional because there are two stable directions; it *contains* the 1-axis and 3-axis, but the structure in between can be (and probably is) highly contorted. A trajectory located on the unstable manifold initially will of course run directly off to infinity, but what happens if it is slightly displaced? To answer this question, we linearize Eqs. (2.1)–(2.3) about the unstable trajectory (2.4b), finding

$$\frac{d}{dt}(A_1, A_3) = (A_1, A_3)\exp(gt) \tag{2.6}$$

Hence neighboring trajectories deviate from the unstable manifold like the exponential of the exponential of time. Clearly, the unstable manifold is a very unstable creature indeed. This prevents the trajectory from shadowing the unstable manifold and running directly off to infinity. However, after deviating from the unstable manifold, the trajectory can come close to the two-dimensional stable manifold and be swept back towards the origin, where it begins the process anew. Thus trajectories are expected to funnel into the origin along the stable manifold and fountain outward from there along the one-dimensional unstable manifold.

Numerical calculations show that this is indeed what happens. The case $d_1 = d_3$ is somewhat pathological, as the quantity F then becomes monotonically decreasing and the trajectories asymptotically collapse onto the crossed planes $A_1 = \pm A_3$, and thence spiral out to infinity along a highly organized trajectory. When $d_1 \neq d_3$ the trajectory is

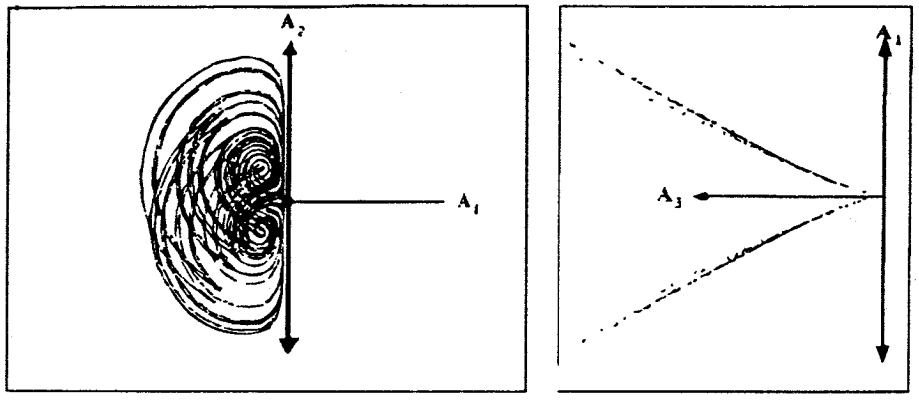

Figure 2(a). Phase space structure of the triad system. Left: projection of trajectories on A_1-A_3 plane. Right: cross section showing intersection of trajectories with the plane $A_2 = 0$. Parameters are $(d_1, g, d_3) = (0.5, 0.25, 1)$.

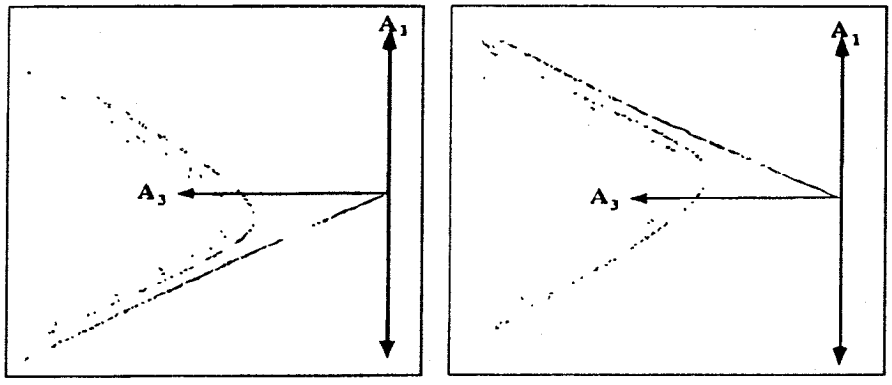

Figure 2(b). Cross sections as in Figure 2(a), but for $A_2 = 1.5$ and $A_2 = -1.5$.

more chaotic and unpredictable. A typical set of points tracing out the attractor for $(d_1, g, d_3) = (0.5, 0.25, 1)$ is shown in Figure 2.

Does the system remain within a bounded region of phase space as the integeration is carried out over ever longer times? This question is equivalent to asking whether the unstable fixed point (0, 0, 0) borders the attractor, since trajectories beginning arbitrarily close to the unstable manifold will attain correspondingly large amplitudes before they

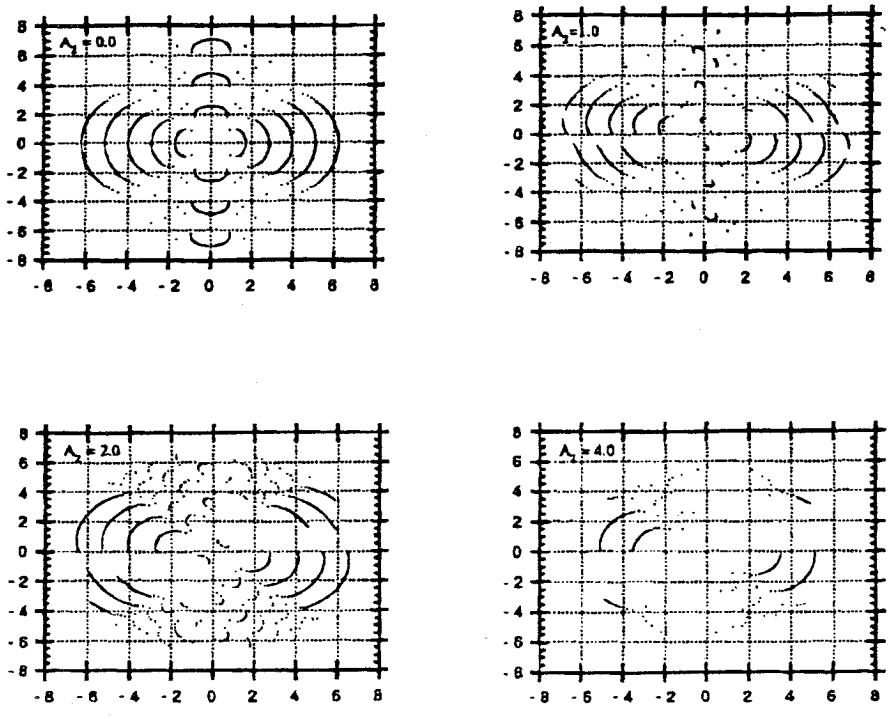

Figure 3. Cross sections through the stable manifold of the triad system for various values of A_2.

deviate. The structure of the *stable* manifold gives some clue of the likelihood of this. Some slices through this 2-D structure are shown in Figure 3; the manifold is highly contorted, rather like a crumpled ball of tissue paper. The series of shell-like structures visible in the cross-sections actually come from sheets winding helically about the A_1 and A_3 axes, giving the stable manifold the appearance of a series of four corkscrews or auger bits with points directed toward the origin along the 1 and 3 axes. Being more or less space-filling, it is rather hard to avoid, and hence trajectories should eventually come near it and be carried near (0, 0, 0). Based on the numerical results (and further analyses to be presented below), we conjecture that the amplitude is indeed unbounded. This conjecture has not yet been proved, however.

The geometric influence of the one-dimensional unstable manifold is clearly present in the trajectories. For this system, which has a three-dimensional phase space, most of what we would like to know about the geometry of the attractor can be inferred directly from looking at plots, a consequence of the fact that we are three-dimensional beings and have brains capable of interpreting three-dimensional geometric structures (or at least inferring them from Platonic shadows on the

two-dimensional retina). In higher dimensions this is not possible, and so one must turn to various numerical descriptions of the geometry. The correlation dimension is the tool we shall employ here. Let us see how the geometric structure of the triad model is reflected in its correlation dimension.

We took a data set of 50,000 points from the triad model trajectories (equally spaced in time), and computed the Euclidean distance in 3-space between each pair of points. The histogram of the set of differences is shown on a log-log plot in Figure 4(a). We show results for the full data set and also for a subset consisting of the first 5,000 points. There is an extensive range of distances for which the attractor is characterized by a correlation dimension of approximately 1.5. This is consonant with what we know about its geometry: the trajectory spends much of its time clustered about the one-dimensional unstable manifold, but during its necessary excursions away from this manifold it executes a trajectory along one of the two planar "ears" sewn to the unstable manifold like pages onto the spine of a book, if we may mix metaphors. Thus, it makes sense for the dimension to be between 1 and 2.

