1
Introduction

The subject of this book is chaos as seen through the filter of topology. The origin of this book lies in the analysis of data generated by a dynamical system operating in a chaotic regime. Throughout this book we develop topological tools for analyzing chaotic data and then show how they are applied to experimental data sets.

More specifically, we describe how to extract, from chaotic data, topological signatures that determine the stretching and squeezing mechanisms that act on flows in phase space and that are responsible for generating chaotic data.

In the first section of this introductory chapter we very briefly review some of the basic ideas from the field of nonlinear dynamics and chaos. This is done to make the work as self-contained as possible. More in-depth treatment of these ideas can be found in the references provided.

In the second section we describe, for purposes of motivation, a laser that has been operated under conditions in which it behaved chaotically. The topological methods of analysis that we describe in this book were developed in response to the challenge of analyzing chaotic data sets generated by this laser.

In the third section we list a number of questions we would like to be able to answer when analyzing a chaotic signal. None of these questions can be addressed by the older tools for analyzing chaotic data. The older methods involve estimates of the spectrum of Lyapunov exponents and estimates of the spectrum of fractal dimensions. The question that we would particularly like to be able to answer is this: How does one model the dynamics? To answer this question we must determine the stretching and squeezing mechanisms that operate together – repeatedly – to generate chaotic data. The stretching mechanism is responsible for *sensitivity to initial conditions* while the squeezing mechanism is responsible for *recurrent nonperiodic behavior*. These two mechanisms operate repeatedly to generate a strange attractor with a self-similar structure.

A new analysis method, topological analysis, has been developed to respond to the fundamental question just stated [7, 8]. At the present time this method is suitable only for strange attractors that can be embedded in three-dimensional spaces. However, for such strange attractors it offers a complete and satisfying resolution to this question. The results are previewed in the fourth section of this chapter. In the final section we provide a brief overview of the organization of this book. In particular, we summarize the organization and content of the following chapters.

The Topology of Chaos, Second Edition. Robert Gilmore, Marc Lefranc

It is astonishing that the topological analysis tools that we describe have provided answers to more questions than we had originally asked. This analysis procedure has also raised more questions than we have answered. We hope that the interaction between experiment and theory and between old questions answered and new questions raised will hasten the evolution of the field of nonlinear dynamics.

1.1 Brief Review of Useful Concepts

There are a number of texts that can serve as excellent introductions to the study of nonlinear dynamics and chaos. These include [9–21]. Any one of these can be used to fill in details that we may pass by a little too quickly in our study. There are also many texts that serve as introductions to topology. All cover far more material than we use here. As a result, we do not recommend that a reader invest time in any one of these texts. We will introduce the topological concepts as needed as we proceed.

For now we review very briefly some of the foundational ideas of chaos.

What Is chaos? We take the following as a useful definition for chaos, or chaotic motion. Chaos is motion that is

1. deterministic
2. bounded
3. nonperiodic
4. sensitive to initial conditions.

Where does this motion take place? It is useful to describe the state of a physical system by a set of coordinates. The most convenient way to do this is to establish a phase space. A point in the phase space describes the state of the physical system. The "motion" described above is that of a point in the phase space. For example, the phase space needed to describe the motion of two particles in a plane is eight-dimensional: two coordinates and two velocity components are required to describe the state of each particle's motion. We will work with smaller phase spaces.

What is "deterministic motion"? The motion of the coordinates x_i of a point in an N-dimensional phase space is governed by a set of N first-order ordinary differential equations:

$$\frac{dx_i}{dt} = f_i(x_1, x_2, \ldots, x_N; c). \tag{1.1}$$

The coordinates $x \in R^N$ are called *phase space coordinates* and the coordinates $c \in R^K$ are called *control parameters*. A set of N equations of this type is called a *dynamical system*.

What is "bounded" motion? The trajectory defined by the equations of motion (1.1) can be parameterized by the time coordinate: $x(t)$. *Bounded* means that the maximum distance between any two points on the trajectory over all times is less than infinity: $\mathrm{Max}_{(t_1,t_2)}||x(t_1) - x(t_2)|| < \infty$.

What is periodic motion? This is motion that returns to its starting point after an elapsed time T: $x(t) = x(t + T)$, for some $T > 0$. T is called the *period* if there is no smaller positive value (for example, $T/2$) for which this equation is true.

What is "sensitivity to initial conditions"? A point x_1 in phase space can serve as an initial condition for a trajectory through it: $x_1(t)$. Two nearby points can serve as initial conditions for two trajectories, one starting at each point. At first the two trajectories remain close to each other. If the distance between them grows exponentially with time, $||x_2(t) - x_1(t)|| \simeq e^{\lambda t}$ with $\lambda > 0$, then the system is said to exhibit sensitivity to initial conditions. This means that, although the evolution is deterministic, the future position of an initial condition becomes unpredictable after some time. The term λ is called a *Lyapunov exponent*.

How do you visualize chaotic motion? A very convenient way to do this is to view the trajectory of an initial condition in the phase space. It often happens that the trajectory "wanders around" in the phase space until it "settles down" onto, or is "attracted to," a reasonably well-defined structure. This structure, when it exists, is called by mathematicians an Ω limit set and by physicists a strange attractor or a chaotic attractor. Plots of strange attractors appear liberally throughout this work: for example, see Figures 6.10 and 7.6.

A strange attractor contains no periodic orbits. However, buried in the strange attractor lies a host of unstable periodic orbits. It is these orbits that form the foundation of the topological analysis methods that are presented in this work. Good approximates to these orbits can be extracted from experimental data.

