

Numerical Analysis of Dynamical Systems

John Guckenheimer *

October 5, 1999

1 Introduction

This paper presents a brief overview of algorithms that aid in the analysis of dynamical systems and their bifurcations. The viewpoint is geometric and the goal is to describe algorithms that reliably compute objects of dynamical significance. Reliability has three facets:

1. the probability that the algorithm returns an answer for different choices of starting data,
2. whether the computed object is qualitatively correct, and
3. the accuracy with which the objects are computed.

Numerical analysis has traditionally concentrated on the third of these topics, but the first two are perhaps more important in numerical studies that seek to delineate the structure of dynamical systems.

This survey concentrates on exposition of fundamental mathematical principles and their application to the numerical analysis of examples. There is a strong interplay between dynamical systems theory and computational analysis of dynamical systems. The theory provides a framework for interpreting numerical observations and foundations for algorithms. Apparent discrepancies between computational output and theoretical expectations point to areas where phenomena have been overlooked in the theory, areas where algorithms produce misleading results, and areas where the relationship between theory and computation is more subtle than anticipated. Several examples of simple systems are used in this article to illustrate seeming differences between computation and theory.

*Mathematics Department, Cornell University, Ithaca, NY 14853. This work was partially supported by grants from the Department of Energy, Air Force Office of Scientific Research and the National Science Foundation

Geometric perspectives have been introduced relatively recently to the numerical analysis of ordinary differential equations. The tension between geometric and more traditional analysis of numerical integration algorithms can be caricatured as the interchange between two limits. The object of study are systems of ordinary differential equations and their flows. Numerical solution of initial value problems for system of ordinary differential equations discretize the equations in time and produce sequences of points that approximate solutions over time intervals. Dynamical systems theory concentrates on questions about long time behavior of the solution trajectories, often investigating intricate geometry in structures formed by the trajectories. The two limits of (1) discretizing the equations with finer and finer resolution in time and (2) letting time tend to infinity do not commute. Classical theories of numerical analysis give little information about the limit behavior of numerical trajectories with increasing time. Extending these theories to do so is feasible only by making the analysis specific to classes of systems with restricted geometric properties. The blend of geometry and numerical analysis that is taking place in current research has begun to produce a subject with lots of detail and richness. Interesting examples from diverse applications infuse the subject and establish mathematical connections between other disciplines. Thus, the development of better algorithms and software can have far reaching consequences. This paper takes a pragmatic view of this research. The focus here is on understanding the mathematical properties observed in numerical computation and on assessing the capability of theory, algorithms and software to elucidate the structure of dynamical models in mathematics, science and engineering. Issues that have been investigated from this perspective are presented and a few pointers are provided to the rapidly growing literature.

2 Numerical Integration

2.1 Classical Theory

Systems of ordinary differential equations

$$\dot{x} = f(x); \quad f : R^n \rightarrow R^n \quad (1)$$

define vector fields. Vector fields on manifolds are also defined by systems of the form (1) in local coordinates [148]. The existence and uniqueness theorem for ODEs [86] states that a Lipschitz continuous vector field (1) has a unique flow $\Phi : R^n \times R \rightarrow R^n$ defined in a neighborhood of $R^n \times 0$ with the properties that $\Phi(x, 0) = x$ and $\dot{\Phi}(x, t) = f(\Phi(x, t))$. The *time t map* $\phi_t : R^n \rightarrow R^n$ is defined by $\phi_t(x) = \Phi(x, t)$. The curves $\Phi(x, t)$ defined by fixing x and letting t vary are *trajectories*, denoted by x_t . There are seldom explicit formulas for Φ in terms of f . Iterative *numerical integration* algorithms are used to compute trajectories with discrete time approximations that march along the trajectories. Numerical integration is a mature subject, but still very active - especially with regard to algorithms designed for special classes of equations. The subject

is supported by extensive theory and abundant software. Several excellent texts and references are Henrici [91], Gear [66], Hairer, Norsett and Wanner [86], Hairer and Wanner [87] and Ascher and Petzold [10]. Part of the intricacy of the subject lies in the fact that no single integration algorithm is suitable for all problems. Different algorithms reflect trade-offs in ease of use, accuracy and complexity. The basic concepts of numerical integration are explained here only briefly. Explicit Runge-Kutta algorithms are described, followed by a short survey of refinements, alternate approaches and terminology for numerical integration of ODEs.

Explicit Runge Kutta methods construct mappings ψ_h from the vector field (1) that depend upon a parameter h , called the *time step* of the method. Several partial steps are taken from an initial point, and the values are combined to produce the map ψ_h . The primary goal is to produce a family of mappings ψ_h depending on h , whose Taylor series expansion in the time step h agrees with that of the flow map ϕ_h to a specified degree. Each function evaluation is called a *stage* of the method. The method is said to have *order* d if the Taylor series agree to degree $d+1$. Each stage is performed at a point that depends upon the preceding stage. The scheme for an s -stage method has the following form:

$$\begin{aligned} k_1 &= f(x) \\ k_2 &= f(x + ha_{21}k_1) \\ &\vdots \\ k_s &= f(x + h(a_{s1}k_1 + a_{s2}k_2 + \cdots + a_{s,s-1}k_{s-1})) \\ \psi_h(x) &= x + h(b_1k_1 + b_2k_2 + \cdots + b_s k_s) \end{aligned} \tag{2}$$

Formulation of higher order Runge-Kutta methods for system (1) is based upon repeated differentiation of this equation. With each differentiation, substitution of $f(x)$ for $\dot{x}$ on the right hand side yields expressions for the derivatives of trajectories in terms of derivatives of f . For example, the second and third derivatives of $x(t)$ are given by

$$\begin{aligned} \ddot{x} &= D_x f \dot{x} = D_x f f \\ x^{(iii)} &= D_{xx} f f^2 + (D_x f)^2 f \end{aligned}$$

Taylor expansion of ψ_h in the system of equations (2) gives expressions in the derivatives of f and the coefficients a_{ij} and b_j of the method. Equating the degree $d+1$ expansions of ψ_h and $x(h)$ obtained from repeated differentiation of system (1) produces a system of polynomial equations for a_{ij} and b_j . The number of equations obtained in this manner grows faster than d . For $d \leq 4$, there are order d methods with d stages [86]. Order d methods with $d > 4$ require more than d stages. As d increases, the complexity of solving the equations for order d methods grows rapidly. The fourth order method

$$\begin{aligned} k_1 &= hf(x_0) \\ k_2 &= hf(x_0 + k_1/2) \\ k_3 &= hf(x_0 + k_2/2) \end{aligned}$$

$$\begin{aligned} k_4 &= hf(x_0 + k_3) \\ x_1 &= x_0 + \frac{1}{6}(k_1 + 2k_2 + 2k_3 + k_4) \end{aligned}$$

is a “standard” choice [86].

The simplest Runge-Kutta method is the *forward Euler method* defined by $E_h(x) = x + hf(x)$. This is a single stage explicit Runge-Kutta method, but it receives limited use for two reasons. First, it only has order one. For example, if $f(x) = x$, then $E_h(x) = (1 + h)x$ and $\phi_h(x) = \exp(h)x = (1 + h + h^2/2)x + O(h^3)$. If $l \rightarrow \infty$ and $h \rightarrow 0$ so that $lh = t$, then the iterates $E_h^n(x)$ converge to $\phi_t(x)$. This expresses the *convergence* of the method. In the example $f(x) = x$ with $x(0) = 1$ and $t = 1$, $E_{1/l}^l(1) = (1 + 1/l)^l$ while $\phi_1(1) = e$. Taking logarithms, we estimate the difference $\log(e) - \log(E_{1/l}^l(1)) = 1 - l \log(1 + 1/l) = -l/2 + o(l)$. The order gives the degree of the lowest order term in the difference between the Taylor series expansion of $\phi_t(x)$ and its computed value from l steps of steplength h with $lh = t$. Computing the value of the vector field to single precision (seven decimal digits) of accuracy for moderate times with a first order method can be expected to take millions of time steps. This can be reduced to a few tens of time steps with a fourth order method. Even with fast computers, the performance of the Euler method is awful.

The second limitation of the Euler method is its instability for stiff systems. This is exemplified by the example $\dot{x} = -\lambda x$ with $E_h(x) = x - h\lambda x$ and $\phi_t(x) = \exp(-\lambda t)$. All of the solution trajectories tend to zero as $t \rightarrow \infty$. However, if $h\lambda > 2$, then the trajectories of the numerical method are unbounded, with oscillating signs. This phenomenon persists in multidimensional linear systems with negative eigenvalues of large magnitude. Many examples, in particular those obtained by discretization of partial differential equations, have rapidly decaying modes whose eigenvalues place stringent limits on the time steps with which the Euler method gives trajectories that even qualitatively resemble the trajectories of a vector field.

Explicit, fixed time-step Runge Kutta methods are only one group of widely used methods. We give here a list of criteria that are used to distinguish and classify numerical integration algorithms.

- **Explicit vs. implicit:**

In explicit methods, the next time step is computed by direct evaluation of function(s) of previously computed data. In implicit methods, the next step is computed by solving a system of equations, often using Newton’s method. The use of implicit methods is motivated by the difficulties of solving stiff systems. Compare the explicit Euler method with the implicit Euler method, defined by $x_{n+1} = x_n + hf(x_{n+1})$. This differs from the explicit Euler method in that the function evaluation takes place at the still unknown next point along the approximate trajectory. The equation for x_{n+1} is implicit since x_{n+1} appears on the right hand side of the equation. For the linear example $\dot{x} = -\lambda x$ with large $\lambda > 0$, the equation is readily solved, giving the formula $x_{n+1} = x_n/(1 + h\lambda)$. When

$\lambda > 0$, $|x_n| \rightarrow 0$ monotonically for any initial condition and any positive step size h . The limitations on step length that were necessary to achieve stability for the explicit Euler method have disappeared, at least for this system. All explicit Runge-Kutta methods applied to a linear equation yield polynomials that are unbounded as the step length increases. Therefore, they all become unstable when applied to the equation $\dot{x} = -\lambda x$ with large enough $\lambda > 0$. When the desired time span for an integration is long enough compared to the step length required for stability by explicit integrators, the differential equation is called *stiff*. Development of stiff integrators was a particularly active research area in the 1970's. The guiding criterion that was applied to this work was Dahlquist's concept of *A-stability* [35], namely that the integrator should remain stable for all positive step sizes when applied to a linear system with negative eigenvalues. Explicit Runge Kutta-methods are not A-stable, a fact that provided strong motivation for improvement of implicit methods.

- **One-step vs multi-step**

One-step methods only use information from the last computed step while multi-step methods use information from several previously computed steps in determining the next step of the integrator. Multi-step methods have the advantage over one-step methods that higher order accuracy can be achieved with a single function evaluation at each time step. On the other hand, theoretical interpretation of one step methods is easier since they can be regarded as giving approximations to the flow map ϕ_h . A k -step multi-step method can be viewed as a discrete mapping on a product of the phase space with itself k times, but it is difficult to single out the class of mappings on this larger space than correspond to multi-step methods. A k -step method also needs a way to compute the first k steps, for example by using a Runge-Kutta algorithm. Implicit multi-step methods called backwards differentiation formulae are used widely as integrators for stiff systems [87, 10].