At very small distances for the 5,000-point plot, the slope approaches unity; here we are beginning to pick up the geometry of individual trajectories, which are one-dimensional curves. For the 50,000-point plot, this behavior sets in only at still smaller distances, as the trajectory has had time to fill out the attractor more densely. At large distances the curve asymptotes to a constant, reflecting the obvious fact that for a finite set of points the distance must have an upper bound. A detailed examination of the behavior at large distances sheds some light on the question of whether the attractor lies in a bounded region of phase space, which is identical to the question of whether the histogram becomes exactly flat above some critical distance. First note that the curve for the 50,000-point data set is almost exactly parallel to that for the first 5,000 data points. To underscore this feature, in Figure 4(b) we have plotted the 5,000-point histogram rescaled by the ratio of the total number of pairs (a bit more than 100). The rescaled curve lies precisely atop the 50,000-point histogram, except at very small distances, indicating that the geometric structure of the attractor has been well defined.

A closer look at the large-distance tail (Figure 4(b)) shows that a larger maximum distance is attained in the 50,000-point case than in the 5,000-point case. In fact, the tail region of the histogram can be well fitted by a curve of the form $1 - ae^{-bd}N$, where N is the

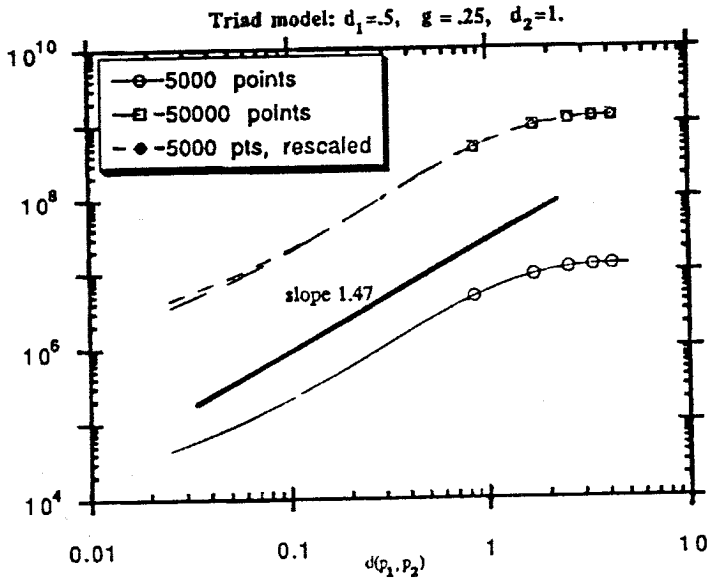

Figure 4 (a). Correlation dimension plot (log-log cumulative histogram) for the triad system.

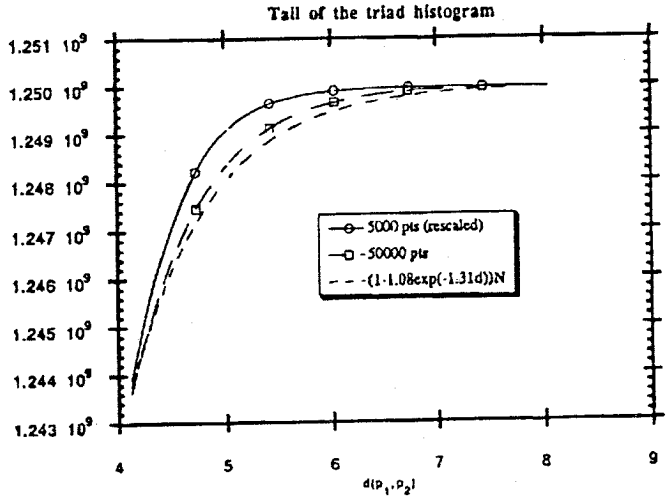

Figure 4(b). As in Figure 4(a), but showing a magnification of the tail region.

number of pairs. From this we infer a finite, though exponentially small, probability of arbitrarily large excursions. However, the distribution is *integrable* at large d, so that averages of any polynomial quantity (such as the variance) converge. This situation contrasts with $1/f$ noise.

The state of the real atmosphere is characterized by an infinite-dimensional phase space (the space of the values of wind, temperature, etc. at each point of the atmosphere), but it may be that relatively few directions in phase space are unstable. It is the signature of such a low-dimensional unstable manifold that we seek in the next section.

3. The Real Atmosphere

The data source for this study is a subset of the operational U.S. National Meteorological Center archive for 1940–1985. It is digested from the thousands of daily inhomogeneous data sources (radiosondes at several hundred stations, surface pressures, aircraft observations, and ship observations) and interpolated on a regular grid of 1,977 points covering the Northern Hemisphere poleward of 30° N latitude. We will focus on the temperature pattern on the 500 millibar (mb) pressure surface (roughly in the middle of the atmosphere, as determined by mass), and the height of the 500mb surface. The latter is of considerable dynamical interest, as it is (approximately) a streamfunction for the upper air flow. This study was made feasible by the availability of the archive on a single 600 megabyte read-only optical disk (CD-ROM) compiled by Professor Clifford Mass of the University of Washington (Seattle). The data analyses can be performed in the comfort of ones' own office with an inexpensive CD-ROM drive attached to a Macintosh II computer.

Although the archive includes data for each day, we will mostly look at the variability of the monthly means. Not only does this keep the computational requirements manageable, it also emphasizes precisely the low-frequency motions that are of primary interest in the effort to probe the feasibility of extended-range forecasting. It is generally agreed that individual storms lose predictability after a week or two. However, if the more ponderous general weather patterns embodied by the monthly means turns out to be governed by a low-dimensional attractor, that would really be news that would shake up the long-range forecasters. To be sure, the monthly mean destroys information on time scales of motion from minutes to weeks that are formally part of the underlying phase space. The premise (or pious hope) is that the *monthly mean rectified fluxes* arising from these motions are nevertheless systematically related to the monthly mean flow itself.

A single 500mb field from the archive is described by a point in a 1,977-dimensional space, and the time series moves in this space. This is not necessarily a phase space of the system, as we have neglected other data (temperature and height at other levels, moisture, etc.) needed to

truly specify the future course of the system; furthermore, the monthly averaging process throws away additional information. Thus, the working space is a projection of the true phase space. Trajectories can cross in the working space, and so there is no guarantee of the existence of a flow determining future evolution from the initial conditions. As is done for the analysis of 1-D climatological time series, we seek to recover some of the true phase space structure by creating synthetic phase spaces of dimension $1,977m$, by taking m-tuples of fields $(Z_1, \ldots, Z_m)$, where each Z_i is a 1,977-point field. The correlation dimension is estimated by taking the Euclidean (ℓ_2) distance between each pair of points in this phase space, and forming the histogram.

Results for the monthly mean 500 millibar height (Z500) are shown in Figure 5 for various values of m. An immediately apparent feature is a subrange at large distances characterized by having slope 1. This indicates that the large-amplitude fluctuations are composed of trajectories that come in one-dimensional bundles. Somewhat smaller amplitude fluctuations appear to be characterized by a three-dimensional structure, and at yet smaller amplitudes the slope increases sharply, suggesting a dimension somewhere in between 20 and 40. This picture shows little sensitivity to the value of m. We have also reproduced similar results for subsets consisting of half the total data.

The 1-D structure at large amplitudes is nothing else than the seasonal cycle. For terrestrial conditions, root mean square differences in Z500 as large as, say, 150 meters occur only between two patterns belonging to different seasons. The results are telling us that the seasonal cycle is dominantly a 1-D entity; to tell what the flow pattern is to lowest order, one only needs to know the time of year. But what of the three-dimensional subrange? The origin of this feature is less apparent.