How do you search for chaotic behavior? A simple, very effective way to do this is to study the motion at a sequence of values of some control parameter. A suitable tool is called a bifurcation diagram. This amounts to a plot of one of the phase space coordinates as a function of one of the control parameters. Phase space plots of the logistic map $x' = \lambda x(1 - x)$ are presented in Figures 2.3 and 2.4. A practiced eye (training time $\sim$ 2 s) can easily distinguish chaotic from nonchaotic behavior.

Is there a "smoking gun" for chaotic behavior? No. The original discovery [22] that launched a thousand studies was that a period-doubling cascade was a prelude to chaos and that a number of invariants were associated with such cascades (see Preface). The possibility of confirming these predictions, or showing that they were not correct, attracted a number of experimentalists into this field (cf. [10]). Having said that, not all "routes to chaos" go by the period-doubling pathway.

What topological tools will be used? We will study how the unstable periodic orbits that exist in plenty in a strange attractor are organized. To do this we will rely on the Gauss linking number of a pair of closed orbits as well as a closely related idea (relative rotation rates). We will also study braids and introduce cardboard-type structures that serve to hold all the periodic orbits in a strange attractor in a very simple way. Toward the end of this work we will introduce more venerable topological ideas: Euler characteristic, genus, isotopy, and so on.

1.2 Laser with Modulated Losses

The possibility of observing chaos in lasers was originally demonstrated by Arecchi *et al*. [23] and by Gioggia and Abraham [24]. The use of lasers as a testbed for generating deterministic chaotic signals has two major advantages over fluid and chemical systems, which until that time had been the principal sources for chaotic data:

1. Time scales intrinsic to a laser (10^{-7}–10^{-3} s) are much shorter than time scales in fluid experiments and oscillating chemical reactions. This is important for experimentalists, since it is possible to explore a very large parameter range during a relatively short time.
2. Reliable laser models exist in terms of a small number of ordinary differential equations whose solutions show close qualitative similarity to the behavior of the lasers that are modeled [25, 26].

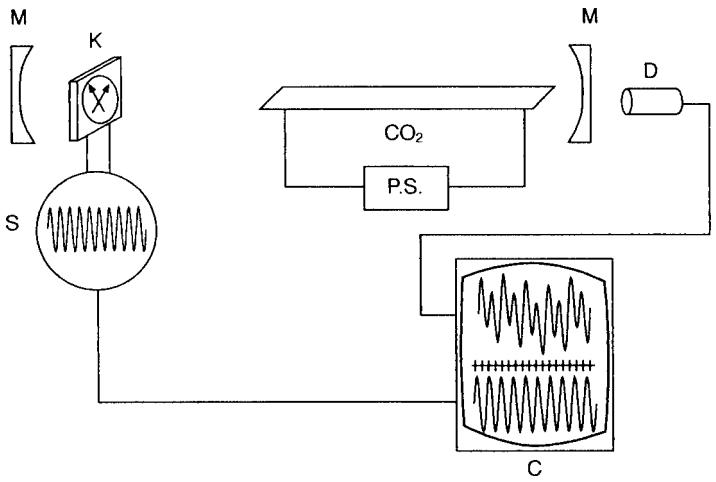

Figure 1.1 This schematic representation of a laser with modulated losses shows the carbon dioxide tube (CO_2); power source (P.S.); mirrors (M); Kerr cell (K); signal generator (S); detector (D); and computer, oscilloscope, and recorder (C). A variable electric field across the Kerr cell rotates its polarization direction and modulates the electric field within the cavity.

The topological methods described in the remainder of this work were originally developed to understand the data generated by a laser with modulated losses [26]. A schematic of this laser is shown in Figure 1.1. A CO_2 gas tube is placed between two infrared mirrors (M). The ends of the tube are terminated by Brewster angle windows, which polarize the electric field in the vertical direction. Under normal operating conditions, the laser is very stable. A Kerr cell (K) is placed inside the laser cavity. A Kerr cell modifies the polarization state of an electromagnetic field. This modification, coupled with the polarization introduced by the Brewster windows, allows one to change the intracavity losses. The Kerr cell is modulated at a frequency determined by the operating conditions of the laser. When the modulation is small, the losses within the cavity are small, and the laser output tracks the input modulation from the signal generator. The input signal (from the signal generator) and the output signal (the measured laser intensity) are both recorded in a computer (C). When the modulation crosses a threshold, the laser output can no longer track the signal input. At first every other output peak has the same height, then every fourth peak, then every eighth peak, and so on.

In Figure 1.2 we present some of the recorded and processed signals from this part of the period-doubling cascade and beyond [26]. The signals were recorded under different operating conditions and are displayed in five lines, as follows: (a) period 1; (b) period 2; (c) period 4; (d) period 8; (e) chaos. Each of the four columns presents a different representation of the data. In the first column the intensity output is displayed as a function of time. In this presentation the period-1 and period-2 behaviors are clear but the higher-period behavior is not.

The second column displays a projection of the dynamics into a two-dimensional plane, the dI/dt vs. $I(t)$ plane. In this projection, periodic orbits appear as closed loops (deformed circles) that go around once, twice, four times, etc. before closing. In this presentation the behavior of periods 1, 2, and 4 is clear. Period 8 and chaotic behavior is less clear. The third column displays the power spectrum. Not only is the periodic behavior clear from this display, but the relative intensity of the various harmonics is also evident. Chaotic behavior is manifest in the broadband power spectrum. Finally, the last column displays a stroboscopic sampling of the output. In this sampling technique, the output intensity is recorded each time the input signal reaches a maximum (or some fixed phase with respect to the maximum). There is one sample per cycle. In period-1 behavior, all samples have the same value. In period-2 behavior, every other sample has the same value. The stroboscopic display clearly distinguishes between periods 1, 2, 4, and 8. It also distinguishes periodic from chaotic behavior. All four of these display modalities are available in real time, during the experiment.