- **Fixed step vs. variable step**

The most common type of adaptation in numerical integrators is the use of prediction-correction to adjust step size. Variable time step algorithms incorporate criteria for assessing the accuracy of each computed time step. With Runge-Kutta methods, accuracy is commonly assessed by formulating methods of different orders that share intermediate time steps. By comparing the solutions with the principal terms in the asymptotic expansions of the truncation error for each method, an estimate of the error for the time step can be made. If the estimate is larger than a predetermined error criterion, the time step is reduced and the step is repeated. Typically, there are also criteria used to determine when time steps can be increased while maintaining the desired error criteria. For the most part, heuristic arguments and tests with sample problems form the basis of adaptive strategies that are used to vary time steps. As with multistep methods, variable step methods are hard to interpret as discrete approximations to a flow.

The use of variable time step methods is an area in which practice is far ahead of theory. There are few theorems describing the qualitative properties of adaptive time step algorithms viewed as dynamical systems.

Extrapolation is a technique that can be used to improve the order of accuracy of either explicit or implicit integrators. Many integrators have asymptotic expansions in step size h for the errors made in computing ϕ_t as h tends to zero. When numerical computations of ϕ_t are performed with different step sizes, the sequence of computed values can be fit to the beginning of the asymptotic expansion for the errors. These can then be extrapolated to the limit $h = 0$, giving a higher order estimate for ϕ_t . The extrapolation process is independent of the integrator that is used, so that high orders are achievable using simple low-order integrators. Most implementations of extrapolation methods are based upon integrators for which only even terms appear in the asymptotic expansion of the error. The methods can readily vary their order adaptively by selecting the number of intermediate time-steps and the segment of the asymptotic expansion that is fit. These properties give these methods more flexibility in automatically adapting to ongoing computations than high order Runge-Kutta methods.

There is renewed interest in Taylor series methods for numerical integration at this time, some of it motivated by work on verification discussed below in Section 2.3. Series solutions of trajectories are easy to construct theoretically: substitution of a power series expansion $x(t) = \sum a_i t^i$ into the system (1) yields a recursive system of equations for a_i . Implementation requires that the series expansion of $f(x(t))$ is computed, and this is not straightforward. On the one hand, finite difference approximations of the derivatives of f are no more accurate than the Runge-Kutta methods. On the other hand, symbolic differentiation of f produces long complicated formulas that are expensive to evaluate. But there is a third way. *Automatic differentiation* [73] is a technique for computing derivatives with only round-off errors that makes the computation of highly accurate Taylor series approximations to solutions practical. The series can be computed readily to sufficiently high order that the radius of convergence and truncation errors of the Taylor polynomials can be estimated, enabling choice of time steps based entirely upon information at the initial point of the step. Moreover, the methods produce *dense output* in that the Taylor polynomials give the value of the trajectory at all intermediate times to uniform order. In Runge Kutta methods, there is no procedure to directly evaluate the trajectory at fractional time steps while maintaining the order of accuracy of the methods. *A posteriori* tests of the Taylor polynomials can form the basis for adaptive reduction of step size. For example, error criteria can be formulated in terms of the difference between the vector field evaluated along the Taylor polynomial approximation to a trajectory and the tangent vectors to these approximate solutions [32]. While Taylor series methods long ago were demonstrated to work extremely well on a broad range of examples [32], they have not yet been widely adopted. The advent of improved programming languages and environments may change this situation.

Complementary to the classification of numerical integrators as explicit/implicit, one-step/multi-step and adaptive/fixed-step are questions about whether numerical integrators preserve mathematical structure found in special classes of problems. The most intense effort has been devoted to the development of symplectic integrators [134]. A Hamiltonian vector field is one that has the form of Hamiltonian's equations in classical mechanics: $\dot{p} = \partial H / \partial q$ and $\dot{q} = -\partial H / \partial p$ with Hamiltonian function $H : R^n \times R^n \rightarrow R$. The flow of a Hamiltonian vector field is symplectic, meaning that it preserves the two form $\sum dp_i \wedge dq_i$ and energy preserving, meaning that H is a constant of the motion. A symplectic integrator is one for which each time step is given by a symplectic map. The differences between symplectic integrators and other methods become most evident when performing very long time integrations. A common feature of non-symplectic integrators is that the value of H changes slowly along trajectories, but eventually drifts far from its original value [151]. Symplectic integrators do not usually preserve energy either, but the fluctuations in H from its original value remain small. On a deeper level, KAM theory implies that quasiperiodic motions are frequently observed in symplectic flows [134]. Symplectic integrators define maps that satisfy assumptions of the KAM theory while nonsymplectic integrators generally do not. The construction of symplectic integrators is still sufficiently new that it is early to tell how prevalent they will become as the methods of choice for investigation of conservative systems.

2.2 Limit Sets

Classical theories of numerical integration give information about how well different methods approximate trajectories for fixed times as step sizes tend to zero. Dynamical systems theory asks questions about asymptotic, i.e. infinite time, behavior. Only recently has there been emphasis upon understanding whether numerical methods produce good approximations to trajectories over arbitrarily long periods of time [143]. We investigate the question as to when the limits of step size tending to zero and time tending to infinity can be interchanged in numerical computations, but there are additional questions that give a different perspective on long time integration. Two phenomena shape our discussion on the limitations of long time integration. The first phenomenon is based upon a slow drift of numerically computed trajectories from those of the underlying vector field. Consider the explicit Euler method applied to the harmonic oscillator

$$\begin{aligned}\dot{x} &= -y \\ \dot{y} &= x\end{aligned}$$

The flow trajectories are circles, but the non-zero trajectories of the numerical method satisfy $x_{n+1}^2 + y_{n+1}^2 = (1+h^2)(x_n^2 + y_n^2)$ and are all unbounded. The second phenomenon is closely related to structural stability of hyperbolic invariant sets. Hyperbolic invariant sets have the property that trajectories do not remain close. There is a bound $\delta > 0$ so that no two distinct trajectories in the invariant set remain within distance δ of each

other. Given this fact, it is unreasonable to expect that a numerical computation will remain close to the trajectory of its initial condition for all time. Nonetheless, there is a sense in which numerical trajectories can give excellent approximations to trajectories within the invariant set. The concepts of pseudoorbits and shadowing described in this section help explain this apparent paradox.

Further discussion of infinite time behavior of flows and approximating numerical methods will be facilitated by the following definitions and concepts from dynamical systems theory:

- **Invariant set**

Λ is an invariant set if $\phi_t(\Lambda) = \Lambda$ for all t . Λ is forward invariant if $\phi_t(\Lambda) \subset \Lambda$ for all $t > 0$. Λ is backward invariant if $\phi_t(\Lambda) \subset \Lambda$ for all $t < 0$.

- **ω -limit set of a trajectory**

y is in the ω -limit set of the trajectory $x(t)$ if there is a sequence $t_i \rightarrow \infty$ so that $x(t_i) \rightarrow y$.

- **α -limit set of a trajectory**

y is in the α -limit set of the trajectory $x(t)$ if there is a sequence $t_i \rightarrow -\infty$ so that $x(t_i) \rightarrow y$.

- **Wandering point**

x is a wandering point if there is a neighborhood U of x and a $T > 0$ so that $t > T$ implies $x(t) \notin U$.

- **Non-wandering set**

The non-wandering set is the complement of the set of wandering points.

- **(Uniformly) hyperbolic structure**

A hyperbolic structure of a compact invariant set Λ is an invariant splitting of tangent spaces $T_\Lambda R^n = E^s \oplus E^u \oplus E^c$ so that E^c is the one dimensional space tangent to the vector field and for t large $D\phi_t$ expands vectors in E^u at an exponential rate while contracting vectors in E^s at an exponential rate.

The wandering set of a flow is open and the nonwandering set is closed. One of the goals of dynamical systems theory is to decompose the nonwandering set into disjoint closed subsets, called *basic sets*, which have dense orbits. When this can be done, the entire phase space can be partitioned into the stable sets of the basic sets. The stable set of a basic set is the set of points whose ω -limit is in the basic set. Similarly, the unstable set of a basic set is the set of points with α -limit set in the basic set. The geometric characterization of structural stable dynamical systems advanced by the seminal work of Smale [140] gives a large class of systems for which these decompositions have a particularly nice form. On a compact manifold M , structurally

stable systems have a finite number of basic sets Λ_i , each of which posses a uniformly hyperbolic structure.

Chaotic dynamical systems display sensitive dependence to initial conditions: nearby trajectories diverge from one another, typically at exponential rates. In the presence of sensitive dependence to initial conditions, it is hardly reasonable to expect that a numerical method will accurately track the trajectory of its initial condition for long times, since trajectories of nearby initial conditions do not remain close to the chosen one. Any error made in a single step of a numerical integration will be amplified by the inherent divergence of trajectories. This fact underlies fundamental limitations in the accuracy of numerical integration over long times. In hyperbolic invariant sets, it is inevitable that errors in the numerical solution of the initial value problem grow exponentially. This is even true for iteration of diffeomorphisms where there is no truncation error of numerical integration, only round-off error in the evaluation of the diffeomorphism.

The effects of sensitivity to initial conditions prompt new perspectives on the initial value problem. Over short times, we expect numerical integration to be accurate. What positive results can be established about long time integration? In the case of hyperbolic invariant sets, there is a satisfying theory whose ultimate conclusion is that numerical trajectories approximate actual trajectories of a different initial condition. The concepts of *shadowing* and *pseudoorbit* [20] have been used to explore these issues. For a discrete dynamical system defined by the mapping $F : R^n \rightarrow R^n$, a δ -pseudoorbit is a sequence of points x_i with the property that $|x_{i+1} - F(x_i)| < \delta$. On each iterate, there is an error of at most δ in the location of the next point relative to the location of the mapping applied to the current point of the pseudoorbit. If there is a point y whose trajectory has the property that $|F^i(y) - x^i| < \epsilon$, then we say that the trajectory of y ϵ -shadows the pseudoorbit. The extension of the shadowing concept and this theorem to a flow Φ requires allowance for time “drift” along trajectories. If $(x_k, h_k) \in R^n \times R$ is a sequence of points and time increments, it is a δ -pseudoorbit if $|\Phi(x_k, t_k) - x_{k+1}| < \delta$ for each k . The pseudoorbit ϵ -shadows the orbit of y if there are times t_k with $|x_k - \Phi(y, t_k)| < \epsilon$ and $|h_k - (t_{k+1} - t_k)| < \epsilon$ [30]. For a numerical iteration or integration, we can view the algorithm as producing a pseudoorbit. The one step accuracy of the method determines a δ for which the numerical trajectory is a pseudoorbit. We can ask for which systems, which one step methods ψ_h , which initial conditions x and which ϵ there is a point y so that $\Phi(y, nh) \epsilon$ -shadows the numerical trajectory $\psi_h^n(x)$.

The qualitative characteristics of the invariant sets of a flow Φ are a big factor in determining whether they satisfy shadowing properties. Hyperbolic basic sets do. Here is the statement of a result for discrete time systems.

Theorem 1 [20] *Let Λ be a hyperbolic invariant set of a C^1 diffeomorphism F . If $\epsilon > 0$, there is a $\delta > 0$ such that every δ -pseudorobit for a trajectory in Λ ϵ -shadows a trajectory of F in Λ*

Numerical trajectories that start near a hyperbolic attractor Λ will stay near Λ and they will be shadowed by trajectories within the attractor. Thus, the shadowing property of

hyperbolic sets enables us to recover long time approximation properties of numerical trajectories when they are computed with sufficient accuracy for fixed, moderate times. This theorem is very satisfying mathematically, but we note with caution that there are few examples of hyperbolic attractors that arise from physical examples. The chaotic attractors that have been observed in applications seem not to be hyperbolic and structurally stable. The discrete Henon mapping is the example that has been studied most intensively [89]. The details of its dynamical properties are much more subtle and complex than those of hyperbolic attractors [20], but they appear to be typical of chaotic attractors with a single unstable direction. Additional complexity is present in partially hyperbolic attractors in which the dimension of the unstable manifolds of points vary within attractors [2, 127]. Hammel et al. [88] have investigated the shadowing properties of one dimensional mappings and the Henon mapping. They demonstrate that very long sections of trajectories have the shadowing property, but that one cannot expect it to hold for infinite time.