In order to remove the seasonal effects, we repeated the calculation using data only from the winter months (December, January and February). Results are shown in Figure 6. Not only has the 1-D subrange disappeared, but the 3-D subrange has disappeared as well. We are left only with the high-dimensional structure at short distances; the slight apparent decrease in its slope cannot in any way be regarded as statistically significant, a point about which we shall have more to say shortly. The results indicate that there is a component of the interseaonal variability that can be characterized by three degrees of freedom, one that is not phase-locked to the seasonal cycle. What is this motion? And what cycles of spatial patterns does it correspond to? It is one of the frustrations of the correlation dimension analysis

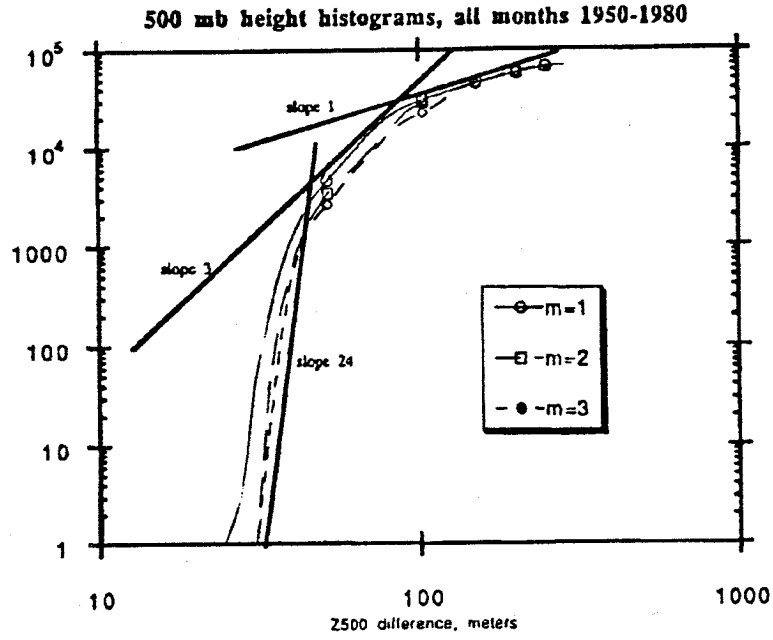

Figure 5. Correlation dimension plot for the 500mb geopotential height field of the atmosphere, using all monthly data.

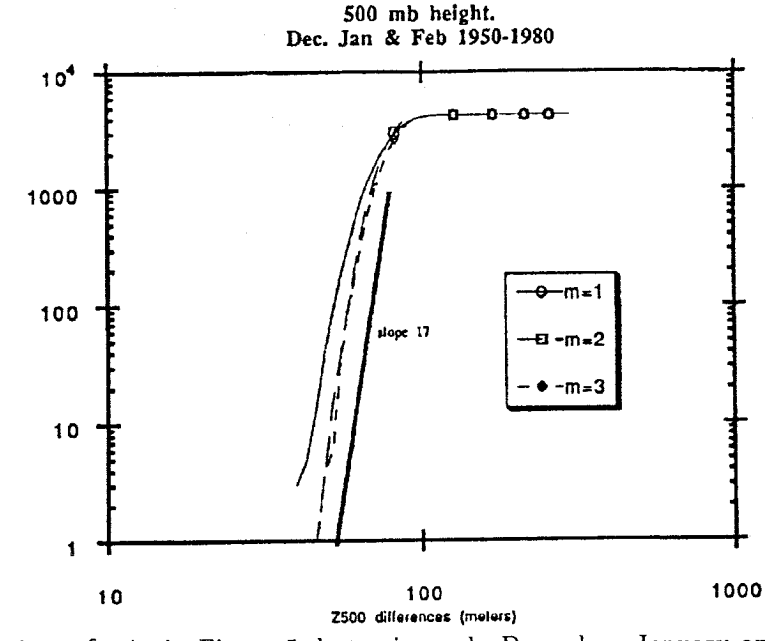

Figure 6. As in Figure 5, but using only December, January and February data.

that it tells us that the motion in infinite-dimensional function space is confined to the vicinity of a three-dimensional manifold, but it gives us no way of constructing this manifold. Specifically, there is (at least locally) a map

$$(\lambda_1, \lambda_2, \lambda_3) \Rightarrow Z(x, y|\lambda_1, \lambda_2, \lambda_3)$$

from triples of real numbers to the function space of all atmospheric flow patterns, which captures a great deal of the atmospheric variability. It would be very useful to know this map, as it would enable the characterization of much of the state of the atmosphere in terms of the specification of the triple of numbers.

An analysis of variability of winter 500mb temperatures is shown in Figure 7. The overall picture is rather similar to that for the height field; there is no indication of a very low-dimensional subrange, and slopes for small temperature differences indicate a dimensionality of around 20. Temperature is governed by very different detailed dynamics from the height field. In particular, it is (in a crude sense) advected around by the streamfunction derived from height data, and for this reason might be expected to show a markedly higher dimensionality. Evidently, the details arising from this process are averaged out in the monthly mean, leaving a temperature variability pattern of similar complexity to that for the height field.

Analysis of daily data would incorporate modes of variability that are filtered out in the monthly mean data set. Since there are 365 height fields per year, and the required number of disk accesses grows as a quadratic function of the number of samples, the rather slow access speed of the optical disk precludes an extensive multi-year study such as we have undertaken for the monthly means (though such a study will be feasible when we acquire enough fast magnetic disk space to upload an appropriate subset of the data). Still it is of interest to see what features can be discerned in the analysis of a single year's height data. Results for the year 1977 are given in Figure 8. Of course, with only one year of data we cannot pick up the 1-D seasonal cycle, as one cannot isolate a periodic signal from the sampling of a single period. Indeed, the 1-D subrange at large distances is not present in this plot. However, there is a distinct 5-dimensional subrange for moderately large distances. Probably three dimensions of this correspond to the interseasonal variability signal detected in the monthly mean analyses. We speculate that the remaining two degrees of freedom are associated with propagating features in the storm tracks. For smaller

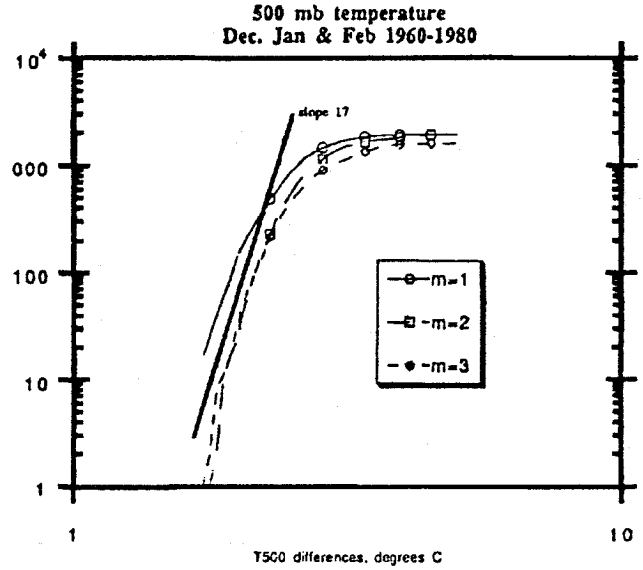

Figure 7. As in Figure 6, but for 500mb temperature fields.

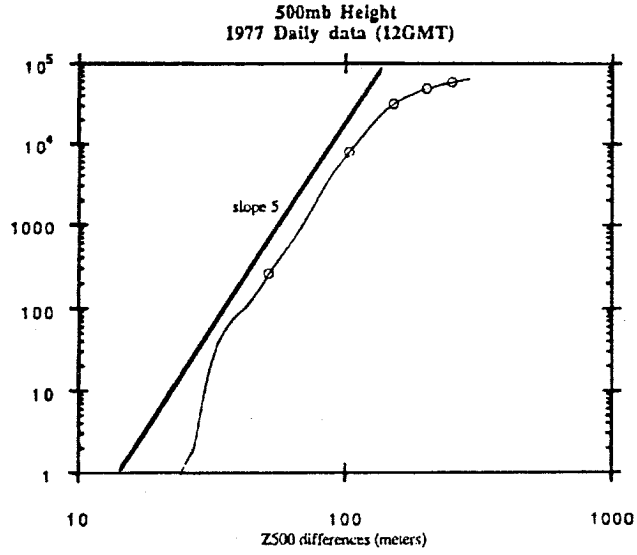

Figure 8. Correlation dimension plot for the 500mb geopotential height based on daily data at 12GMT for the year 1977.

height differences, the slope indicates a high-dimensional subrange similar to that picked up in the monthly mean analyses, suggesting that the higher frequency motion adds little to the spatial complexity of the signal; however, the number of samples in this steep subrange is rather

small, and so drawing conclusions of this nature is perilous.