The laser with modulated losses has been studied extensively both experimentally [23–29] and theoretically [30–32]. The rate equations governing the laser intensity I and the population inversion N are

$$\begin{aligned} \frac{dI}{dt} &= -k_0 I \left[(1 - N) + m \cos(\omega t) \right] \\ \frac{dN}{dt} &= -\gamma \left[(N - N_0) + (N_0 - 1) I N \right] \end{aligned} \tag{1.2}$$

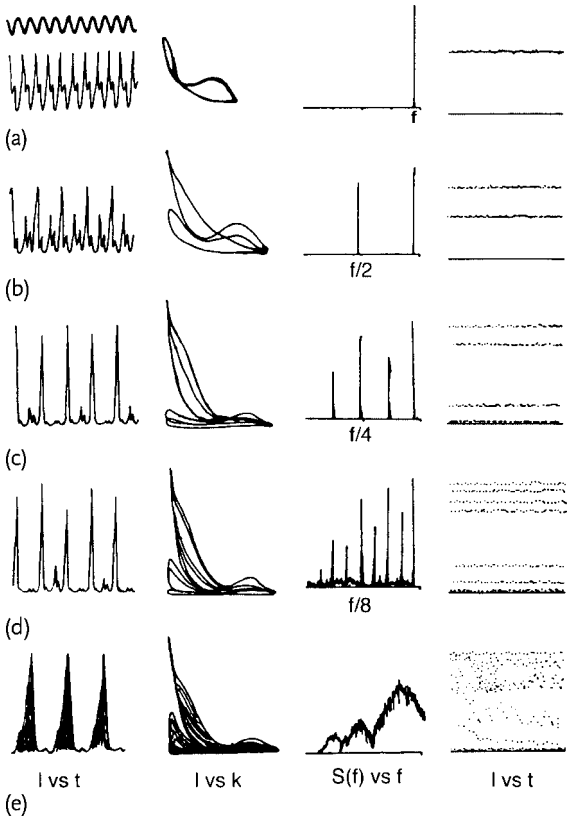

Figure 1.2 Each column provides a different representation of the experimental data. Each row describes different experimental conditions. The first column shows the recorded intensity time signal, $I(t)$. The second column presents the phase-space projection, $dI(t)/dt$ vs. $I(t)$. The third column shows the power spectrum of the recorded intensity signal. The frequencies of the Fourier components in the signal, and their relative amplitudes, jump out of this plot. The last column presents a stroboscopic plot (Poincaré section). This is a record of the intensity output at each successive peak (or, more generally, at some constant phase) of the input signal. The data sets were recorded under the following experimental conditions: (a) period 1; (b) period 2; (c) period 4; (d) period 8; (e) chaotic. Reprinted with permission from Tredicce *et al.* [26].

Here m and ω are the modulation amplitude and angular frequency, respectively, of the signal to the Kerr cell; N_0 is the pump parameter, normalized to $N_0 = 1$ at the threshold for laser activity, and k_0 and γ are loss rates. In dimensionless, scaled form this equation is

$$\begin{aligned} \frac{du}{d\tau} &= \left[z - A\cos(\Omega\tau)\right] u \\ \frac{dz}{d\tau} &= (1 - \epsilon_1 z) - (1 + \epsilon_2 z)\, u \end{aligned} \tag{1.3}$$

The scaled variables are $u = I$, $z = k_0\kappa(N-1)$, $t = \kappa\tau$, $A = k_0 m$, $\epsilon_1 = 1/\kappa k_0$, and $\kappa^2 = 1/\gamma k_0(N_0 - 1)$. The bifurcation behavior exhibited by the simple models (1.2) and (1.3) is qualitatively, if not quantitatively, in agreement with the experimentally observed behavior of this laser.

The pair of equations (1.3) comprise a dynamical system. This dynamical system is explicitly time dependent through the term $A\cos(\Omega\tau)u$. Its time evolution can be represented geometrically in a state space, called the phase space, whose coordinates are the state variables of the system. The state variables for this dynamical system are u and z living in the half plane $R^2_+ = (u \geq 0, z)$ and the time coordinate τ. The phase space is the three-dimensional torus $R^2_+ \times S^1$, where times t and $t + T$ are identified since $\cos[\Omega(\tau + T)] = \cos(\Omega\tau)$ when the period of the drive $T = 2\pi/\Omega$. As time flows, the representative point $(u, z, \tau \mod T)$ follows a trajectory in phase space determined by (1.3). These equations define a *flow* in phase space. A flow maps the state at time t to the state at a later time. This dynamical system has four control parameters: $(A, \Omega, \epsilon_1, \epsilon_2)$.

The stroboscopic sampling technique is equivalent to the construction of a *Poincaré section* for this periodically driven dynamical system. A Poincaré section is a hypersurface of phase space that is intersected transversely by trajectories and serves to simplify the analysis of the dynamics. Instead of entire trajectories, we can consider only their intersections with the Poincaré sections. If the dynamics is periodic, there will be a finite number of intersections, and their number may be used to define the period. If the dynamics is irregular, there is an infinite number of intersections. In both cases, the dynamics is described by the *Poincaré map* (or first-return map), which maps the location of an intersection to the location of the next intersection when following a trajectory.