Coombes et al. [30] implemented methods for shadowing trajectories of vector fields. Convergence of a numerical integration algorithm proves that it provides shadowing of trajectories for fixed, finite times. Nonetheless, the theory of these algorithms says little about the precision of a pseudoorbit required to provide an ϵ -shadow of a trajectory. The ability to shadow trajectories for long periods of time is closely related to their Liapunov exponents (discussed in Section 2.3) and exponential dichotomies. If there are no Liapunov exponents close to zero, then an infinitesimal neighborhood of a trajectory $x(t)$ can be decomposed into unstable directions that diverge from $x(t)$ and stable directions that converge towards $x(t)$. Deviations from $x(t)$ in the stable directions (eventually) become smaller in the forwards direction while deviations from $x(t)$ in the unstable directions become smaller in the backwards direction. Heuristically, to find an orbit shadowed by a pseudo-orbit, take the trajectory that matches the projection of the pseudoorbit onto the stable directions at its beginning and onto the unstable directions at its end. Several authors, including [28, 29, 30, 31, 88], have formalized this conceptual framework to give explicit estimates for the shadowing constants of a vector field.

Consider a set of points (x_k, t_k) , $k = 0, \dots, N$ along a trajectory of the vector field f with flow Φ . These points satisfy the equations

$$\Phi(x_k, t_k) - x_{k+1} = 0 \tag{3}$$

for $k = 0, \dots, N - 1$. We can view the left hand side of these equations as a map $F : R^{(n+1)(N+1)} \rightarrow R^N$. A one-step numerical method for integrating the vector field gives an approximation to this map. The analysis of Coombes et al. [30] apply Newton's method and its extensions to analyze how much the solutions of the system (3) change with perturbations of Φ . To deal with the flow direction itself, these authors constrain the points x_k to lie on a fixed set of cross-sections to the vector field. The important quantities in determining the shadowing data are the magnitude of a right inverse to the map F and the C^2 norm of the vector field in a neighborhood of the trajectory. The right inverse of F is not determined uniquely: there are essentially n degrees of

freedom that specify a trajectory. To obtain a right inverse whose norm is relatively small, the trajectory is decomposed into its expanding and contracting directions. The contracting coordinates are chosen at the beginning of the trajectory and the expanding coordinates are chosen at its end. Using their methods Coombes et al. [30] demonstrate shadowing of very long trajectories in the Lorenz system [113].

2.3 Error Estimation and Verified Computation

Numerical integration algorithms are fundamental tools for the investigation of dynamical systems, but the results they produce are seldom subjected to verification or rigorous error estimation. Indeed, the exponential divergence of trajectories in systems with sensitive dependence to initial conditions sets limits on the time for which one can expect a numerical trajectory to remain close to the actual trajectory with the same initial condition. Naive attempts at estimating the errors of numerical integration tend to introduce artificial instability coming from varied sources such as the rectilinear geometry inherent in interval arithmetic [33]. This *wrapping* phenomenon amplifies the expected exponential growth of errors, typically producing pessimistic results. This state of affairs creates a tension between simulation and the mathematical theory of dynamical systems. On the one hand, numerical integration seems necessary in the investigation of systems that don't have analytically explicit solutions. On the other hand, the difficulty in estimating the errors in these integrations makes it hard to use simulations in rigorous analysis. There has been recent progress in attacking this issue, and the number of successful examples in which numerical computation gives rigorous results about dynamical systems is growing steadily.

The variational equations $\dot{\xi}(t) = Df_{x(t)}\xi(t)$ of a trajectory $x(t)$ for system (1) give an infinitesimal picture of how nearby trajectories differ. The exponential growth rates of solutions to this time varying linear system of differential equations are the *Liapunov exponents* of the trajectories. Their existence is discussed in the next section. The computation of Liapunov exponents must contend with two phenomena. The first phenomenon is that the rates of expansion and contraction along a trajectory may vary dramatically. This is particularly evident in canard solutions to systems of differential equations with multiple time scales [47, 48, 57], also discussed in the next section. The second phenomenon is that the directions of expansion and contraction may twist along a trajectory. Changes in twist are intimately involved in the bifurcations that take place in chaotic invariant sets. Both of these phenomena are common, so general algorithms for computing Liapunov exponents should take them into account. Suppose that (x_k, t_k) is a sequence of points along a trajectory and that D_k is the Jacobian of the flow map from (x_k, t_k) to (x_{k+1}, t_{k+1}) . Denoting by $\sigma_1, \dots, \sigma_n$ the singular values of $J_N = D_N \cdots D_2 D_0$, the Liapunov exponents are defined as the limits of $\log(\sigma_i)/t_N$ as $N \rightarrow \infty$. Thus computation of the Liapunov exponents requires computation of the singular values of J_N . Performing this computation by first computing the product and then computing the singular values is ill-conditioned in general. If the largest Liapunov exponent is separated from the remaining ones, the J_N tend to rank one matrices and

small perturbations of the D_i produce large relative changes in the magnitudes of its smaller singular values. To accurately compute the smaller singular values, two basic strategies have been proposed [136, 152]. The first is to form exterior powers of the Jacobians that represent its action on subspaces. The dominant singular values of exterior power i will be the product of the largest i singular values, so the Liapunov exponents can be recovered from the ratios of the largest singular values of the exterior powers. The second strategy is perform a matrix decomposition of J_N by working with its factors iteratively. For example, the QR factorization of J_N can be computed by first computing the QR factorization $D_0 = Q_0 R_0$. Then one computes the QR factorization of $D_1 Q_0 = Q_1 R_1$. Proceeding inductively, one computes $D_k \cdots D_0 = Q_k R_k \cdots R_0$. The inductive step is to compute the QR factorization $D_{k+1} Q_k = Q_{k+1} R_{k+1}$. Dieci et al. [41] have proposed a continuous analog for this factorization of J_N based upon solving Riccati equations. The goal is produce a frame, i.e., a smoothly varying set of orthogonal coordinate systems along the trajectory, so that the variational equations become triangular when transformed to this frame. The differential equations satisfied by the frame are a *Riccati* system for which there are special methods of numerical integration [39, 111].

The most direct approaches to error estimation for numerical integration are based upon interval arithmetic. Numbers are replaced by intervals and operations are replaced by enclosures. For example, the sum of two intervals is an interval that contains the sum of any numbers contained in the two summands. For calculations involving a moderate number of operations, interval arithmetic is often an effective means of obtaining rigorous estimates for calculations. As an example, interval implementations of Newton's method often work well to give precise estimates on the location of all the zeros of a function, including proofs of their existence. Within the context of dynamical systems, Lanford's computer based proof for the existence of a fixed point for the period doubling renormalization operator on unimodal functions [108, 109] exploits the application of interval arithmetic to Newton's method. Application of interval arithmetic to numerical integration is an old idea [33], but the results are frequently disappointing, leading to poor bounds on the computation of a trajectory compared to the apparent accuracy of the calculation. Unless the flow is uniformly contracting in the phase space, each step of an interval based numerical integration tends to produce larger enclosing intervals. The continued growth of enclosing intervals limits severely the number of time steps that can be taken before the bounds become useless. The work on shadowing described above and three additional examples illustrate ways of circumventing these limitations of interval arithmetic.

The first example comes from the work of Löhner. He replaces intervals by Taylor series (or jets) augmented by bounds that enclose function values. These Taylor series with bounds are the fundamental objects with which computations are performed. Thus, a function $f : R^n \rightarrow R$ is approximated on a domain $D \subset R^n$ by a polynomial P and $\epsilon > 0$ with the property that $|P(x) - f(x)| < \epsilon$ for all $x \in D$. This is a much richer class of objects than intervals, and it is possible to construct a precise numerical calculus for these objects in the context of floating point arithmetic. Lanford's numerical proof

of the Feigenbaum conjectures uses these data structures [108]. Berz has implemented this calculus in his COSY software, using the term differential remainder algebra [15], and applied it to normal form calculations of Hamiltonian systems to achieve strong estimates of the stability properties of accelerator designs [115]. Löhner uses algorithms that compute Taylor series approximations of trajectories for differential equations with automatic differentiation, and then obtains error estimates for these approximations. The error estimates come from an adaptation of the contraction operator used in the Picard proof of existence of solutions to ordinary differential equations. If $f : R^n \rightarrow R^n$ is Lipschitz continuous and T is small enough, then the operator

$$\Pi(g)(t) = x_0 + \int_0^t f(g(s))ds$$

acting on continuous functions $g : [0, T] \rightarrow R^n$ with $g(0) = x_0$ has a contracting fixed point at the solution to the differential equation $\dot{x} = f(x)$ with initial condition x_0 . Bounds on T are readily computed in terms of the domain, magnitude and Lipschitz constant of f . Löhner applies this operator in the context of Taylor series with remainders to obtain good bounds on Taylor series approximations to the solution of the differential equation. This work complements the use of automatic differentiation to produce high order numerical integration algorithms based upon Taylor series.

The second example of error estimates for solutions of differential equations exploits transversality in the context of planar dynamical systems. Guckenheimer and Malo [80] observed that numerical integration of rotated vector fields [54] can be used to compute curves that are transverse to the trajectories of a planar vector field. Using the terminology of Hubbard and West [94], we expect the numerical trajectories of the rotated vector fields to be fences that provide barriers the trajectories of the original vector field can cross only once. Interval arithmetic can be used to verify that the trajectories of the rotated vector fields are indeed transverse to the trajectories of the original vector field. The advantage of this method compared to direct error estimation of trajectories is that the transversality computations are all local to individual time steps of the rotated vector field. The interval estimates for each time step are independent of one another, so the growth of the estimates does not propagate from one time step to the next. Using these ideas, Guckenheimer [75] described an algorithm to rigorously verify the correctness of phase portraits of structurally stable planar vector fields defined by functions for which interval evaluations have been implemented.

The fourth example of rigorous results based upon numerical error estimates is the recent analysis of the Lorenz system by Tucker [146]. In 1963, Lorenz described the first strange attractor that was observed via numerical computation [113]. A more complete geometric model of the Lorenz attractor was formulated fifteen [84] years later. Verification that the assumptions underlying the geometric model are satisfied by the Lorenz system has been a benchmark problem in the numerical analysis of dynamical systems. The key assumption can be expressed as the statement that there exists a family of cones in the tangent spaces of points in the attractor that are forward invariant. Tucker has solved this problem. His work is based upon a careful dissection

of a neighborhood of the attractor into regions on which the behavior of the variational equations can be described with interval computations. The interval analysis itself does not need to be precise once a suitable covering of the attractor has been constructed.