This leads us to the question of statistical significance of the results. Grassberger (1986) has treated the question of how much data is needed to reliably estimate dimensionality. In a nutshell, the answer is that the data requirements increase precipitously as the underlying dimension increases, since the volume of an N-dimensional sphere increases rapidly with N and one must have enough points to more or less fill up the sphere. To provide a more graphic indication of statistical reliability, which is nonetheless tied to the same idea, we have constructed sequences of random fields of N points each, filled-in using a random number generator. If the random number generator is truly random, the underlying dimensionality of the sequence will be N. We then constructed graphs of the correlation dimension from finite sequences of various lengths n, to see how well we could infer the true dimension. The results are shown in Figure 9. For a 20-dimensional underlying data set, an analysis with 120 points (similar to the size of the sample in our winter analyses) indicates a dimensionality in the 'teens, which indeed is an underestimate. Doubling the amount of data does not improve the situation much; it is not until we go to 10,000 data points and very small distances that we begin to see much improvement. Hence, we are not likely to ever know whether the true slope for the atmospheric data set is 10, 20, or 30. However, 120 point samples from data with underlying dimensionalities of 100 and 2,000 show markedly steeper slopes. From this, we conclude that the dimensionality of the atmospheric data *as present in this data set and for the ranges of difference discussed* is unlikely to be as much as 100. The issue remains as to how much of this apparent low dimensionality is in the real atmosphere, and how much is an artifact of the spatial smoothing and interpolation techniques used to derive the gridded data set from the raw observations. This question can only be answered with recourse to higher resolution data sets derived from the raw data by different methods.

In dealing with the height field (which is practically a streamfunction for mid-latitude flow), we are actually already applying a spatial filter to the underlying dynamics. To specify the future course of a two-dimensional fluid, one needs to know the *vorticity* accurately, not the streamfunction. Similar considerations obtain for potential vorticity in the atmosphere. Since the vorticity ζ and streamfunction ψ are related by $\nabla^2\psi = \zeta$, the streamfunction represents a highly smoothed picture of the vorticity. It is true that knowing the streamfunction *exactly,* one can infer the vorticity exactly. However, small amplitude

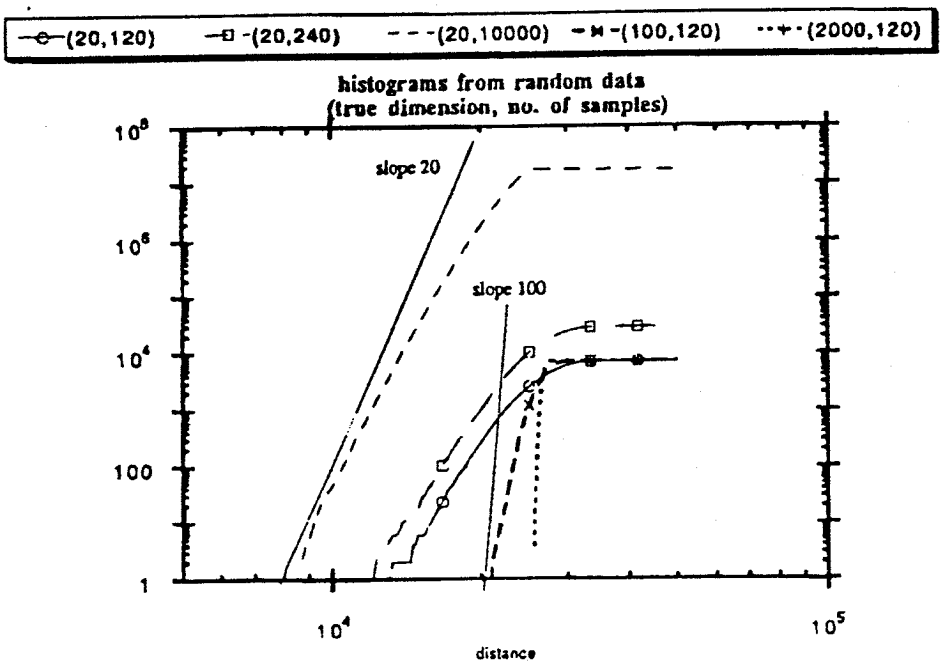

Figure 9. Correlation dimension plot for random fields of various dimensions.

streamfunction errors at small scales destroy the relation. To bring this point home, consider the functions

$$f_1(x) = \cos x, \quad f_2(x) = \cos x + \epsilon \cos\left(\frac{x}{\epsilon}\right)$$

The root-mean-square mean difference between these two functions is small if ϵ is small. However the root-mean-square mean differences of $d^2 f/dx^2$ (which is like vorticity, if one thinks of f as the streamfunction field) diverge as ϵ approaches zero. This phenomenon is a reflection of the fact that in function spaces (unlike finite-dimensional systems) not all norms are equivalent. Distances that approach zero with respect to the streamfunction norm do not approach zero with respect to the vorticity norm. The relevance for the correlation dimension analysis, using data sampled from function spaces, is that we could get a very different dimension from vorticity data than we would from streamfunction data. We would need to have accurate data down to very small distances between streamfunctions to pick up the effects of vorticity fluctuations with small spatial scales. Yet, according to prognostic equations such as Eq. (2.5), fields with differing small-scale

vorticity patterns, but nearly identical streamfunction patterns, will ultimately evolve in time along very different paths. Hence apparent low dimensionality in streamfunction data does not mean that the future course of the system can be accurately predicted by projection onto this low-dimensional manifold.

Another possible source of apparent low dimensionality is filtering of a high-dimensional process by a linear subsystem. Consider the first-order Markov process $A_n = \rho A_{n-1} + \epsilon_n$, where $\rho < 1$ and ϵ_n is uncorrelated white noise. The underlying dimensionality of the system is infinite, in the sense that no finite amount of initial data admits the prediction of the future course of the system over an indefinitely long time period. However, examination of large amplitude fluctuations will generally pick up the one-dimensionality associated with the exponential decay of A. The correlation graph will show a one-dimensional subrange for large A differences, with marked steepening as small distances are approached. However, the apparent one-dimensionality is not entirely without dynamical significance, as it does indeed reflect a certain limited measure of predictability of large amplitude fluctuations. A similar process may occur in the atmosphere, where the linear subsystem could be provided by large-scale, low-frequency Rossby waves.

The low dimensionality of the variability of the height and temperature patterns has some interesting consequences for the reccurence of weather patterns. If the dimensionality of the monthly mean fields were as high as, say, 1,000, one would have to wait $(dZ/dZ_0)^{1,000}$ months to see two fields that differed by dZ, where dZ_0 is some reference value. If this were the case, each weather pattern would for all intents and purposes be unique. With a dimensionality of 20 or so, however, recurrence of similar patterns would be rather common. In fact, long before fractal geometry had been invented, Lorenz (1969) studied recurrence in atmospheric flow. His intent was to analyze atmospheric predictability by first searching for similar initial states in the historical record, and then examining the divergence of their subsequent evolution. Lorentz's attempt was only partially successful, for reasons evident in Figures 5–8—close analogues, while not exactly nonexistent, are nonetheless rather rare. The entire 40-year historical record yields only a score or so of analogues close enough to be useful in predictability studies. Even if the dimension did not increase beyond 20 at small distances, halving the threshold dZ at which we are willing to accept two flows as analogues still would mean waiting 2^{20} times as many months (or about 40 million years) for suitable analogues to appear.

4. *Genesis of Spatial Complexity: Chaotic Mixing*

Regardless of the implications (or lack thereof) of the results of §3 for the presence of low-dimensional chaos in the atmosphere, it's clear that the space of all streamline patterns visited by points of the atmospheric attractor can, with considerable accuracy, be characterized by a rather small number of coordinates. On the other hand, observations of tracers such as potential vorticity (e.g., McIntyre and Palmer 1983) typically show a great deal of spatial complexity. There is no contradiction in this state of affairs. Recent results on mixing by two-dimensional, time-dependent flows show that the particle trajectories can be chaotic and exhibit predictability decay even if the velocity field moving the particles around is perfectly predictable—in fact, even if it is periodic in time (Ottino et al. 1988). In this section, we analyze an example of this kind of mixing in an idealized atmospheric flow. The correlation-dimension analysis will reappear here as a measure of the spatial complexity of the tracer field.