A bifurcation diagram for the laser model (1.3) is shown in Figure 1.3. The bifurcation diagram is constructed by varying the modulation amplitude A and keeping the other three control parameters fixed. It indicates the nature of the behavior observed for different values of A. Periodic solutions are represented by a characteristic quantity (e.g., their amplitude) plotted as a function of A, with a solid line if they are stable and with a dotted line otherwise. Chaotic solutions are represented by plotting all values obtained by the stroboscopic sampling for the given value of A (in general, this covers one or several intervals). It may happen that the system is multistable and that we have several possible stable solutions for the same value of A. The overall structures of the bifurcation diagram in Figure 1.3 are similar to experimentally observed bifurcation diagrams.

Figure 1.4 shows a bifurcation diagram observed experimentally in another laser system [33, 34]. By means of a simple electronic device, stroboscopic samples are plotted as a function of modulation amplitude on the screen of an oscilloscope so that a period-n orbit is characterized by n different samples. The structure is simpler than in Figure 1.3 because only one solution may be shown for a given modulation amplitude but the sequence of bifurcations observed is compatible with that of Figure 1.3.

This figure shows that a period-1 solution exists above the laser threshold ($N_0 > 1$) for $A = 0$ and remains stable as A is increased until $A \sim 0.8$. It

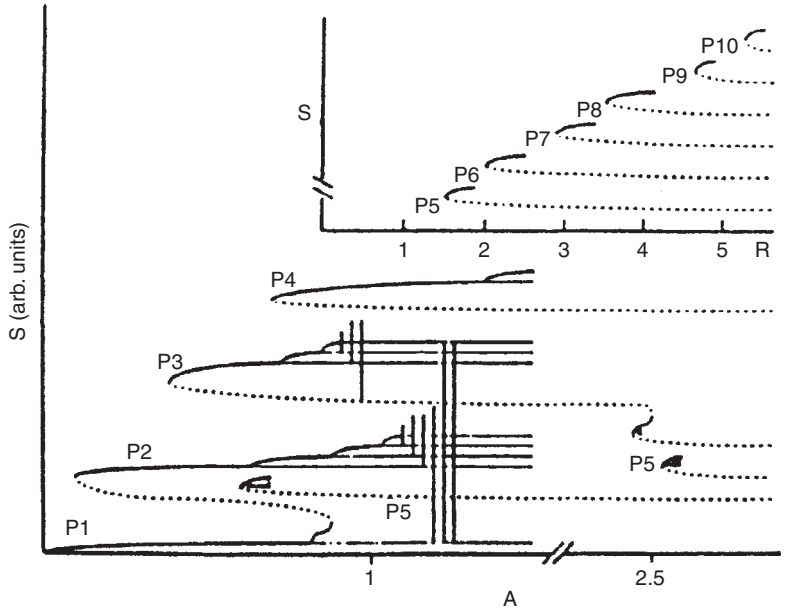

Figure 1.3 The bifurcation diagram for the laser model (1.3) is computed by varying the modulation amplitude A. Stable periodic orbits (solid lines), regular saddles (dotted lines), and strange attractors are shown. Period-n branches ($P_n, n \geq 2$) are created in saddle–node bifurcations and evolve through the Feigenbaum period-doubling cascade as the modulation amplitude increases. There are two apparently distinct stable period-2 orbits. However, these are connected by an unstable period-2 orbit (dotted, extending from $A \simeq 0.1$ to $A \simeq 0.8$) and thus constitute a single period-2 orbit. A period-3 orbit also winds its way through the bifurcation diagram. Two distinct stable period-4 orbits are present and coexist over a short range of parameter values ($0.7 < A < 0.8$). The inset shows a sequence of period-n orbits (*Newhouse orbits*) for $n \geq 5$. The Smale horseshoe mechanism predicts that as many as three nonequivalent pairs of period-5 orbits could exist. The locations of the two additional pairs are shown in this diagram at $A \simeq 0.65$ and $A \simeq 2.5$. Parameter values: $\epsilon_1 = 0.03, \epsilon_2 = 0.009$, $\Omega = 1.5$.

becomes unstable above $A \sim 0.8$, with a stable period-2 orbit emerging from it in a period-doubling bifurcation. In a period-doubling bifurcation, a periodic orbit becomes unstable while giving birth to another periodic orbit of twice the period. At the bifurcation, the newborn orbit is superimposed on the original orbit.

Contrary to what might be expected, this is not the early stage of a period-doubling cascade, for the period-2 orbit is annihilated at $A \sim 0.85$ in an inverse saddle–node bifurcation (cf. Section 2.7.5.1) with a period-2 regular saddle. In a saddle–node (resp. inverse saddle–node) bifurcation, two periodic orbits appear (resp. disappear) simultaneously and are degenerate at the bifurcation. The saddle–node bifurcation at $A \sim 0.85$ destroys the basin of attraction of the period-2 orbit. The basin of attraction of a solution is the set of points whose orbits converge to the solution. Any point in that basin is dumped into the basin of a period

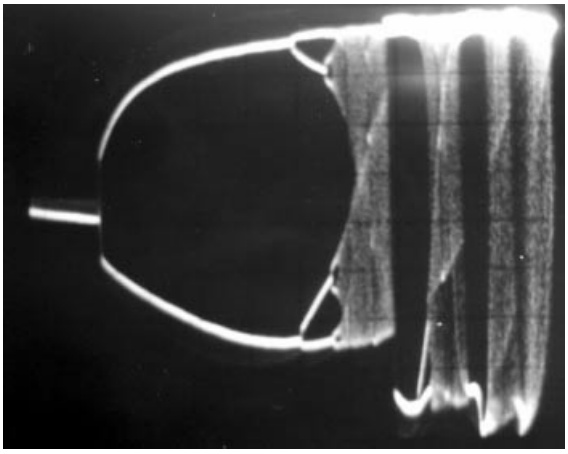

Figure 1.4 Bifurcation diagram of the chaotic laser used in [33, 34]. The laser output intensity is sampled stroboscopically at the modulation period. The modulation amplitude A is swept slowly and successive samples are plotted on the oscilloscope screen as a function of modulation amplitude. Thus, a period-n periodic window is characterized by n different samples whereas continuous ranges are explored in the chaotic windows.