3 Computation of Invariant Sets

Invariant sets with complex geometry are common in dynamical systems. The complexity comes both in the local structure of the sets and in convoluted shapes of smooth objects. Smale's horseshoe [141] is an important example of a fractal invariant set for a discrete time dynamical system. The analog of the horseshoe in continuous time dynamical systems is the *solenoid* [140]. The stable and unstable manifolds of periodic orbits in the horseshoe and solenoid are folded, with regions of arbitrarily large curvature. The Lorenz system [113] has been a rich source of complex geometric objects, including the convoluted, two dimensional stable manifold of the origin and its fractal, chaotic attractor. The algorithmic aspects of computing invariant sets is a relatively new subject compared to numerical integration. The list of successful methods for directly computing invariant sets has been growing. This section surveys research on computing four types of invariant sets: periodic orbits, invariant tori, stable and unstable manifolds of equilibria and chaotic invariant sets.

3.1 Periodic Orbits

A *periodic orbit* of a flow ϕ is a non-equilibrium trajectory $x(t)$ with $x(0) = x(T)$ for some $T > 0$. The minimal value of T is the period of the orbit. Perhaps the most common way of finding stable periodic orbits is to identify them as the limit sets of trajectories computed by numerical integration. However, there are circumstances in which numerical integration does not yield an accurate representation of a stable periodic orbit. An example is given below. There are also circumstances in which theoretical considerations suggest that the computation of periodic orbits with numerical integration may fail. A one-step numerical integration method with fixed steps is a diffeomorphism, and periodic orbits of flows become invariant curves of the time h map ϕ_h of the flow. The theory of normal hyperbolicity implies that if the periodic orbit is hyperbolic, then the map defined by a numerical integration algorithm will have an invariant curve near the periodic orbit if it is a sufficiently accurate approximation of ϕ_h . On this invariant curve there may be resonance, with numerical trajectories converging to a stable, discrete periodic orbit with a finite number of points rather than filling the periodic orbit densely. This cannot happen for ϕ_h if h is incommensurate with the period. When T is a multiple of h , the periodic orbit is a continuous family of periodic orbits of ϕ_h . This qualitative discrepancy between the generic behavior of a numerical integration algorithm and the flow map seldom impedes the identification of stable periodic orbits as the limits of numerical trajectories. Nonetheless, when trying to compute periodic orbits whose stability is weak enough, normal hyperbolicity

breaks down and the numerical algorithm may acquire more complex limit sets close to the periodic orbit. In particular, using numerical integration to accurately identify the location of saddle-node bifurcations of periodic orbits is problematic.

It is desirable to have direct methods for locating periodic orbits for at least three reasons:

- Numerical integration may fail as described above.
- Direct methods may be more efficient than numerical integration for computing periodic orbits.
- Unstable periodic orbits are not readily obtainable as limit sets of trajectories.

The equations defining periodic orbits are *boundary value problems* for ordinary differential equations. Most of the extensive literature and software dealing with boundary value problems applies to two point boundary value problems with separated boundary conditions [9]. While the equations for periodic orbits can be recast in this form by enlarging the dimension of the phase space, this approach has not been successfully applied to many problems. Instead, most studies of periodic orbits use algorithms that are specifically designed for the solution of boundary value problems with periodic boundary conditions. There are two methods that dominate these studies: simple shooting methods and the collocation method implemented in the computer code AUTO [50]. For these methods to work well, the periodic orbit should itself be robust with respect to perturbation: the periodic orbit should vary continuously with deformations of the vector field. A sufficient condition for this robustness can be formulated in terms of the monodromy matrix of the orbit γ . The monodromy matrix A of a point $x \in \gamma$ is the Jacobian of the time T flow map at x . The monodromy matrix A always has 1 as an eigenvalue (with $f(x)$ as eigenvector), but if 1 is a simple eigenvalue of A , then the periodic orbit perturbs smoothly with perturbations of f . A periodic orbit for which 1 is a simple eigenvalue of A is called *regular*.

The periodic orbit equations of a vector field do not have isolated solutions. If $\gamma(t)$ is a periodic orbit and $c \in R$, then $\gamma(t + c)$ is also a periodic orbit. To obtain a regular system of equations for points approximating a periodic orbit, this degeneracy coming from translation in time must be removed. In *simple shooting* methods, this is typically done by restricting initial conditions to lie in a cross section to the periodic orbit. The *return map* for the cross section is defined by mapping each point to the next point on its trajectory lying in the cross section. Simple shooting methods compute the return map using a numerical integration algorithm. This is augmented by using a root finding algorithm such as Newton's method to compute a fixed point for the return map. Implementing simple shooting methods is straightforward, but their performance is subject to numerical limitations that are explained below. There have been many implementations of simple shooting methods, for example in the code LOCBIF [102]. There is one detail worth noting. Computing the return map requires

that the intersection of a trajectory with the cross-section be computed. Many numerical integration algorithms do not yield values at intermediate points of a trajectory between time steps with the same accuracy as those at the time steps. In this case, interpolation using several computed points of the trajectory or a change of time step that yields a point on the trajectory can be used to complete the calculation of the return map. If the cross-section is given by the equation $x_k = c$, E. Lorenz observed that one can rescale the vector field by dividing by its k th component to obtain a vector field near the cross-section in which the k th component evolves with unit speed. Choosing the time step for the rescaled vector field to be $c - x_k$ gives a time step that ends on the cross-section.

The numerical difficulties with simple shooting come from two sources. The first source of difficulty is the accuracy of the return map. Newton's method requires the Jacobian J of the return map. If J is computed with finite differences, this can make application of Newton's method to the return map erratic. The second source of difficulty with simple shooting is due to the potential ill-conditioning of the problem. The return map may have a Jacobian with very large norm at the intersection of a periodic orbit with a cross section. This norm can readily become large enough that changes in initial condition of unit precision produce changes in the value of a return map that are larger than a desired error tolerance in the fixed point procedure used to locate the periodic orbit. In the canard example discussed below, integration over part of the cycle produces an extremely ill-conditioned flow map. While simple shooting works with many problems despite these potential difficulties, more elaborate methods for computing periodic orbits are often required.

Multiple shooting algorithms address the difficulties associated with ill-conditioning of the return map. Instead of solving the single equation $\Phi(x, t) = x$ for a fixed point of the return map, one seeks a set of points (x_k, t_k) , $k = 0, \dots, N$ with $t_0 = 0$, $x_N = x_0$ and $\Phi(x_k, t_{k+1} - t_k) = x_{k+1}$ for $k = 0, \dots, N - 1$. This forms a system of nN equations in $(n + 1)N$ variables. For a regular periodic orbit, there is a smooth N dimensional manifold of solutions coming from different choices of x_k on the periodic orbit. Multiple shooting algorithms either constrain the (x_k, t_k) to lie in a set that yields a unique point on the periodic orbit, or they augment the system of equations to yield a regular system of $(n + 1)N$ equations. Conceptually, multiple shooting is simple. By judicious choice of the length of segments along the periodic orbit, the condition number of the system of equations $\Phi(x_k, t_{k+1} - t_k) = x_{k+1}$ can be reduced to manageable levels on problems where simple shooting fails completely. We illustrate this with an example.

Problems with multiple time scales are especially prone to ill-conditioned return maps since the periods of the periodic orbits may be very long when measured in the faster time scales of the problem. The following example exhibits stable periodic orbits that cannot be computed readily with numerical integration or a simple shooting method. These orbits are examples of trajectories called *canards* due to their visual appearance in a family of vector fields that generalizes the van der Pol equation [47, 48].

Figure 1: Periodic orbits of the extended van der Pol equation. The orbits which follow the middle, unstable branch of the “slow” manifold are canards.

Figure 2: Fourth order Runge-Kutta integration of the canard vector field. Due to the massive instability of the unstable branch of the slow manifold, the integration is unable to compute trajectories that follow this branch. Compare with the family of canards shown in Figure 1.

The vector field

$$\begin{aligned}\dot{x} &= \frac{1}{\epsilon}(y - x^2 - x^3) \\ \dot{y} &= a - x\end{aligned}$$

is a translate of the van der Pol equation when $a = \frac{-1}{3}$ and has two time scales when $\epsilon > 0$ is small. The single equilibrium point at $(a, a^2 + a^3)$ undergoes supercritical Hopf bifurcation at $a = 0$ with decreasing a . Figure 1 shows periodic orbits from this family with $\epsilon = 0.001$, computed with a multiple shooting algorithm based upon automatic differentiation. The periodic orbits that emerge grow very quickly in a to a relaxation oscillation approximated by a pair of segments that follow the nullcline $y = x^2 + x^3$ and a pair of horizontal segments that are tangent to the nullcline. The growth of the periodic orbits occurs during an interval of a whose width shrinks to zero proportionally to $\exp(-1/\epsilon)$ [47]. The cycles of intermediate size in the family have a segment that follows the unstable middle portion of the nullcline with $x \in (-2/3, 0)$. In this region, trajectories of the vector field diverge rapidly from one another at a rate comparable to $\exp(-1/\epsilon)$. This divergence makes it essentially impossible to compute a trajectory that follows the nullcline using numerical integration forwards in time with double-precision (64 bit) floating point arithmetic. Starting near the local minimum of the nullcline, an error in the computation of one step comparable to the unit precision of the floating point arithmetic produces a point whose trajectory will leave a moderate sized neighborhood of the nullcline in a time comparable to 50ϵ , the factor 50 being approximately the logarithm of the unit precision. For $\epsilon = 0.001$, this time is approximately 0.05, only long enough for y to travel a short distance up the nullcline at the rate $|\dot{y}| = a - x$. Consequently, forward numerical integrations of trajectories appear to “peel off” from the unstable portion of the nullcline without traveling very far along it. Varying a appears to produce almost a discontinuity in the numerical ω -limit sets. Attempts to find the canard solutions with Runge-Kutta algorithms find a narrow interval of a of width comparable to 10^{-14} in which the integration algorithm becomes chaotic, erratically producing a combination of small and large loops along the same numerical trajectory. Such a trajectory is illustrated in Figure 2. Thus, despite the fact that the canard cycles are stable, they cannot be computed readily with numerical integration. Guckenheimer et al. [77, 81] have been developing new algorithms for the computation of periodic orbits based upon

the use of Taylor series and automatic differentiation. Automatic differentiation is used to achieve high orders of accuracy while maintaining coarse meshes and relatively small systems of equations to solve. We describe here the multiple shooting algorithm used to produce Figure 1. Approximations to the periodic orbit are parametrized by a discrete mesh of points (x_k, t_k) near the periodic orbit, including the times associated to these points. Numerical integration from one mesh point to the next is performed using Taylor polynomial approximations to the trajectories, with step sizes chosen adaptively so that estimated errors of the trajectory lie below a predetermined threshold. The Taylor series and the Jacobians of the Taylor series coefficients with respect to variation of the mesh point are computed using procedures that are part of the program ADOL-C developed at Argonne National Laboratory [73]. The variational equations for the orbit are solved along with the original differential equation. The system of equations $\Phi(x_k, t_{k+1} - t_k) = x_{k+1}$, $k = 0, \dots, N$ is augmented with N additional equations. These additional equations express constraints that force updates of (x_k, t_k) to be orthogonal to the trajectory of (x_k, t_k) in R^{n+1} . The tangent vector to this trajectory is given by $(f(x_k), 1)$. (In the case of $k = 0$ where we fix $t_0 = 0$, the times $t_1, \dots, t_N - 1$ are varied instead of t_0 .) For a regular periodic orbit, the augmented system of equations is regular and Newton's method is used to solve it. The algorithm requires an initial approximation to a periodic orbit as starting conditions, say a set of points obtained from a numerical integration. New mesh points are created adaptively when the norm of the Jacobian from the previous mesh point exceeds a specified bound. Mesh points are removed when the remaining points adequately represent the periodic orbit. Continuation, described in the next section, is easily implemented in this algorithm. Automatic differentiation is used to compute derivatives of the Taylor series coefficients with respect to parameters as well as with respect to the phase space variables. A parameter is regarded as an additional independent variable in the augmented system of equations. Continuation steps are taken by computing the tangent vector to the curve of solutions of the augmented equations and adding an increment of this tangent vector $v \in R^{(n+1)N+1}$ to the current solution. Yet one more equation is added to the augmented system, constraining Newton updates to be orthogonal to the tangent vector v .