Consider the two-dimensional incompressible velocity field defined by the streamfunction

$$\psi = A\cos(k_1(x - c_1t))\sin(l_1y) + \epsilon\cos(k_2(x - c_2t))\sin(l_2y) \qquad (4.1)$$

This represents a primary traveling wave with amplitude A and phase speed c_1, disturbed by a perturbation with speed c_2 and (small) amplitude ϵ. Most of the results shown below are insensitive to the specific kind of waves chosen, but for the sake of concreteness let's consider solutions of the barotropic β-plane equation. This equation is a crude approximation to planetary-scale flows in a shallow atmosphere, and consists of Eq. (2.5) with $g = d = 0$, the vorticity being replaced by the potential vorticity $\nabla^2\psi + \beta y$. It is essentially a tangent-plane approximation to incompressible 2-D flow on the surface of the sphere, with x being approximately the longitude and y being approximately the latitude. The waves in this system are known as Rossby waves, and satisfy the dispersion relation

$$c = \frac{-\beta k}{k^2 + l^2} \qquad (4.2)$$

when the equations are linearized about a state of rest. In fact, with c_1 given by Eq. (4.2), Eq. (4.1) is an exact nonlinear solution of the equations for arbitrary A provided $\epsilon = 0$. The quantity c_2 is also taken to be governed by the Rossby wave dispersion relation, though there

is no particular mathematical justification in doing so. The mixing results are governed primarily by the streamline geometry of the unperturbed flow and the overall magnitude of the perturbation, so this crude assumption on the form of the perturbing wave isn't particularly dangerous.

In the comoving frame defined by $X = x - c_1 t$, the velocity field has the streamfunction

$$\psi = A\cos(k_1 X)\sin(l_1 y) + c_1 y + \epsilon\cos(k_2 X - k_2(c_2 - c_1)t)\sin(l_2 y) \quad (4.3)$$

and corresponds to a steady velocity field disturbed by a perturbation that's periodic in time, the Poincaré period being $T = 2\pi/(k_2(c_2 - c_1))$. The problem of particle motion in this flow field is a perturbed planar Hamiltonian system. Figure 10 shows the unperturbed streamlines in the comoving frame for $A = 1$. The stagnation points P and P' form what is known as a "heteroclinic cycle," being connected by the upper arc PP' (the unstable manifold of P and the stable manifold of P') and the lower segment $P'P$ (the unstable manifold of P' and the stable manifold of P). The latter is preserved under perturbation, but powerful theorems of dynamical systems theory imply that the former breaks up generically into a chaotic set. It is the coexistence of closed and open streamlines that is at the heart of the chaos, leading to the inevitability of chaotic mixing. Indeed, Knobloch and Weiss (1987) have presented evidence of chaotic mixing in thermosolutal convection, which involves modulated traveling waves having an identical streamline geometry to that arising in the Rossby wave case.

In Figure 11 we show the time evolution of an initially small cloud of particles placed near the bottom of the recirculating eddy at the lower left of Figure 10. The particles were moved around in the velocity field defined by Eq. (4.3), taking $A = 1$ and $\epsilon = 0.15$. The number of time periods T that have passed are indicated on each frame. The cloud is stretched out rather gently for a time, but then "finds" the unstable manifold and expands in a burst. Thereafter, the mixing becomes roughly diffusive, displaying rapid stretching and folding and efficient cross-streamline mixing. The final state of a mixed cloud of 10,000 particles is plotted in Figure 10, in order to show the intimate association between the mixed region and the heteroclinic streamline structure.

It is common practice to characterize observed or simulated concentration distributions by their Fourier power spectra. What is the power spectrum of a distribution such as shown in Figure 10? An el-

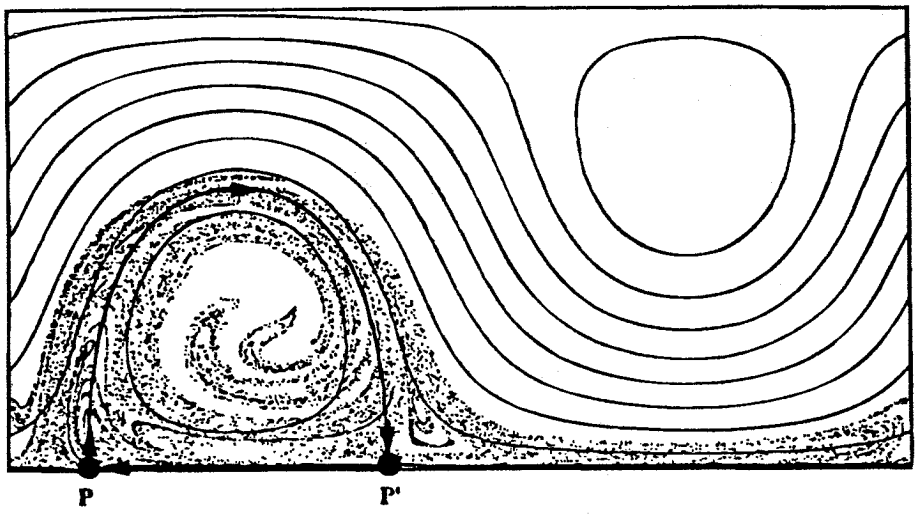

Figure 10. Streamline pattern of the primary Rossby wave in the comoving reference frame. The domain goes from 0 to $2p$ in x and from 0 to p in y. The cloud of points illustrates the dispersal of an initially small cloud of particles when they are advected by a perturbed Rossby wave with $e = 0.25$.

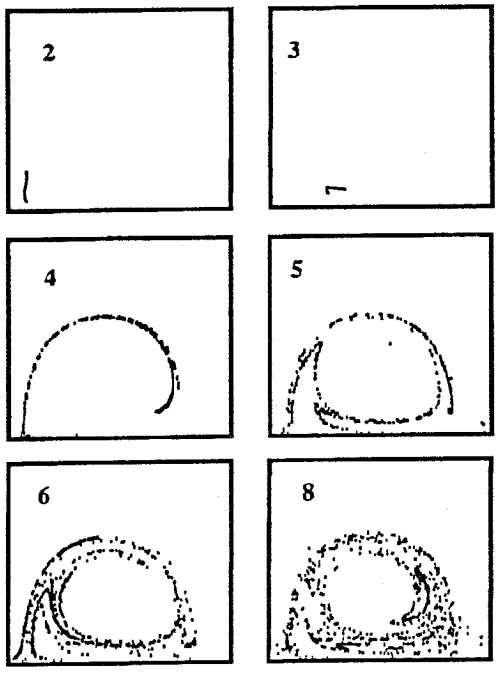

Figure 11. Time sequence of particle dispersion.

ementary argument reveals a direct relation between spectra and the correlation dimension of the particle cloud.