$4 = 2 \times 2^1$ orbit, even though there are two other coexisting basins of attraction for stable orbits of periods $6 = 3 \times 2^1$ and four at this value of A.

Subharmonics of period n (P_n, $n \geq 2$) are created in saddle–node bifurcations at increasing values of A and I (P_2 at $A \sim 0.1$, P_3 at $A \sim 0.3$, P_4 at $A \sim 0.7$, P_5 and higher shown in the inset). All subharmonics in this series up to period $n = 11$ have been seen both experimentally and in simulations of (1.3). The evolution (*perestroika* [35]) of each of these subharmonics follows a standard scenario as T increases [36] (cf. Chapter 9):

1. A saddle–node bifurcation creates an unstable saddle and a node that is initially stable.
2. Each node becomes unstable and initiates a period-doubling cascade as A increases. The cascade follows the standard Feigenbaum scenario [22, 37–39]. The ratio of A intervals between successive bifurcations, and of geometric sizes of the stable nodes of periods $n \times 2^k$, has been estimated up to $k \leq 6$ for some of these subharmonics, both from experimental data and from the simulations. These ratios are compatible with the universal scaling ratios.
3. Beyond accumulation, there is a series of noisy orbits of period $n \times 2^k$ that undergo inverse period-halving bifurcations. This scenario was predicted by Lorenz [40].

Additional systematic behavior has been observed. Higher subharmonics are generally created at larger values of A. They are created with smaller basins of attraction. The range of A values over which the Feigenbaum scenario is played out becomes smaller as the period n increases. In addition, the subharmonics show an

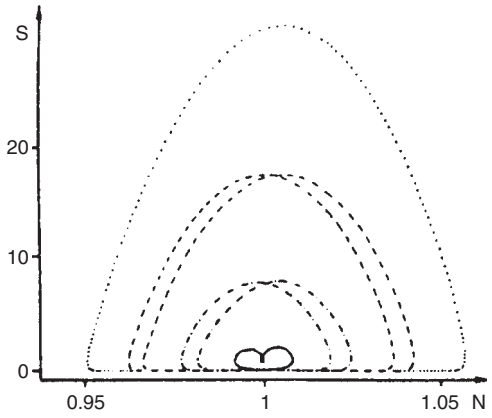

Figure 1.5 Multiple basins of attraction can coexist over a broad range of parameter values. The stable periodic orbits and the strange attractors within these basins have a characteristic organization. The coexisting orbits shown above are, from inside to outside: period-2 bifurcated from a period-1 branch; period 2; period 3; period 4. The two inner orbits are separated by an unstable period-2 orbit (not shown).

ordered pattern in phase space. In Figure 1.5 we show four stable periodic orbits that coexist under certain operating conditions. Roughly speaking, the larger period orbits exist "outside" the smaller period orbits. These orbits share many other systematics, which have been described by Eschenazi, Solari, and Gilmore [36]. In Figure 1.6 we show an example of a chaotic time series taken for $A \sim 1.3$. The chaotic attractor based on the period-2 (period-1) orbit has just collided with the period-3 regular saddle.

This experiment, and many others like it, was designed to test the Feigenbaum scaling predictions made for period-doubling cascades. These tests focus on a very small part of the bifurcation diagram. Much of the data recorded did not have a direct bearing on the scaling question. They did have a direct bearing on understanding the behavior of chaotic dynamical systems. Based on the theory that it is criminal to discard hard-won experimental data, we sought to analyze these "surplus" data.

Bifurcation diagrams had been observed for a variety of physical systems at that time: other lasers [41, 42], electric circuits [43–46], a biological model [47], and a bouncing ball [14]. They all supported the Feigenbaum scaling relations at the accumulation point of the period-doubling bifurcations. Their bifurcation diagrams away from the accumulation points are similar but not identical to those shown above. This raised the question of whether similar processes were governing the description of this large variety of physical systems.

During these analyses, it became clear that the standard tools for analyzing chaotic data – estimates of the spectrum of Lyapunov exponents and estimates of the various fractal dimensions – were not sufficient for a satisfying understanding of the stretching and squeezing processes that occur in phase space and that are responsible for generating chaotic behavior. In the laser we found many coexisting basins

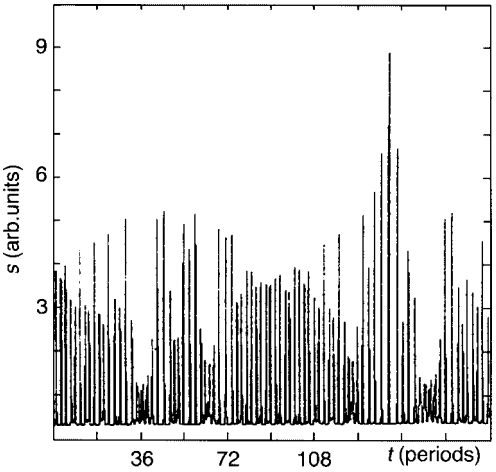

Figure 1.6 This time series from a laser with modulated losses was taken at a value of $A \sim 1.3$, which is just beyond the collision (crisis) of the strange attractors based on the period-2 and period-3 orbits. There is an alternation in this time series between noisy period-2 and noisy period-3 behavior. Essentially the same experiment was repeated later with a more forgiving (logarithmic) detector. A segment of the time series from this experiment is shown in Figure 7.23. The minima are not saturated using the logarithmic detector.

of attraction, some containing a periodic attractor, others containing a strange attractor. The rapid alternation between periodic and chaotic behavior as control parameters (e.g., A and Ω) were changed meant that Lyapunov exponents and fractal dimensions depended on the basins and varied at least as rapidly.