Global boundary value methods for computing periodic orbits project the equation $\dot{\gamma}(t) = (1/T)f(\gamma(t))$ onto finite dimensional approximations of the space Γ of smooth functions $\gamma : S^1 \rightarrow R^n$. Alternatively, we can try to solve $\dot{\gamma}(t) = f(\gamma(t))$ for curves with $\gamma(0) = \gamma(T)$. In both alternatives, the period T is an unknown that is part of the equations to be solved. The global method used in the code AUTO [52, 53] is collocation. The approximations to Γ consist of continuous, piecewise polynomial functions. A discretization of γ is determined by a mesh of N points $(x_1, \dots, x_N)$ and N time intervals $t_1, \dots, t_N$. Inside each of the N mesh intervals k collocation points are chosen. In AUTO, as well as the boundary value solver colsys and its descendants [9, 52, 53], the collocation times for each mesh interval are at Gauss points in order to produce *superconvergence* of the method. The polynomial function on each mesh interval is required to satisfy the differential equation at the collocation points and to be continuous at the endpoints of the mesh intervals. For an n dimensional

system, this yields a total of $n * N * (k + 1)$ independent equations on the same number of variables. The orbit period is one additional independent variable, and it is balanced by an equation that removes the degeneracy associated with time translation. This can be done by restricting a point of the orbit to a cross section or it can be given by an integral *phase condition* in the form $\int g(\gamma(s))ds = 0$ for a function $g : R^n \rightarrow R$. The system of $n * N * (k + 1) + 1$ equations defined by the collocation algorithm is quite large. Fortunately, the Jacobian of this system is sparse. AUTO exploits the sparsity using a special Gaussian elimination procedure in its Newton iterations to obtain solutions of the periodic orbit equations.

The accuracy and efficiency of the Taylor series methods are demonstrated in the following example. The planar vector field

$$\begin{aligned}\dot{x} &= y - y^2 - x(x^2 - y^2 + 2y^3/3 + c) \\ \dot{y} &= x + (y - y^2)(x^2 - y^2 + 2y^3/3 + c)\end{aligned}$$

has a periodic orbit that lies in the zero set of the polynomial $h = x^2 - y^2 + 2y^3/3 + c$ when $c \in (0, 1/3)$. The value of h along a computed curve measures the distance of the curve from the periodic orbit. Figure 3 displays values of h along approximate periodic orbits with $c = 0.07$ and period approximately 7.7 computed in three ways. The top panel presents values of h at the 60 mesh points of an AUTO calculation of the periodic orbit. The data is representative of the most accurate approximations to the orbit produced by AUTO while varying the number of mesh points, the number of collocation points and the error tolerances allowed by the algorithm. The middle panel shows three numerical trajectories computed with a fourth order Runge-Kutta algorithm using step sizes of fixed lengths 0.00125, 0.001 and 0.0001. The step size 0.001 appears to provide close to optimal accuracy for this method, with round-off errors apparently dominating truncation errors for smaller step sizes. The lower panel displays the results of a calculation using a multiple shooting algorithm based on automatic differentiation. There are five mesh intervals with the solution approximated on each half of a mesh interval by the degree 16 Taylor series polynomial at the boundary mesh points. Convergence was obtained in 6 Newton steps starting with mesh points that are far from the computed solution. The maximum value of $|h|$ in the three calculations is approximately 6×10^{-11} , 8×10^{-15} and 6×10^{-16} . The Taylor series methods achieve the best accuracy, comparable to the inherent precision of the floating point arithmetic, with surprisingly coarse meshes. Further development of methods based upon Taylor series appears to be very promising.

3.2 Invariant Tori

Invariant tori are a prominent feature of symplectic flows and also arise through Hopf bifurcation of periodic orbits in dissipative systems. For flows with a global cross-section on a two dimensional torus, a fundamental invariant is the winding number,

Figure 3: Three different methods have been used to compute a periodic orbit that lies along the level curve $h = 0.07$ of the polynomial $h = x^2 - y^2 + 2y^3/3$. Values of h are plotted as a function of time during one traversal of the periodic orbit. The numerical periodic orbit in the top panel was computed with AUTO, the three in the middle panel with a fourth order Runge-Kutta method, and the bottom one with a multiple shooting algorithm employing automatic differentiation.

or equivalently the rotation number of a return map [90]. This invariant is rational if and only if the flow has periodic orbits. If the flow is C^2 and the winding number is irrational, then all trajectories of the flow are dense [38]. KAM theory [90] and the theory of normal hyperbolicity [93] provide theoretical tools for the analysis of invariant tori. There is a modest body of research on algorithms for computing these objects, much of it framed in the context of invariant curves of discrete time systems. Three different approaches, quite different from one another, are discussed here.

The first approach to computing invariant tori of discrete maps has been to represent one dimensional tori as graphs of functions and to formulate a system of equations that gives a finite dimensional approximation to the invariance of these curves. This approach has been pursued in different ways. KAM Theory restricts attention to invariant tori on which the motion is conjugate to irrational rotation and solves for the Fourier series of the conjugacy. In the case of invariant curves, piecewise polynomial approximations of invariant curves lead to general algorithms that apply to invariant curves that contain periodic orbits as well as tori that have irrational rotation numbers. Implementations of such methods have been described by Kevrekidis et al. [101] and van Veldhuizen [147]. Their results indicate that it is difficult numerically to follow a family of invariant tori to the point at which they begin to lose smoothness and disappear. Aronson et al. [8] give a comprehensive description of ways in which tori with rational winding numbers can lose smoothness.

The second approach to computing invariant tori was pioneered by Greene [72]. This method seeks to compute invariant tori in symplectic systems by approximation with periodic orbits. Most of the research has concentrated on area preserving diffeomorphism of the plane. KAM theory proves that each periodic orbit of elliptic type is surrounded by a family of invariant curves with irrational rotation numbers satisfying diophantine inequalities [72, 155]. Each of these invariant curves is the limit of periodic orbits whose rotation numbers are obtained by truncating the continued fraction expansion of the irrational rotation number. Periodic orbits of high period are computed with root finding algorithms analogous to either shooting or global boundary value methods. Estimates of the convergence of the approximating periodic orbits to the (unique) invariant measure of the invariant curve gives information about the structure of the invariant curve. If the diffeomorphism depends upon a parameter, some of the invariant curves may evolve into Cantor sets. Renormalization methods have been applied to study this transition, especially for invariant curves for the golden mean and other rotation numbers with periodic continued fraction expansions [110]. The numer-

ical computations of these “last” invariant curves have been based upon computations of approximating periodic orbits.

Dieci et al. [45, 42, 44, 43] have investigated the computation of invariant tori for vector fields. Their starting point has been the formulation of a partial differential equation that implies the invariance of the torus. This partial differential equation states that the vector field is tangent to the torus. The innovative aspects of these studies lie in using algorithms for solving PDE’s to address this problem. The torus is represented as the image of a mapping on a discrete grid, and then partial differential equations are approximated to yield a set of equations for this mapping analogous to global methods for computing periodic orbits. Implementations of the algorithms have been tested on a few examples like the forced van der Pol equation, but experience as to the domain of problems for which these algorithms work well remains limited.

3.3 Stable and Unstable Manifolds

Stable and unstable manifolds of equilibrium points and periodic orbits are important objects in phase portraits. In physical systems subject to disturbances, the distance of a stable equilibrium point to the boundary of its stable manifold provides an estimate for the robustness of the equilibrium point. The closer the boundary, the more likely disturbances will kick the system out of the basin of attraction of the equilibrium. In the simplest situations, these boundaries are formed by stable and unstable manifolds of saddles. In more complex situations, the basin boundaries are fractal, chaotic invariant sets containing large numbers of periodic orbits and their stable manifolds. Thus, there is great interest from both theoretical and practical perspectives in computations of stable and unstable manifolds.

From a naive perspective, it would appear that the computation of stable and unstable manifolds of equilibria is no more difficult than numerical integration. For one dimensional manifolds this is true. One dimensional stable and unstable manifolds of equilibria of flows consist of pairs of trajectories, so their computation can be implemented by applying an initial value solver to a well chosen initial condition. Higher dimensional stable and unstable manifolds are harder to compute. The two dimensional stable manifold of the origin for the Lorenz system [113]

$$\begin{aligned}\dot{x} &= 10(y - x) \\ \dot{y} &= 28x - y - xz \\ \dot{z} &= -\frac{8}{3}z + xy\end{aligned}$$

has served as a benchmark problem. There are two difficulties in computing this manifold. First, the stable eigenvalues at the origin of this system are approximately -2.67 and -22.8 with a ratio that is approximately 8.56. As a result, backwards trajectories in the manifold tend to flow parallel to the strong stable direction. Numerical integration of initial conditions in the stable manifold uniformly clustered near the origin

produces only a strip along the strong stable direction. The second difficulty in computing this stable manifold is that it becomes highly convoluted far from the origin. Part of the manifold spirals around the z -axis while part of it curls around the stable manifolds of the equilibria located at $(\pm 6\sqrt{(2)}, \pm 6\sqrt{(2)}, 27)$ [126].

Symbolic methods can be used to compute high order approximations to the Taylor series of stable and unstable manifolds at equilibrium points. One approach to these algebraic calculations is to subsume the computation of stable and unstable manifolds of equilibria into the linearization problem: finding a smooth coordinate transformation that transforms the system $\dot{x} = f(x)$ into a linear system of equations near an equilibrium. In the transformed coordinates, the stable and unstable manifolds are linear subspaces. Formally, the linearization problem can be reduced to a sequence of systems of linear equations for the Taylor series of the coordinate transformation [78]. These linear systems degenerate if the eigenvalues λ_i at the equilibrium satisfy *resonance* conditions of the form

$$\lambda_i = \sum_{j=1}^n a_j \lambda_j$$

with non-negative integer coefficients a_j . The order of the resonance condition is $\sum a_j$. When resonance conditions are satisfied, transformation to *normal forms* containing only nonlinear terms associated with the resonance conditions can still be accomplished but the system can only be linearized with finite smoothness related to the order of the resonance conditions [96]. Transformation to simpler nonlinear systems, called *normal forms*, is used extensively in the analysis of bifurcations [78, 131].

Algebraic computation of linearizations and normal forms are readily implemented in symbolic systems for vector fields of moderate size [133]. Nonetheless, the complexity of these computations grows quickly with problem size. For large problems, instead of computing a full linearization, one would like to extract more limited information. Problems are common for which almost all modes are highly damped and a low dimensional submanifold in the phase is attracting. These problems often arise from investigations of instability when a system is driven by external forces until its attractors are time dependent, but not highly disordered. Beyn and Kless [17] have examined the computation of low dimensional invariant manifolds within this context. They study the use of iterative methods in linear algebra to compute the location and normal forms of invariant manifolds while avoiding such operations as the factorization or inversion of the full system Jacobian at an equilibrium.