Consider a concentration pattern made up of an assemblage of N circular blobs located at points $\mathbf{r}_j$ in the two-dimensional plane, with $j = 1, \ldots, N$. Each blob contributes a concentration $G(|\mathbf{r} - \mathbf{r}_j|)$, where $G(0) = 1$. Thus, the concentration is

$$C(\mathbf{r}) = \sum_{j=1}^{N} G(|\mathbf{r} - \mathbf{r}_j|)) \tag{4.4}$$

and its 2-D Fourier transform is

$$c(\mathbf{k}) = \frac{g(|\mathbf{k}|)}{2\pi} \sum_{j=1}^{N} e^{i\mathbf{k}\cdot\mathbf{r}_j} \tag{4.5}$$

where $g(\cdot)$ is the Fourier transform of $G(\cdot)$. Thus, the concentration variance spectrum is

$$|c(\mathbf{k})|^2 = \frac{|g(|\mathbf{k}|)|^2}{2\pi} \sum_{j,j'=1}^{N} e^{i\mathbf{k}\cdot(\mathbf{r}_j - \mathbf{r}_{j'})} \tag{4.6}$$

Now suppose that there are a large number of particles, and that the distribution of interparticle distances is governed by a probability density $P(\delta\mathbf{r}) = P(\delta x, \delta y)$, where $\delta\mathbf{r} = (\delta x, \delta y)$ is the distance between particles. Then the number of particle pairs separated by vector distance $\delta\mathbf{r}$ is approximately $N^2 P(\delta\mathbf{r})$, whence Eq. (4.6) becomes

$$|c(\mathbf{k})|^2 = N^2 \frac{|g(|\mathbf{k}|)|^2}{2\pi} \iint P(\mathbf{r}) e^{i\mathbf{k}\cdot\mathbf{r}}\, dx dy \tag{4.7}$$

in which the double integral extends over all space. For wavelengths much longer than the individual blob radius, $g = 1$. Thus, the concentration variance spectrum is proportional to the Fourier transform of the probability distribution of the interparticle distance. If we further assume the probability distribution to be isotropic, then $P(\mathbf{r}) = P(r)$, where $r = |\mathbf{r}|$, whence, upon integrating over angles θ, Eq. (4.7) becomes

$$|c(\mathbf{k})|^2 = \text{constant} \cdot \int_0^\infty r P(r) J_0(kr)\, dr \tag{4.8}$$

where J_0 is the zero-order Bessel function of the first kind. This gives us the desired connection with the correlation dimension, since

$rP(r) = dH/dr$, H being the cumulative histogram we have been dealing with in our dimension computations. If there is an extensive range of distances for which $H = r^d$, then letting (kr) be the integration variable in Eq. (4.8), we find

$$C(k) = k|c(\mathbf{k})|^2 = \text{constant} \cdot k^{1-d} \tag{4.9}$$

where $C(k)$ is the isotropic power spectrum as conventionally defined.

As the concentration pattern is embedded in the two-dimensional plane, we must have $0 < d < 2$. The case $d = 0$ corresponds to clustering of the concentration into a few isolated pointlike regions, and yields a power spectrum that increases toward short waves. The situation $d = 1$ corresponds to well-separated strands, and yields a flat spectrum. The steepest possible spectrum is k^{-1}, corresponding to mixing of the tracer over a two-dimensional region. Steeper spectra cannot be associated with self-similar fractal concentration patterns, since there is not enough variation at small scales. A more traditional way of looking at this derives from the fact that k^{-1} has a divergent integral from (any) k_0 to ∞, corresponding to an infinite total concentration variance. Steeper spectra, however, have convergent variance, whence one can define the integral scale

$$L^2 = \frac{\int_{k_0}^{\infty} k^{-2} C \, dk}{\int_{k_0}^{\infty} C \, dk} \tag{4.10}$$

This defines a preferred scale for structures in the concentration pattern. Obviously, such distributions cannot exhibit self-similarity.

It is a simple matter to apply this calculation to particle dispersion data from the Rossby wave problem. A computation of the correlation dimension for a cloud of particles like that shown in Figure 10, but recomputed with 40,000 particles, shows an extensive range of scales characterized by $d = 2$ (or perhaps a bit less). This implies a subrange with a k^{-1} spectrum. In Figure 12 we show the spectrum, calculated from a Bessel transform of the histogram. It indeed shows the predicted spectrum for the intermediate length scales. This spectrum is the same as the classical Bachelor spectrum predicted for tracers in homogeneous, isotropic, two-dimensional turbulence, on the basis of scaling arguments (e.g., see Rhines 1979).

How universal is this result? All parameter settings we have tried in the Rossby wave case ultimately lead to a mixed cloud characterized by $d = 2$. We have also found $d = 2$ for chaotic mixing by the

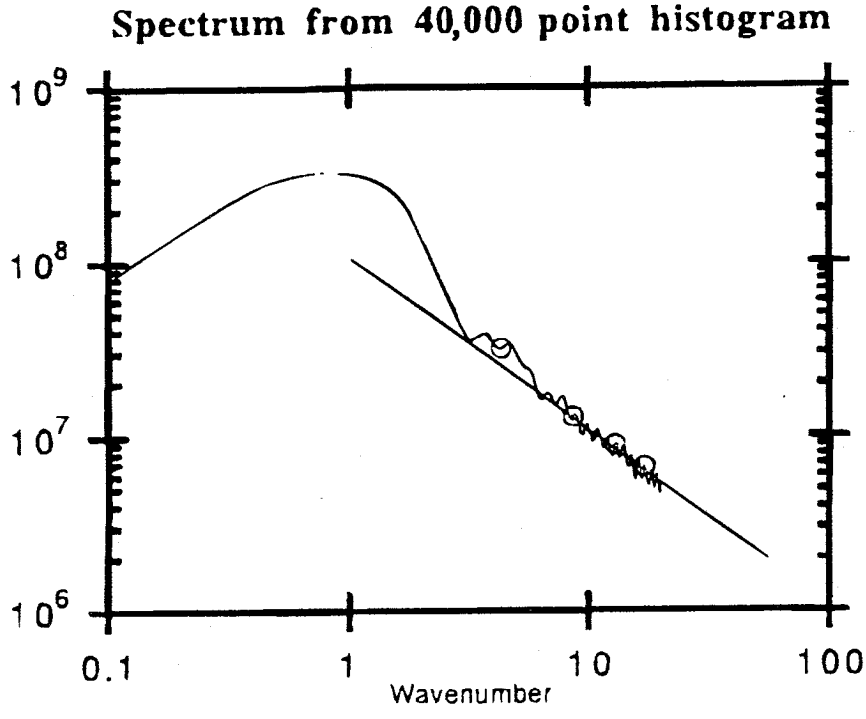

Figure 12. Power spectrum of the concentration distribution in Figure 10, computed by Bessel transform of the histogram.

Kida vortex (Polvani and Pierrehumbert 1989). Visual analysis of the numerous cases considered by Ottino and others reveals no obvious contradictions to the $d = 2$ law. While it seems plausible that chaotic mixing should always proceed to $d = 2$, we have been unable to produce a proof of this conjecture. If we take this conjecture as true, however, it implies an ubiquity for the Bachelor spectrum that transcends the restrictive assumptions of the original scaling arguments. The k^{-1} spectrum would arise naturally in any flow with fluctuating large eddies, regardless of whether the eddy field is itself homogeneous or isotropic. The spectrum can be obtained as a result of advection *by the large scale component of the flow field alone.* This is a dramatic illustration of the well-known nonlocality of two-dimensional turbulence in spectral space. The concentration spectrum at length scale L (no matter how small) is governed by the velocity field of the large eddies, rather than by advection of eddies having scales near L.

Our results on concentration spectra have some implications for the dynamics of two-dimensional turbulence as well. Vorticity (or more generally potential vorticity), is a tracer in an inviscid fluid, as it is con-

served following a fluid parcel. The power spectrum of vorticity is distinguished by the name "enstrophy spectrum." Owing to the relation between vorticity and velocity, a $k^{-\alpha}$ enstrophy spectrum corresponds to a $k^{-(\alpha+2)}$ energy spectrum. The same scaling arguments that predict a k^{-1} concentration power spectrum in homogeneous, isotropic 2-D turbulence yield a k^{-1} enstrophy spectrum and corresponding k^{-3} energy spectrum. Assuming vorticity acts qualitatively like a passive tracer, the chaotic mixing properties imply that the k^{-1} enstrophy spectrum can arise in much more general circumstances. How does vorticity differ from a passive tracer?

Unlike a passive tracer, it changes the velocity field. From our standpoint, the most important manifestation of this is that as a circular vorticity blob is stretched out into an ellipse, its self-induced velocity causes it to rotate or roll up, both of which act to hold the blob together. Thus, we expect that isolated strong vortices will survive, while sufficiently weak vortices will be dispersed in the manner of a passive tracer. Chaotic mixing is generic in the sense that it is not particularly sensitive to the details of the fluctuating velocity field. This fact leads us to conjecture that the strength of "weak" vortices need not be *infinitesimal* in order to be dispersed like a passive tracer. It should be sufficient that they not be strong enough to hold together.