For this reason we sought to develop additional tools for the analysis of data generated by dynamical systems that exhibit chaotic behavior. The objective was to develop measures that were invariant under control parameter changes.

1.3 Objectives of a New Analysis Procedure

In view of the experiments just described and the data that they generated, we hoped to develop a procedure for analyzing data that achieved a number of objectives. These included an ability to answer the following questions:

1. Is it possible to develop a procedure for understanding dynamical systems *and their evolution* (perestroikas) as the control parameters (e.g., k_0, m, γ, or $A, \Omega, \epsilon_1, \epsilon_2$) change?
2. Is it possible to identify a dynamical system by means of topological invariants, following suggestions proposed by Poincaré?
3. Can selection rules be constructed under which it is possible to determine the order in which periodic orbits can be created or annihilated by standard bifurcations?

4. Is it possible to determine when two strange attractors are (a) equivalent in the sense that one can be transformed into the other without creating or annihilating orbits, (b) adiabatically equivalent (one can be deformed into the other by changing parameters to create or annihilate only a small number of orbit pairs below any period), or (c) inequivalent (there is no way to transform one into the other)?

1.4 Preview of Results

A new topological analysis procedure was developed in response to the questions asked of the data initially. These questions are summarized in Section 1.3. The remarkable result is that there is now a positive and constructive answer to the question: How can I look at experimental data, such as those shown in Figure 1.2 or 1.6, and extract useful information, let alone information about stretching and squeezing, let alone a small set of integers?

This new analysis procedure answered more questions than were asked originally. It also raised a great many additional questions. This is one of the ways we know that we are on the right track.

The results of this new topological analysis procedure are presented throughout this book. Below we provide a succinct preview of the major accomplishments of this topological analysis tool.

- There is now a simple, algorithmic procedure for:
  - Classifying strange attractors (Chapter 5).
  - Extracting classification information from experimental data (Chapters 6–8).
- The data requirements are not heavy. Data sets of limited length are required.
- The data need not be exceptionally clean. Only a modest signal-to-noise level is required. The analysis method degrades gracefully with noise. Specifically, as the noise level degrades the data, it becomes more difficult to identify the higher-period orbits, which are the least important for this analysis. The most important orbits, those of lowest period, persist longest with increasing noise. As a result, "Murphy is on vacation" (author of the famous law).
- The data analysis method comes endowed with a rejection criterion.
- The branched manifold identifies the stretching and squeezing mechanisms that generate chaotic behavior.
- Thus, this doubly discrete classification describes "how to model the dynamics."
- There is now a classification theory for low-dimensional strange attractors.
  - It is topological.
  - It has a hierarchy of four levels.
  - Each is discrete.
  - There is rigidity and there are degrees of freedom.
  - It is applicable to R^3 – for now.

- The four levels of structure are:
  - Basis sets of orbits (Chapter 9)
  - Branched manifolds (Chapter 5)
  - Bounding tori (Chapter 12)
  - Embeddings of bounding tori (Chapter 13).
- There is a poetry to the organization of this hierarchy:

Links of periodic orbits
organize
Bounding tori
organize
Branched manifolds
organize
Links of periodic orbits

- We have come across many unexpected results:
  - Routes to chaos are paths through orbit forcing diagrams (Chapter 9, Figure 9.8).
  - Perestroikas of orbits are constrained by branched manifolds (Chapter 9).
  - Perestroikas of branched manifolds are constrained by bounding tori (Chapter 11).
  - Perestroikas of bounding tori are also constrained (Chapter 11).
  - The global Poincaré surface of section is the union of certain disks in a trinion decomposition of a bounding torus (Chapter 11).
  - There are systematic ways to enlarge/reduce the size of a strange attractor using symmetry (Chapter 10).
  - The number of inequivalent ways of embedding a low-dimensional strange attractor in R^3 has been determined (Chapter 12).
  - There is a representation theory for strange attractors (Chapter 12).
- Periodic orbits play a fundamental role in the analysis of physical data.
  - *Linear systems:* Fourier showed in 1822 that periodic (circular) orbits of different frequencies (in fact, their projections onto the horizontal and vertical axes: sines and cosines) could be used to analyze the signals from any linear system. This became an industrial-strength tool with the development of the fast Fourier transform algorithm by Cooley and Tukey in 1965.
  - *Nonlinear systems:* Unstable periodic orbits of different periods can be used to analyze the signals from nonlinear systems. This analysis procedure will become an industrial-strength tool when this ability is extended from three- to higher-dimensional nonlinear dynamical systems.
- We have asked many questions.
- Each question that we have answered has raised two new questions.
- We must be on the right track!

1.5 Organization of This Work

The best way to summarize the organization of this book is to provide a brief summary of each chapter.

Chapter 1 – Introduction. Data generated by a laser with modulated losses are used to illustrate the complexity of the problem. The problem – in a nutshell – is this: How do you let data speak to you and tell you what mechanisms are generating chaotic behavior? The basic objectives of a new analysis method are outlined. Also summarized are the accomplishments of the new topological analysis method.

Chapter 2 – Discrete Dynamical Systems: Maps. Mappings have been excellent surrogates for continuous dynamical systems. They were introduced by Poincaré as a way to reduce the dimension of the space needed to study the properties of chaotic dynamical systems. We introduce the idea of a map (of a Poincaré section to itself) as a dynamical system in this chapter. The chapter contains many familiar results useful for the topological analyses that follow. These include review of the logistic and circle maps, the interplay between periodic orbits and chaotic behavior, the implications of topology on the order in which periodic orbits are created as control parameters vary, symbolic dynamics, kneading theory, and several ways to compute topological entropy. We emphasize that noninvertible maps are characterized by their singularity structure. Invertible maps are not, but if they are dissipative, their large n iterates are also characterized by their singularities.