The most complete work on computing stable and unstable manifolds has been done in the context of one dimensional stable and unstable manifolds of fixed points of discrete time dynamical systems. These methods have been applied to the return maps of periodically forced continuous systems and to computation of two dimensional stable and unstable manifolds of periodic orbits [24]. One dimensional stable and unstable manifolds of fixed points for maps have fundamental domains: if the eigenvalue of the manifold is positive, each half of the manifold is the union of iterates of a segment

joining a point to its image. Moreover, the manifold lies close to its tangent near the fixed point. Thus an initial approximation of the manifold can be obtained by iterating points that lie in a small fundamental domain of the linearized map of the fixed point. However, this procedure does not always give a well resolved approximation to the manifold because the points may separate from one another as they iterate away from the fixed point. Algorithms that avoid this difficulty have been implemented [104]. Yorke et al. [114] have used a divide and conquer algorithm to compute one dimensional stable and unstable manifolds of saddles in two dimensional maps. These *straddle* algorithms locate a stable manifold by finding segments whose endpoints iterate towards the saddle point and then proceed in opposite directions along the unstable manifold of the saddle. Continuity implies that a point of the segment lies in the stable manifold. Iteratively following the midpoint of the segment and selecting the half that straddles the stable manifold, the intersection of the stable manifold with the segment can be located precisely. The method is inherently very robust, but it does not emphasize computational efficiency. Recently, Osinga and Krauskopf [104] have described a different procedure to compute one dimensional stable and unstable manifolds.

Some research has been done on the global computation of two dimensional stable and unstable manifolds of equilibria for flows. Several different strategies have been used with reasonable success on such problems, all tested with the stable manifold of the origin in the Lorenz system. Johnson et al. [97] rescaled the vector field so that it had constant length. This approach makes trajectories advance at uniform speed, but their direction continues to follow the strong unstable manifold. Guckenheimer [85] experimented with computation of the geodesic rays in the induced metric of the stable manifold. This procedure appeared to work well, but developed numerical instabilities far from the equilibrium. Osinga and Osinga and Krauskopf [125, 104] have developed methods based upon the *graph transform*. The graph transform is an operator that is used to prove the stable manifold theorem, and Osinga in her thesis implements methods that follow closely the proof. As the Lorenz system stable manifold grows, it acquires complex folds and twists [126]. Tracking the manifold through these folds and twists has been difficult. The graph transform methods are based upon a decomposition of the phase space into a product of linear stable and unstable manifolds near the equilibrium, but the manifold does not remain transverse to the unstable manifold of the equilibrium. Therefore, an adaptive set of coordinate systems is required to track the manifold as it turns. In the methods of Guckenheimer and Johnson, the manifold is computed as a set of curves that bound a growing disk in the manifold. These curves grow in length quickly enough that an interpolation procedure that places new points on the curves as they grow is required to resolve the stable manifold adequately. In places where the manifold develops sharp folds, it becomes difficult to perform this interpolation accurately.

A complementary method for computing the stable and unstable manifolds of low dimensional systems similar to the straddle algorithms of Yorke et al. was implemented by Dellnitz and Hohmann [37]. To compute a compact portion of the manifold, a region in phase space is partitioned and each partition element is marked as to whether it

might intersect the desired manifold. Starting with a coarse partition, many partition elements can usually be marked as not containing an intersection point. These are discarded, and the remaining elements of the partition are refined and then tested to see whether they intersect the manifold. The number of rectangles in successive refinements that must be tested depends on the dimension of the manifold being computed rather than on the dimension of the phase space, so the methods appear feasible for two dimensional manifolds of rather large systems.

Doedel [51] has suggested yet another procedure for computing stable and unstable manifolds based upon the solution of boundary value problems. The idea advanced by Doedel is to formulate an iterative procedure in which each step is the solution of a two point boundary value problem. If W is the invariant manifold and $U \subset W$ is a neighborhood of the equilibrium point that has been determined, then one wants to compute a larger neighborhood of the equilibrium in W . The boundary value solver end point conditions for one end of the interval will be chosen so that the end point of the desired trajectory lies on the boundary of U . If W has dimension d , then these boundary conditions have dimension $d - 1$. The other end point is required to lie on a specified manifold V transverse to W . If V has complementary dimension to W , one more boundary condition is needed. This can be obtained either by enlarging the dimension of V or by allowing the transit time from one end point to the other to vary. The latter strategy is similar to that used by a boundary value solver to obtain the period of a periodic orbit by fixing its length in time while rescaling the vector field with a free parameter.

Robust implementations of algorithms to compute two dimensional stable and unstable manifolds of equilibria have not yet been achieved. The work described above reveals some of the obstacles that have been discovered. These obstacles appear surmountable. Recent improvements in computers should make methods feasible that previously required too much floating point computation or memory use. Better adaptive methods to discretize the intricate geometry of two dimensional stable and unstable manifolds are needed before we will have general purpose codes that reliably compute two dimensional stable and unstable manifolds of equilibria.

3.4 Chaotic Invariant Sets

Chaotic invariant sets have been the focus of a large amount of dynamical systems research. Chaos is a term that has come to mean any type of asymptotic dynamics more irregular than quasiperiodicity. Numerous papers have made the claim that chaos occurs in a particular system, but most of these claims are based only upon visual observation of numerical trajectories. Infrequently, the claims are substantiated with arguments demonstrating that the system has a property that implies the existence of chaos. The strongest criterion for the existence of chaos is the existence of horseshoes in discrete systems or solenoids in flows [140]. These are invariant sets which are topologically equivalent to *subshifts of finite type* in the case of discrete time and their suspensions in the case of flows. All of these objects have been extensively studied from

measure theoretic and statistical viewpoints. They carry invariant measures which are ergodic and have positive entropy and Liapunov exponents [59, 158].

Horseshoes and solenoids vary continuously with C^1 perturbations of a map or flow, lending credence to numerical observations of chaotic structure. The Smale-Birkhoff Homoclinic Theorem [140] gives a necessary and sufficient criterion for the existence of horseshoes for diffeomorphisms, namely that there are transversal intersections of stable and unstable manifolds of a periodic orbit. Application of this theorem to return maps of a flow gives the same result there. When stable and unstable manifolds of periodic orbits can be calculated, this result gives a procedure for verifying the existence of chaotic dynamics in a system. The Melnikov method [78, 132] gives criteria for perturbations of nontransversal homoclinic orbits of periodic orbits to become transversal as a system is deformed. Note, however, that many examples have intersections of stable and unstable manifolds in which the angles of intersection are small, making numerical verification of chaotic dynamics difficult [78]. This is especially true in Hamiltonian systems [123] where Melnikov theory applied to resonant layers of nearly integrable systems fails. Asymptotic analysis of these systems reveals that the angles between stable and unstable manifolds in these layers is “beyond all orders” of the perturbation theory [135]. Simo and his collaborators have investigated carefully several Hamiltonian systems arising in celestial mechanics, including the restricted three body problem [139]. They have made significant strides in demonstrating the existence of very small transversal intersections between stable and unstable manifolds of periodic orbits.

Although numerical evidence is often used to substantiate claims of chaotic behavior, this evidence can be unreliable. One step numerical integration algorithms with fixed time step h define maps that approximate the time h maps of flows. There is a notable qualitative difference between these objects, namely that the trajectories of flows are one dimensional curves while the trajectories of the numerical integrators are sequences of points. Homoclinic orbits of an equilibrium point for a flow cannot be transverse because the stable and unstable manifolds have complementary dimension and any intersection has dimension at least one. The numerical method will have a fixed point near the equilibrium with stable and unstable manifolds of the same dimensions as those of the flow. However, their intersection can be zero dimensional since the trajectories of the numerical method are sequences of points rather than curves. Indeed, the Kupka-Smale Theorem states that for a generic set of maps, homoclinic intersections of periodic points will be transverse [140]. Thus numerical integration can be expected to introduce chaotic behavior to simulations of dynamical systems that cannot have chaotic behavior. The canard example in the previous section displays this property in a slightly different setting. The scale on which such chaos occurs is frequently small, but claims for chaos in a dynamical system based upon observations of a numerical simulation should be bolstered by additional analysis.

Conversely, chaotic dynamics is sometimes difficult to observe in simulations of systems that are indeed believed to be chaotic. Guckenheimer, Kim and MacKay [103] studied an example of this kind in a family of diffeomorphisms of the two dimensional

torus. Investigating resonances in these maps, they discovered the presence of codimension two Takens-Bogdanov bifurcations. Generic two parameter families of maps that undergo Takens-Bogdanov bifurcation have nearby parameters at which a saddle has transversal intersections of its stable and unstable manifolds. In the example investigated by Guckenheimer et al., the region in which this behavior was found was very, very small – a strip of width less than 10^{-10} in a problem for which the parameter space is naturally the unit square. Moreover, the angle between the manifolds became large only in very small neighborhoods of the periodic orbit. Without a systematic search, the chaotic behavior in this family is difficult to find. Similar phenomena occur in the analysis of unfoldings of bifurcations of flows. Two parameter families of flows near codimension two bifurcations of equilibria with a zero eigenvalue and a pair of pure-imaginary eigenvalues have chaotic dynamics in a persistent manner. However, truncated normal forms of these bifurcations do not have chaotic dynamics and once again the angles of transverse intersections of stable and unstable manifolds are initially very small. Thus the failure to detect chaos in numerical simulations does not always mean that it is not present. Compare the study of toral maps by Yorke et al. [71] with the torus maps described above as an example where chaos is almost certainly present but hard to stumble across.

The existence of chaotic attractors has been a subject of intense theoretical investigation. Structurally stable chaotic attractors have uniform hyperbolic structures [140], but it is apparent that many examples which appear to have chaotic attractors cannot have uniform hyperbolic structures on these attractors. The most studied discrete system of this kind is the Henon attractor [89]. Beginning with the theory of iterations of one dimensional mappings [36], an understanding of the properties of chaotic attractors that do not have uniform hyperbolic structures has begun to emerge. Benedicks and Carleson [11] proved that there are families in which non-uniformly hyperbolic attractors occur on parameter sets of positive measure. Their theory and its extensions [12] lend credibility to the belief that chaotic dynamics observed in numerical simulations does indicate that the underlying system has a chaotic attractor.

One of the principal theoretical tools for investigating uniformly hyperbolic invariant sets has been the concept of *Markov partitions* [20]. These partitions lead immediately to representations of these invariant sets as images of subshifts of finite type by maps for which most points have a single preimage. From a statistical perspective, the invariant sets behave like subshifts of finite type. Anosov diffeomorphisms are defined to be diffeomorphisms with dense trajectories and uniform hyperbolic structures on a compact manifold. For two dimensional Anosov diffeomorphisms, the elements of Markov partitions are rectangles whose boundaries are smooth segments of stable and unstable manifolds [3]. In odd dimensions, the boundaries of Markov partitions of Anosov diffeomorphisms are always fractal [21, 26]. Algorithms to compute these partitions have only recently been studied [100, 137]. The work thus far has been restricted to linear Anosov diffeomorphism of the torus and is heavily dependent on algebraic constructions.

There are several phenomena that occur in other examples of chaotic attractors that highlight the bewildering complexity of dynamical systems. One such phenomenon is partial hyperbolicity. Higher dimensional attractors may exhibit *partial hyperbolicity* in which the dimensions of the unstable manifolds of points are always positive, but vary from point to point. Abraham and Smale [2] described an early example of this phenomenon. More recently, Pugh and Shub [127] and others have devoted renewed attention to the analysis of partial hyperbolicity. Numerical investigations of partial hyperbolicity have hardly begun. A second complex phenomenon is *riddled* basins of attraction [4] in which two or more invariant sets have positive measure “domains” of attraction that are densely intertwined. Every open set that contains points tending to one of these invariant sets also contains points tending to another invariant set. In these circumstances, there appears to be an inherent unpredictability about the limit behavior of initial conditions in large regions of phase space.