For many years the k^{-1} shortwave enstrophy spectrum was assumed to be a fact of life in 2-D turbulence. However, calculations by McWilliams (1984) and many others since tend to show a steeper spectrum. The steepness is generally associated with the freezing out of isolated eddies that have individually smooth vorticity distributions. The incompatibility of this situation with self similarity is in accord with the analysis presented above. When the tracer in question is vorticity, the integral scale of Eq. (4.10) represents a characteristic vortex size. Since the k^{-1} spectrum appears to be an inevitable consequence of chaotic dispersal, the steeper spectra are expected to occur only when the initial condition contains small eddies intense enough to resist being torn apart by the large-scale flow field. This hypothesis is currently being tested in fully nonlinear simulations, the results of which will be reported elsewhere.

The above considerations apply to the small-scale spectrum of two-dimensional turbulence. At large scales, two-dimensional turbulence exhibits a different universal spectrum, characterized by a $k^{-5/3}$ energy spectrum and a flux of energy from small scales to large scales. It is interesting to note that this spectrum is also compatible with a fractal vorticity structure. The energy spectrum implies a $k^{1/3}$ enstrophy spec-

trum, which yields a dimensionality $d = \frac{2}{3}$. Since $d = 0$ corresponds to "vortex clumps" and $d = 1$ corresponds to "vortex strands," the indicated geometry of the large-scale vorticity structures is "clumpy strands." The chaotic mixing results provide some dynamical understanding of the reasons for the small-scale vorticity geometry, but we do not yet have analogous explanation for the large-scale geometry.

5. The Predictability of Predictability: Fractal Structure of Lyapunov Exponents

We have seen that the trajectories associated with deterministic flow fields can exhibit chaos and loss of predictability. The mixing seen in Figure 10 is *prima facia* evidence that, given a small initial error in the position of a particle, after a sufficiently long period of time it will be more-or-less equally probable to find the particle anywhere within a broad two-dimensional region. Can we at least predict the rate at which predictability will be lost? If, for example, a radioactive cloud is released over Chernobyl, can we say with confidence how long its trajectory can be accurately predicted? In this section we take up this question, using the system described in §4 as an example.

The basis of this analysis is the computation of a variant of the *Lyapunov exponent.* This number measures the exponential growth rate of the separation of neighboring trajectories. To compute the Lyapunov exponent, we take two initial conditions separated by a distance δ and integrate the corresponding trajectories forward in time for a long period τ. The distance $d(t)$ between the trajectories is computed, with the (most unstable) Lyapunov exponent then being given by $\log d/\tau$. For a two-dimensional system, such as the Rossby wave example, there are only two Lyapunov exponents; one is positive and the other negative, their sum being zero, in accordance with the incompressibility of the velocity field. A technical detail of the calculation is that we actually use *three* trajectories, taking the maximum separation in order to ensure against the unlikely possibility of the initial condition projecting exactly onto the contracting direction rather than that of expansion.

The true Lyapunov exponent for a given trajectory is obtained in the joint limit as $\delta \to 0$ and $\tau \to \infty$. Of course, this is an impossible limiting operation to carry out numerically. Instead we deal with estimates of λ for finite times τ, which are in any case of greater physical interest since we're limited to only a finite time in which to observe a particular experiment anyway. It's generally assumed that any trajectory within the chaotic zone will yield the same Lyapunov exponent. But we take nothing for granted, especially since we are computing only

a finite-time estimate. Hence, we cover the region of interest with a grid of initial conditions, and compute λ for trajectories starting from each point on the grid. In the subsequent discussion, we confine our attention to the finite-time results, and do not attempt to treat the formidable question of what happens as $\tau \to \infty$.

The spatial distribution of the finite-time estimates of the Lyapunov exponent are shown at three different scales in Figure 13,

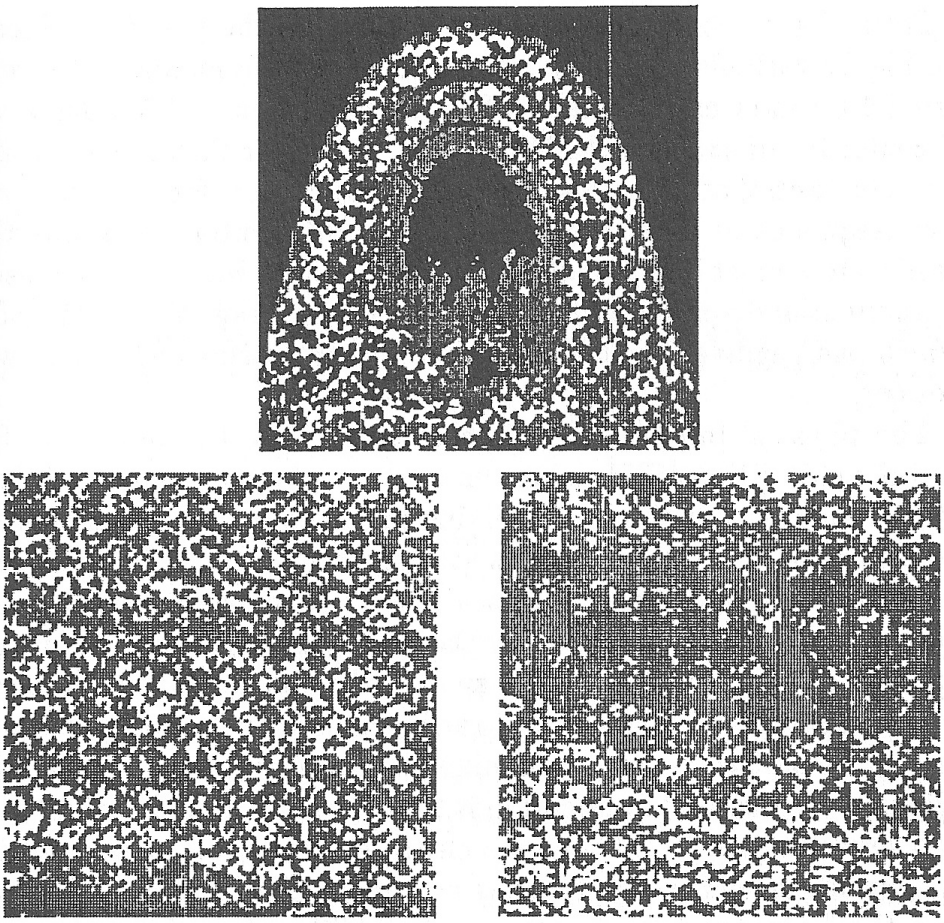

Figure 13. Spatial distribution of the Lyapunov exponents based on 50 Poincaré time periods. Black regions are nonchaotic ($\lambda = 0$), grey regions correspond to values of λ between 0 and the median of all positive values, and white regions correspond to above median values. The top panel shows the macroscopic structure of the chaotic zone, and covers the domain $x \in [0, \pi]$, $y \in [0, 2]$. The lower-left panel zooms in by roughly a factor of 10, covering the domain $x \in [1.520795, 1.620795]$, $y \in [0.2, 0.3]$. The lower-right panel zooms in by yet another factor of one thousand, covering the domain $x \in [1.570695, 1.570895]$, $y \in [0.2499, 0.2501]$.

computed for $\epsilon = 0.15$ and $\tau = 50T$. This analysis makes it easy to discern the macroscopic structure of the chaotic zone (see top panel), as non-chaotic regions are characterized by $\lambda = 0$. A striking feature of the distribution is that the boundary of the chaotic zone is not a smooth curve, but has the typical coastline-like appearance of a fractal. We leave it to the reader to decide whether or not this is surprising. Moreover, there is a great deal of fine-grained variability of λ within the chaotic zone. The lower panels of Figure 13 show that this variablity is present over at least four orders of magnitude of spatial scales.

From Figure 13, it's difficult to tell whether the positive values of λ are highly variable or sharply clustered about the median. Hence, in Figure 14 we show cross sections of actual pixel values of λ along a vertical centerline of each panel of Figure 13. It is clear that the variability of λ is substantial at all scales considered. Of course, for any fixed τ, it is a consequence of regularity theorems for differential equations that the estimate λ must be a smooth function of space. However, because of the exponential divergence of neighboring trajectories, the spatial scale at which one begins to see the smoothness decreases exponentially with increasing τ.