Chapter 3 – Continuous Dynamical Systems: Flows. The principal subject of this book is the analysis of dynamical systems. These are systems of coupled (nonlinear) first-order ordinary differential equations. We review the properties of such systems in this chapter, such as the laser equations (1.2) and (1.3). In particular, we introduce the existence and uniqueness theorem, on which all else rests. All the other important concepts related to flows are also introduced here. These include especially the spectrum of Lyapunov exponents and the concepts of stretching and squeezing in phase space. Four of the basic theoretical testbeds used in the study of dynamical systems theory are introduced here. These are the nonautonomous Duffing and van der Pol equations for periodically driven two-dimensional nonlinear oscillators and the autonomous Lorenz and Rössler equations.

Chapter 4 – Topological Invariants. If we want to understand chaos in dynamical systems, we must understand strange attractors. In particular, we must understand the (unstable) periodic orbits "in" strange attractors. How do we know this? Poincaré [48] told us so over a century ago in an oft-quoted and more often neglected statement:

"[periodic orbits] yield us the solutions so precious, that is to say, they are the only breach through which we can penetrate into a place which up to now has been reputed to be inaccessible."

Our contribution to this understanding is the introduction of Gauss linking numbers between pairs of orbits into the study of chaotic dynamical systems. These, and an even more refined invariant for orbits and pairs of orbits, the relative rotation rates, are introduced in this chapter. These topological invariants exist only in R^3 or, more generally, three-dimensional manifolds. This presently limits the results of the topological analysis method to strange attractors that live in three-dimensional manifolds.

Chapter 5 – Branched Manifolds. Different stretching and squeezing mechanisms generate different (inequivalent) strange attractors. The topological organization of the unstable periodic orbits in a strange attractor provides a clear fingerprint for the strange attractor. In fact, the organization identifies the stretching and squeezing mechanisms that generate the strange attractor while organizing all the unstable periodic orbits in the strange attractor in a unique way. There is a geometric structure that supports all the unstable periodic orbits in a strange attractor with the same unique organization that they possess in the strange attractor. This structure is variously called a knot holder, a branched manifold, or a template. These structures are introduced in this chapter. They can be identified by a set of integers. Therefore, they are discretely classifiable. As a result, strange attractors are discretely classifiable.

Chapter 6 – Topological Analysis Program. There is a straightforward procedure for extracting the signature of a strange attractor from experimental data. This consists of a number of simple and easily implementable steps. The input to this analysis procedure consists of experimental time series. The output consists of a branched manifold or, more abstractly, a set of integers. The results of this analysis are subject to a follow-on rejection or "confirmation" test. In Chapter 6 we present a step-by-step account of the topological analysis method.

Chapter 7 – Folding Mechanisms: A_2. The topological analysis procedure has been applied to a large number of data sets. Most of them revealed stretching and squeezing mechanisms that were variations on a single theme. The basic theme in its most elementary form is the Smale horseshoe stretch-and-fold mechanism (without global torsion). Other variations include reverse horseshoes, horseshoes with global torsion, and stretch-and-roll mechanisms, variously called gâteau roulé or jellyroll mechanisms. The "S" mechanism is also predicted to be possible and was recently been discovered in laser experiments carried out in Zaragoza. These mechanisms form a large subset of the possible mechanisms that can occur when the phase space containing the attractor is $R^2 \times S^1$. The topological structure of the phase space is evident by inspection of the attractor: If it has a hole in the mid-

dle, it has this property and the dynamics can be formed only by stretch-and-fold mechanisms, although the folding can sometimes take on imaginative forms.

Chapter 8 – Tearing Mechanisms: A_3. Dynamical systems that possess twofold symmetry sometimes exhibit a different stretching and squeezing mechanism. This mechanism involves tearing the flow apart in phase space. Two physical systems whose strange attractors are generated by this mechanism are analyzed in this chapter. The Lorenz system is a standard representation for this process. Folding and tearing are intimately related to the two simplest catastrophes [49], the fold A_2 and the cusp A_3. Other mechanisms are related to other catastrophes and singularities.

Chapter 9 – Unfoldings. One of our objectives is to be able to predict what will – or can – happen as control parameters are changed. This question reduces to the problem of determining the unfoldings of the "germs" of dynamical systems. If control parameters are changed slightly, there is a perturbation in the spectrum of orbits in a strange attractor, but the underlying branched manifold remains unchanged (is robust). The change in the spectrum of orbits can occur in a number of ways limited by topological considerations. If the change in control parameters is very large, the branched manifold itself can change. In the latter case, new branches can be added or old branches removed in a number of ways limited by topological and continuity considerations. As a result, the perestroikas in strange attractors can take place in a number of ways that are large but constrained. Put another way: Given any point in this multiply discrete classification representing a strange attractor with a specific spectrum of periodic orbits, under perturbation of the control parameters it can only move to its neighboring points in this classification. Its neighbors are large in number but constrained by both continuity and topological arguments.