3.5 Statistical Analysis of Time Series

The numerical analysis of chaotic dynamics has dealt with the statistical properties of invariant sets as well as with algorithms for locating the sets and describing their basins of attraction. Most of the statistical methods are based upon ergodic theory [99] and formulated in terms of invariant measures. From this perspective, the analog of topological transitivity for invariant sets is ergodicity of invariant measures. Hyperbolic invariant sets support many ergodic invariant measures, including the measure theoretic limits $\lim \frac{1}{N} \sum_{i=0}^{N-1} \delta(F^i(x))$ of atomic measures along trajectories tending to the invariant set. These limit measures are frequently called the *time averages* of the trajectories. They do not always exist, even for almost all initial conditions, as has been demonstrated for one dimensional mappings [79]. For attractors, special emphasis has been given to *natural* measures, defined as the limits attained from sets of initial conditions having positive Lebesgue measure. For hyperbolic attractors, these natural measures are the *SBR* (Sinai-Bowen-Ruelle) measures characterized by a variational principle [58]. The convergence of trajectories to measure theoretic limits has been investigated for various examples [14, 13].

Three important statistics of ergodic attractors are their entropy, Liapunov exponents and dimension [156, 157]. Computation of entropy has received relatively little attention compared to computation of Liapunov exponents and dimension. Most algorithms to compute these quantities use data from trajectories, and the methods have been applied to observational data as well as simulations. Nonlinear time series analysis based on these methods provides tools that help assess whether a system might be modeled effectively by one with a low dimensional chaotic attractor. Reconstruction of attractors and construction of models from a scalar time series is a topic that has been extensively studied. The theoretical basis for methods of recovering attractors from one dimensional data was studied by Takens [145] who formulated adaptations of the Whitney Embedding Theorem [92]. This theorem states that generic mappings of an n dimensional manifold into a manifold of dimension $2n + 1$ are embeddings.

Extensions of the theorem have been used to justify the view that the method of *time delays* can be used to embed a chaotic attractor of dimension d into R^k when $k > 2d$. The method begins with a scalar time series of observations y_i that are assumed to be values of the function y at points $x(i\Delta)$ sampled along a trajectory of the attractor. Vectors of the form $(y_i, y_{i+l}, \dots, y_{i+lk})$ are used as observations of the map $E_k(x) = (y(x), y(x(l\Delta)), \dots, y(x(lk\Delta)))$. Takens [145] demonstrates that for a generic observable y , attractor of dimension d and $k > 2d$, E_k is a 1 – 1 map of the attractor into R^k . Procedures for choosing l and k to obtain reliable estimates of the dimension of an observed attractor have been extensively investigated. The accuracy of the methods tends to degrade rapidly with the dimension of the attractor [74].

Liapunov exponents measure the exponential rates of growth of solutions of the variational equations of a vector field. Oseledec [124] proves their existence as measurable functions with respect to any invariant measure. Consequently, the Liapunov exponents of an ergodic measure are invariants of the measure. If x_t is a trajectory of the vector field $\dot{x} = f(x)$ in Euclidean space R^n , then the variational equations of f along x_t are $\dot{\xi} = Df_{x_t}(\xi)$, a nonautonomous system of linear differential equations. Its fundamental solution $\Xi(t)$ is the matrix solution with initial condition $\Xi(0) = I$, the identity matrix. Denoting $\sigma_i(t)$ the singular values of $\Xi(t)$ in decreasing order, the Liapunov exponents of x_t are defined to be $\liminf(\frac{1}{t} \log(\sigma_i(t)))$. Positive Liapunov exponents indicate that there are nearby trajectories that diverge from x_t at an exponential rate. Computation of the largest Liapunov exponent is straightforward: $\sigma_1(t)$ is comparable to $\|\Xi(t)\|$. Determining smaller Liapunov exponents is more difficult because, when the largest Liapunov exponent is simple, $\Xi(t)/\|\Xi(t)\|$ tends to a rank one matrix and round-off errors interfere with the calculation of the smaller Liapunov exponents. This problem has been addressed by reorthogonalization of the solutions of the variational equation. The time interval $[0, T]$ is subdivided into k segments of length δ_j , and $\Xi(t)$ is written as a product $\Xi_k \cdots \Xi_1$ of the fundamental solutions for each of these segments. A series of QR factorizations is then calculated so that $\Xi_j \cdots \Xi_1 = Q_j R_j$. At step j of this iteration, the QR factorization of the matrix $\Xi_{j+1} Q_j$ is needed. The matrices R_j are products of right triangular matrices of moderate size, so their singular values are expected to be more reliable estimates of σ_i than those obtained from a singular value decomposition of the matrix obtained for $\Xi(t)$ by numerical integration. The continuous methods of Dieci et al. [46] provide an alternate approach to this decomposition.

There are several distinct definitions of the dimension of an attractor. The two that have been most used the most in analysis of numerical and observational data are the *pointwise* dimension and the *correlation* dimension [59, 58]. Both of these concepts implicitly rely upon an invariant measure of the attractor. The starting part for their computation is a time series of a trajectory $x(t)$ sampled at N discrete times t_i . Appealing to the ergodic theorem, one assumes that the discrete measure $\frac{1}{N} \sum_{i=0}^{N-1} \delta(x(t_i))$ approximates an invariant measure μ of the attractor. For any measurable set U , the proportion of the points $x(t_i)$ that lie in U is then approximately $\mu(U)$. The calculation of the pointwise dimension of μ and of the correlation dimension of the attractor use

interpoint distances $d_{i,j} = d(x(t_i), x(t_j))$ with respect to a metric d on the phase space. The pointwise dimension gives the rate at which the volumes of balls shrink as their radius tends to zero. Denoting $B_x(r)$ the ball of radius r centered at x , the point-wise dimension of the measure μ is $\lim(\log(\mu(B_x(r))/\log(r))$ as $r \rightarrow 0$ for μ -almost all x in the attractor [156]. To estimate the pointwise dimension, x is chosen to be one of the points $x(t_i)$ in the time series and the sequence $d_{i,j}, j \neq i$, is sorted to produce an increasing sequence r_s . A proportion $s/(N-1)$ of the points $x(t_s)$ lie in the (closed) ball $B_x(r_s)$, so $\mu(B_x(r_s))$ is estimated to be $s/(N-1)$. Extrapolating the slope of $\log(s)$ vs $\log(r_s)$ as $r_s \rightarrow 0$ yields an estimate of the pointwise dimension. There are statistical fluctuations in this estimate that depend on N and the choice of $x = x_i$ [74]. Noise in data affects these calculations, so judgment must be exercised about the range of scales above which deterministic dynamics dominates the location of the observed points and below which noise dominates. If the points on a trajectory are regarded as a random sample of points drawn from the measure μ , variance due to sampling can be reduced by averaging the results for several choices of initial points x_i . Estimating the volume of small balls with feasible amounts of data is problematic for attractors of large dimension α , since the volume decreases like r^α . The correlation dimension is computed with a similar calculation to the pointwise dimension, but instead of sorting the sequence of $N-1$ numbers $d_{i,j}, j \neq i$, all of the $N(N-1)/2$ interpoint distances $d_{i,j}, j < i$ are used.

“Nonlinear” methods for the analysis of time series data have been extensively investigated since the early 1980’s [149]. This research is an inverse problem to the numerical analysis of dynamical systems models, in that it seeks to identify models that fit data. The research began with the observation that linear time series analysis methods did not readily distinguish characteristics of data produced from low dimensional attractors from data produced by systems with large random fluctuations or from systems with high dimensional attractors. The Ruelle-Takens theory of transition to turbulence [130] motivated this research, prompting careful scrutiny of time series data from fluid systems as flows evolved from steady states to turbulent flow [76, 22]. A rich set of methods has been developed using many of the ideas described above, as well as others such as *multi-fractal* analysis [6]. There have been attempts to reconstruct dynamical models directly from data [1]. The mathematical foundations for most of this work is poor compared to the remainder of the material reviewed in this survey. Nonetheless, from a practical standpoint, this area of research has great potential to enhance industrial design and scientific study of systems that can be adequately represented by dynamical systems with low dimensional attractors. The methods are less appropriate for systems that have high dimensional attractors because the amount of data required to reconstruct attractors grows very rapidly with the dimension of the attractors.

4 Bifurcations

Bifurcation theory is the study of how phase portraits of families of dynamical systems change qualitatively as parameters of the family vary. It is a subject filled with complex detail. Singularity theory [67] is an analog to bifurcation theory, providing a framework that has been partially transplanted to the setting of dynamical systems. These efforts have produced a wealth of valuable information, but some of the mathematical completeness and elegance of singularity theory does not carry over to bifurcation theory. The intricacy of dynamical phenomena act as a barrier to the formulation of a theory that classifies all bifurcations that occur in generic families of dynamical systems. Nonetheless, the mathematical concepts adapted from differential topology and singularity theory provide the foundations for successful algorithms. The focus here is upon describing those concepts that are used in numerical methods. Less attention is devoted to results concerning structural stability or genericity.

4.1 Bifurcation Theory

Let $f : R^n \times R^k \rightarrow R^n$ be a k -parameter family of vector fields on R^n . *Equilibrium points* (x, λ) of f are the solutions of $f(x, \lambda) = 0$. The goal of local bifurcation theory is to analyze the set of equilibrium points and their stability, taking into account the dependence upon the parameters [68]. We discuss the location of equilibrium points first and then consider their stability. Near equilibrium points (x, λ) where $D_x f$ is regular or, equivalently, has full rank n , the Implicit Function Theorem states that the solutions of the equilibrium equations form a k -dimensional submanifold of $R^n \times R^k$ that can be parametrized as the graph of a function $x_e : R^k \rightarrow R^n$ from the parameter space to the phase space. Continuation methods implement the computation of x_e . By changing coordinate systems to mix parameters and phase space variables, equilibrium point manifolds that are not graphs from phase space to parameter space can be computed.

Local bifurcations include all points where $D_x f$ is singular. To use equation solvers that rely upon the regularity of the system being solved, we require reformulation of the problem at bifurcation points. A fundamental example, saddle-node bifurcation, introduces the methods used to do so. Saddle-node bifurcation occurs at equilibrium points (x, λ) where $D_x f$ has a simple eigenvalue zero. Thus $(n+1)$ *defining equations* in $R^n \times R^k$ for saddle-node bifurcations are given by zeros of the map

$$F(x, \lambda) = \begin{pmatrix} f(x, \lambda) \\ \det(D_x f)(x, \lambda) \end{pmatrix}$$

If the system of equations $F = 0$ is regular, then Newton's method can be used to locate the points of saddle-node bifurcation. Consider the case in which there is a single parameter: $k = 1$. Assume that (x_0, λ_0) is an equilibrium point at which the defining equations are satisfied. The derivative of DF is then a square $(n+1) \times (n+1)$

matrix with block structure

$$DF = \begin{pmatrix} D_x f & D_\lambda f \\ D_x \det(D_x f) & D_\lambda \det(D_x f) \end{pmatrix}$$

and $D_x f$ singular. In order for this matrix to have full rank, $D_x f$ must have rank at least $n - 1$ since the addition of a single row or column to a matrix increases its rank by at most 1. Here, DF can be obtained from $D_x f$ by the successive addition of one column and one row, so the difference between the ranks of DF and $D_x f$ is at most 2. If $D_x f$ has rank $(n - 1)$, then it has unique left and right eigenvectors $w^T, v \in R^n$ up to scalar multiples. The regularity of DF implies that the products $(w^T 0)DF$ and $DF \begin{pmatrix} v \\ 0 \end{pmatrix}$ are non-zero, yielding that $w D_\lambda f \neq 0$ and $D_x(\det(D_x f))v \neq 0$. The second of these equations is satisfied if $w D_{xx} f(v, v) \neq 0$. The inequalities $w D_\lambda f \neq 0$ and $w D_{xx} f(v, v) \neq 0$ are *nondegeneracy* conditions for saddle-node bifurcation. Together with the assumption that $D_x f$ has rank $(n - 1)$, they give sufficient conditions that the defining equations $F = 0$ for the saddle-node bifurcation are regular.