The physical implications of these results are twofold. The first concerns mixing times, a crucially important physical characteristic of the flow. Figure 11 showed that mixing proceeds in two stages, becoming diffusive only after the blob is stretched out over a distance comparable to the scale of the large-scale flow. The Lyapunov exponent approximates the time occupied by the first stage of mixing, our results suggesting a fractal pattern in the spatial distribution of Lyapunov exponents. Thus, in the limit of arbitrarily small initial blobs, the time it takes for the blob to become well mixed is arbitrarily sensitive to initial conditions. The second implication is that the rate of predictability loss is not very predictable. There are clearly macroscopic regions of tori in Figure 13, within which there is no predictability decay. However, within the fractal region of λ, the predictability of an given trajectory can differ wildly from that of its nearby neighbor.

6. Conclusions

We have given several examples here illustrating the power of the correlation dimension analysis as a tool for probing the patterns of variability in the atmosphere and in atmospheric models. Along the way, we have encountered a number of interesting dynamical systems, not least the atmosphere itself. The analysis picks up an expected one-dimensional component of atmospheric variability associated with the periodic sea-

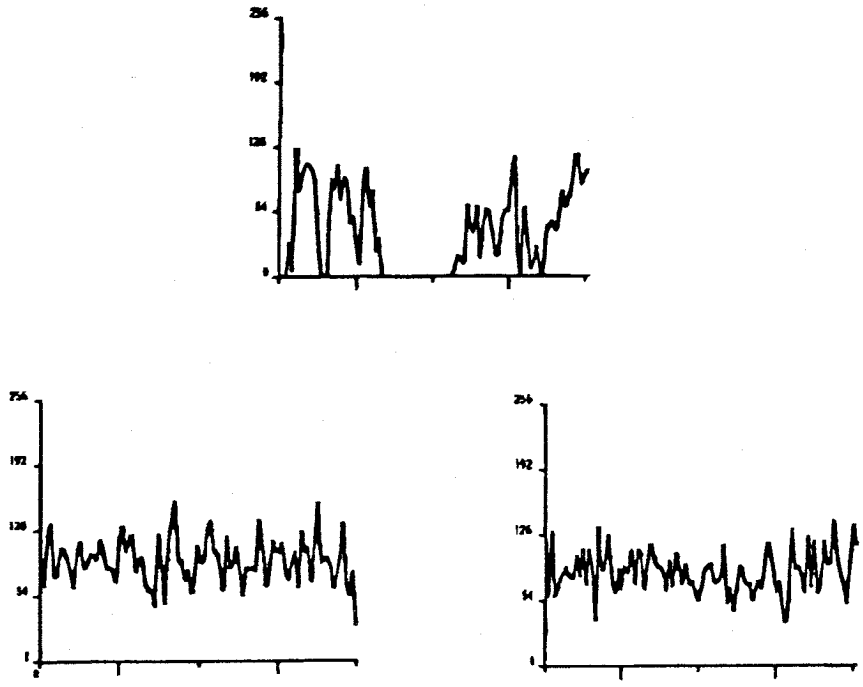

Figure 14. Cross section of values of λ along a vertical line at the centers of the three panels of Figure 13.

sonal cycle. For somewhat smaller amplitudes of fluctuation, it also reveals an additional two degrees of freedom linked to the seasonal cycle. The nature of this motion is obscure, and the next challenge will be to find some way to isolate the sequence of spatial patterns associated with it. For still smaller amplitudes of fluctuations, the indicated dimensionality is perhaps 20, and certainly less than 100. This regime is not tied to the seasonal cycle. A much longer record would be necessary to determine whether and how fast the dimensionality increases as the amplitude approaches zero. It would also be illuminating to carry out the analysis for higher-resolution data sets from computer models of the atmosphere's general circulation.

We are reluctant to conclude on the basis of these results that the atmospheric evolution could be reproduced with a model containing a few dozen variables (even assuming one had some way of identifying them.) Our analysis does not convincingly demonstrate that the embedding space we have dealt with is indeed a phase space. Considerations of vorticity versus streamfunction suggest that, on the micro-

scopic level, it is indeed unlikely to be a phase space. The picture that emerges of the geometry of the atmospheric variability is the following: When we stand far back in the multimillion-dimensional embedding space of the atmospheric system, the trajectory appears to be confined to a hypersurface of rather low dimensionality (let us say it is 20). When we get closer to this hypersurface and take out our magnifying glass, we see that it is actually thick in several thousand orthogonal directions, and that the detailed course of the system trajectory depends not just on the projection onto the hypersurface, but also on knowlege of the value of these additional variables. This leads to several intriguing questions: Can one define a flow on the 20-dimensional hypersurface whose trajectories approximate those of the true system? As a function of time, what is the uncertainty in the future course of the true system, given knowlege only of the projection on the 20-dimensional hypersurface? And, finally, can one define a flow that at least reproduces the *climate,* i.e., the statistical behavior, of the true system projected on the hypersurface?

The correlation dimension analysis does, however, show that the state of the atmosphere can, in principal, be specified with considerable accuracy in terms of a few dozen variables. In order to carry out this data compression, one again comes up against the need to find a way to determine the *modes* of atmospheric variability, rather than just its dimensionality.

We have seen further that the rather low dimensionality characterizing the large-amplitude fluctuations of the wind fields is not incompatible with the observed spatial complexity of various conserved tracers in the atmosphere. An example was presented showing that even a precisely periodic large-scale flow field can produce spatial chaos in the tracer field it advects around. These tracer distributions on the plane were found to be characterized by a correlation dimension of very nearly 2. Further analysis revealed a simple connection between the correlation dimension of the spatial pattern and power spectrum of the concentration. This connection allows one to rephrase in geometric terms the search for universal spectra of concentration or vorticity in two-dimensional turbulence. It provides some insight into the prevalence of the classical spectra and the circumstances leading to deviations from the classical spectra.

Finally, there are some striking indications that the very rate at which trajectories lose their predictability is unpredictable, assuming there are even small errors in the initial point of the trajectory. The number characterizing predictability appears to be distributed in a spa-

tially complex manner. We have made no attempt to formally characterize its variability. The fractal tools employed earlier in this essay were useful primarily for characterizing the geometry of objects. In the "predictability of predictability" problem we are dealing instead with a map assigning a real number to each point in space—or more precisely a measure assigning the average of the quantity to each small subset of the space—so the appropriate tool for effecting the characterization is multifractal analysis.

References

Essex, C., Lookman, T., and Nerenberg, M. A. H. 1987: "The Climate Attractor over Short Timescales," *Nature,* 326, 64–66.

Grassberger, P. 1986: "Do Climatic Attractors Exist?" *Nature,* 323, 609–612.

Grassberger, P. and Procaccia, I. 1983: "Measuring the Strangeness of Strange Attractors," *Physica D,* 9D, 189–208.

Knobloch, E. and Weiss, J. B. 1987: "Chaotic Advection by Modulated Traveling Waves," *Phys. Rev. A,* 36, 1522–1524.

Lorenz, E. N. 1969: "Atmospheric Predictability as Revealed by Naturally Occurring Analogues," *J. Atmos. Sci.,* 26, 636–646.

Maasch, K. A. 1989: "Calculating Climate Attractor Dimension from $\delta^{18}O$ Records by the Grassberger-Procaccia Algorithm," *Climate Dynamics,* 4, 45–55.

McIntyre, M. E. and Palmer, T. N. 1983: "Breaking Planetary Waves in the Stratosphere," *Nature,* 305, 593–600.

McWilliams, J. 1984: "The Emergence of Isolated Coherent Vortices in Turbulent Flow," *J. Fluid Mech.,* 146, 21–43.

Ottino, J. M., Leong, C. W., Rising, H. and Swanson, P. D. 1988: "Morphological Structures Produced by Mixing in Chaotic Flows," *Nature,* 333, 419–425.

Pedlosky, J. 1979: *Geophysical Fluid Dynamics,* Springer-Verlag, New York.

Polvani, L. and Pierrehumbert, R. T. 1989: "Characterization of Mixing by the Kida Vortex," Paper presented at the 1989 American Physical Society Symposium on Fluid Mechanics, Anaheim, CA.

Rhines, P. B. 1979: "Geostrophic Turbulence," *Ann. Rev. Fluid Mech.,* 11, 401–442.

Tsonis, A. A. and Elsner, J. B. 1988: "The Weather Attractor Over Very Short Timescales," *Nature,* 333, 545–547.