Chapter 10 – Symmetry. The measurements described in the first section are of the laser output intensity. We understand that the electric field amplitude $E(t)$ is more fundamental than the intensity since $I(t) = E(t)^2$. We also feel (correctly) that if the intensity behaves chaotically, so must the amplitude. But is the chaos the same? In other words, is the strange attractor generated by the amplitude $E(t)$ (were we able to measure it) equivalent to the strange attractor generated by the intensity $I(t)$? This chapter is devoted to an analysis of dynamical systems that are locally equivalent but not globally so (E is a two-branch cover of I: $E = \text{“}\pm\text{”}\sqrt{I}$). We present unexpected, elegant, and powerful results relating cover and image dynamical systems, even when the symmetry group relating them has only two group elements. These results allow us, for example, to take the square root of the Duffing and van der Pol dynamical systems, simplifying the analyses of these two systems dramatically.

Chapter 11 – Bounding Tori. A strange attractor that is contained in R^3 can be immersed in a three-dimensional manifold in R^3. This manifold is a "handlebody"

with g holes, and its boundary is a torus of genus g. The flow that generates the strange attractor contained in the genus-g torus "dresses" the two-dimensional surface with a flow field. Together, the surface and the flow field on it allow us to refine our classification of strange attractors. The genus-g tori that serve to classify strange attractors are built up from fundamental building blocks just as branched manifolds are built up from fundamental building blocks. These were called Y junctions in the era of microwave technology; today, with the ascendance of field theory, they are called trinions. A genus-g torus is built up from $(g-1)$ pairs of trinious – each pair consists of a splitting and a joining trinion. The Poincaré surface of section for a strange attractor contained within a genus-g torus is the union of $g-1$ disks: each disk occurs at the connection between the output port of a joining trinion with the input port of a splitting trinion. Bounding tori are classified in group-theoretical terms. The transition matrix that describes the connectivity of the genus-g torus is the sum of two permutation matrices for the permutation group P_{g-1}.

Chapter 12 – Representation Theory for Strange Attractors. A strange attractor can be embedded in a phase space in many different ways. Some embeddings are equivalent to others. Equivalence is by isotopy. Other embeddings may be inequivalent. Inequivalent embeddings of dynamical systems are in some sense like inequivalent representations of groups and algebras. A parallel suite of questions poses itself for embeddings of strange attractors: Can we find representation labels for inequivalent embeddings? Do we have a complete set? If we immerse two inequivalent embeddings into higher dimensions, do they become equivalent because there is more room in the higher-dimensional phase space in which to move around and avoid obstructions to isotopy? It is possible to answer all these questions for "low-dimensional" strange attractors: those that exist in R^3. The result depends on the genus of the attractor in R^3. In R^5 all representations of a strange attractor become equivalent. This means there is a single universal representation of any given three-dimensional strange attractor in R^5, and that any information gleaned by studying a five-dimensional embedding of a three-dimensional strange attractor depends entirely on the original attractor and is totally independent of the embedding.

Chapter 13 – Flows in Higher Dimensions. What happens when a strange attractor cannot be shoehorned into a three-dimensional manifold? This question is asked for four-dimensional strange attractors. Although periodic orbits can no longer be the working tool in this dimension, it appears that the singularity structure of a return map of a certain manifold into itself provides information needed to classify higher-dimensional attractors. The classification can be discrete when the number of unstable directions in the flow is large and the number of stable directions is small. We survey the rich spectrum of possibilities that exist for four-dimensional chaotic flows with two unstable directions and one stable direction. The results in this chapter are few – rather, the chapter lays the groundwork for future developments in this field.

Chapter 14 – Program for Dynamical Systems Theory. When this work was initiated, one of the authors was surprised that a classification theory for strange attractors did not exist. He was very frustrated by the fact that nobody was even asking the question: How do you classify strange attractors? The first thing you do when confronted with new physics is to try to understand and to classify. The situation has not improved with time. Many are the works devoted to observation or the analysis of small problems. The response of many in the field of dynamical systems to the lack of rapid understanding and development of the field in low dimensions was to look for and to describe more complicated problems in higher dimensions! If we don't have a theory in low dimensions, how are we to find one in higher dimensions? Chapter 14 addresses this problem in an indirect way. There are many similarities between the older fields of Lie group theory and singularity theory (catastrophe theory) and the newer field of dynamical systems theory. It is a hope that simply asking the same questions of the newer field that were asked of the older fields will at least provide a structure in which advances can be made more rapidly. The entire chapter is devoted to comparing similar questions in these three fields and suggesting how some answers might hasten the development of a theory for dynamical systems that is as mature as Lie group theory and singularity theory are today.

Appendix A – Determining Templates from Topological Invariants. A characterization method is most useful when it can be routinely and effortlessly applied to many different systems. In Appendix A, we describe computational techniques useful for determining the simplest template compatible with a given set of topological invariants. These techniques have been implemented in computer programs. We also discuss how the information about the symbolic dynamics of the periodic orbits extracted from invariants can be used to construct symbolic encodings for chaotic attractors.

Appendix B – Embeddings. The first essential step in the analysis of experimental data is the creation of a representation of the attracting set in a suitable phase space. A representation of the attractor that is 1 : 1 without self-intersections is called an embedding. Creating an embedding is an art form. More precisely, there are many ways to create mappings of an unseen attractor into a phase space. Determining whether a mapping is an embedding is the hard part. In the past, tests of whether a mapping was an embedding were based on geometric measures, such as fractal dimensions, or dynamical measures, such as the spectrum of Lyapunov exponents. In this chapter we review many of the mapping techniques and then test the embedding tests. We find that tests based on geometric and dynamical measures give incorrect predictions more often than correct predictions of whether a mapping is an embedding on two test dynamical systems. A topological test for embeddings, though limited to mappings into three-dimensional phase spaces, is far more reliable. We summarize the findings with the following warning: "All previous claims that a mapping is an embedding based solely on geometric or dynamical tests should be taken with a grain of salt unless other independent methods have been used to confirm that the mapping lacks self-intersections."