Regularity of the defining equations for saddle-node bifurcation are not quite enough to characterize the dynamics of a family in the neighborhood of the bifurcation point. However, if the nondegeneracy conditions are strengthened to the requirement that the only eigenvalue on the imaginary axis at the bifurcation point is a simple zero eigenvalue, then the local dynamics near the bifurcation are determined up to topological equivalence and perturbations of the family will have topologically equivalent phase portraits in the neighborhood of the bifurcation. The following theorem summarizes this discussion.

Theorem 2 *Let $\dot{x} = f(x, \lambda)$ be a smooth n -dimensional vector field depending upon a scalar parameter λ . Let (x_0, λ_0) be a solution of the system of equations*

$$\begin{aligned} f(x, \lambda) &= 0 \\ \det(D_x f)(x, \lambda) &= 0 \end{aligned}$$

This system of equations is regular at (x_0, λ_0) if

1. $D_x f$ has rank $n - 1$. Denote the left and right zero eigenvectors of zero by v and w .
2. $w D_{xx} f(v, v) \neq 0$
3. $w D_\lambda f \neq 0$

If properties (1)-(3) are satisfied, the curve γ of equilibrium points $f(x, \lambda) = 0$ is smooth. Furthermore, if zero is a simple eigenvalue of $D_x f$ and $D_x f$ has no pure imaginary eigenvalues, then there is a neighborhood U of (x_0, λ_0) such that all trajectories that remain in U for all time are equilibrium points on γ .

Singularity theory [7, 67] provides a set of tools for analyzing the variation of equilibrium points with respect to parameters in generic families of vector fields. Before tackling the general theory, consider one more example, the family of scalar vector fields $\dot{x} = \lambda_1 + \lambda_2 x + x^3$. For fixed λ_2 and varying λ_1 , this family fails to satisfy the condition $wD_{xx}f(v, v) \neq 0$ when $\lambda_2 = 0$. However, as a two parameter family, the system of equations

$$\begin{aligned} f(x, \lambda) &= 0 \\ D_x f(x, \lambda) &= 0 \\ D_{xx} f(x, \lambda) &= 0 \end{aligned}$$

is regular. The curve on which saddle-node bifurcation occurs in this family is obtained by eliminating x from the pair of equations $\lambda_1 + \lambda_2 x + x^3 = 0$ and $Df(x) = \lambda_2 + 3x^2 = 0$. Parametrically, the curve is given by $\lambda_2 = -3x^2$ and $\lambda_1 = 2x^3$. This implies that $(\frac{\lambda_1}{2})^2 + (\frac{\lambda_2}{3})^3 = 0$. Solutions of this equation form a cusp in the (λ_1, λ_2) parameter plane. It is not smooth at the origin. The example illustrates that the locus of local bifurcation in a generic multi-parameter family may not be a smooth manifold. Nonetheless, singularity theory gives a set of geometric tools that can be used to formulate regular systems of defining equations for local bifurcations. The theory is typically applied at the level of *germs* [67], but the description here avoids this language.

Jets are objects that give coordinate-free expressions for the Taylor series of smooth maps between manifolds. The r -jet extension of a map f associates to each point of the domain of f the Taylor expansion of degree r , viewed as an object in a suitable jet space. *Thom's transversality theorem* [92] states that if P is a submanifold of a jet space, then any smooth map can be perturbed so that its jet extension is transverse to P . In the example of the cusp, the two-jet extension of the family of maps $g(x, \lambda_1, \lambda_2) = \lambda_1 + \lambda_2 x + x^3$ is given by $J^2 g(x, \lambda_1, \lambda_2) = (g, g', g'') = (\lambda_1 + \lambda_2 x + x^3, \lambda_2 + 3x^2, 6x)$ which vanishes at the origin. The Jacobian of $J^2 g$ at the origin is the non-singular matrix

$$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 6 & 0 & 0 \end{pmatrix}$$

Therefore, the cusp gives a family of maps whose two jet extension is transverse to the zero dimensional manifold consisting of the origin. Local bifurcations determined by smooth submanifolds of the jet spaces have regular systems of defining equations in those jet spaces. The transversality theorem implies that solutions of these defining equations yield smooth submanifolds of the product of parameter and phase spaces in generic families of vector fields. The *Thom-Boardman stratification* in singularity theory illustrates concretely how these procedures work.

The Thom-Boardman decomposition of a map $g : R^m \rightarrow R^n$ is constructed as follows. Partition R^m into the sets Σ^i on which Dg has rank $\min(m, n) - i$ or corank i . For generic maps g , the sets Σ^i are submanifolds. The defining equations for Σ^i can be expressed (locally) in the space of 1-jets in terms of minors of Dg . Partition each of the Σ^i

by restricting g to Σ^i and repeating the construction. This produces sets $\Sigma^{i,j}$ on which $g|_{\Sigma^i}$ has rank j . For non-increasing sequences of integers $(i_1, i_2, \dots, i_k)$, Thom defined $\Sigma^{i_1, i_2, \dots, i_k}$ inductively as the set on which the derivative of g restricted to $\Sigma^{i_1, i_2, \dots, i_{k-1}}$ has corank i_k . Boardman proved that, for generic maps g , these sets are submanifolds of R^m . The saddle-nodes and cusps described above correspond to the singularities Σ^1 and $\Sigma^{1,1}$. Interest in the Thom-Boardman stratification was motivated by its relationship to the stability of mappings. The groups of diffeomorphisms of R^m and R^n act on $C^\infty(R^m, R^n)$ by composition on the left and right: $(h, k) \in \text{Diff}^\infty(R^m) \times \text{Diff}^\infty(R^n)$ send g to $h g k$. If g is an interior point of its orbit with respect to this action, then it is stable. The action clearly preserves the Thom-Boardman stratification, so transversality with respect to this stratification is necessary for stability. In a seminal series of papers [117, 119, 118, 120, 121, 116], Mather formulated necessary and sufficient conditions for stability of a mapping. In some cases, transversality with respect to the Thom-Boardman stratification is sufficient, but in other cases it is not.

Local bifurcation theory seeks stratifications of the jet spaces of families of dynamical systems that are analogous to the Thom-Boardman stratification. These stratifications are expected to give necessary conditions for the structural stability of a family, but they will give sufficient conditions in only a limited number of cases. The definition of a stratification used here is naive: a stratification of a closed set V is a sequence of closed subsets $V = V_l \supset V_{l-1} \supset \dots \supset V_0 \supset V_{-1} = \emptyset$ such that each difference $S_i = V_i - V_{i-1}$ is a smooth manifold of dimension i , called a stratum, or empty. The *codimension* of the stratum S_i is $k - i$. Locally, there are regular systems of $k - i$ *defining equations* that define S_i as a subset of R^k . These are *defining equations* for the bifurcations in S_i . Ideally, S_i has a finite number of components, each consisting of vector fields with similar properties near their bifurcation points. A $k - i$ parameter family that is transverse to S_i is an *unfolding*. A particular choice of a point in S_i and an unfolding is a *normal form*. In the “best” circumstances, the normal forms are structurally stable $k - i$ parameter families. Even when this is true, it can be difficult to prove and each case requires a separate analysis.

The analyses of local bifurcations has tended to follow a common pattern. The first step is to identify submanifolds of the space of vector fields that fit into a stratification. These submanifolds should be preserved by topological equivalence or other equivalence relations that are used to describe when vector fields are qualitatively similar to one another. Once the submanifolds are identified, the next step is to choose normal forms for each submanifold. The choice of normal form is usually based upon polynomial coordinate transformations that simplify the analytic expression of the vector field near the bifurcation. The third task in the analysis of local bifurcations is to study the dynamics of the normal form families, seeking to establish their structural stability. The unfoldings of a bifurcation of codimension j will contain in their parameter spaces submanifolds of bifurcations of codimension smaller than j . Many of these lower codimension bifurcations are global bifurcations, making it awkward to maintain a separation between the theories of local and global bifurcation. The primary distinction from a computational perspective is that the defining equations for local

bifurcations are formulated directly in terms of the Taylor series of the vector field rather than in terms of the flow of the vector field. Part of the bifurcation analysis is to identify geometric properties of how strata of smaller codimension limit on the codimension j bifurcation. When the normal forms do not produce structurally stable families, there are two possible scenarios. The first possibility is that a more refined analysis with normal forms of higher degree and additional nondegeneracy conditions on the normal form leads to a structurally stable family. The second possibility is that normal families defined by finite Taylor expansions never produce structurally stable families. As in some cases of double Hopf bifurcation, there may be an infinite number of families of bifurcations that intersect the neighborhood of a bifurcation of codimension j and no simple decomposition of the bifurcation set as a stratified set is possible.

Takens-Bogdanov bifurcation [144, 19] provides a good illustration of the analysis of a local bifurcation of codimension two. This bifurcation occurs at equilibrium points of a vector field for which zero is an eigenvalue of (algebraic) multiplicity two and no eigenvalues are pure imaginary. The defining equations can be expressed easily in terms of the characteristic polynomial of the Jacobian at an equilibrium. Near a point of Takens-Bogdanov bifurcation, the Taylor expansion of degree two for an unfolding can be transformed to

$$\begin{aligned}\dot{x} &= y \\ \dot{y} &= \lambda_1 + ax^2 + y(\lambda_2 + bx)\end{aligned}$$

in the plane corresponding to the generalized eigenspace of zero. This is a normal form for Takens-Bogdanov bifurcation. The normal form is a structurally stable family and the phase portraits near the bifurcation are determined if neither a or b is zero. In the two dimensional parameter plane there are three bifurcation curves that meet at the Takens-Bogdanov point: a curve of saddle-node bifurcations that passes through the TB point, a curve of Hopf bifurcations that terminates at the TB point and a curve of homoclinic bifurcations that terminates at the TB point. These three curves meet with a quadratic tangency, and in the region of parameters between the Hopf and homoclinic bifurcation curves, the vector field has a periodic orbit. This picture illustrates that global bifurcations can appear in the neighborhood of local bifurcations.

The classification of local bifurcations up to topological equivalence of their unfoldings is hardly complete, even for relatively low codimension. There are examples beginning with codimension two bifurcations of three dimensional vector fields in which chaotic dynamics appear in the unfoldings. These examples have an infinite number of bifurcation curves that terminate at the codimension two point in the parameter space, and the families are never structurally stable. Kuznetsov [106] gives a comprehensive summary of information about codimension one and two local bifurcations. Dumortier, Roussarie and Sotomayor [55, 56] have analyzed codimension three local bifurcations of planar vector fields. Their work is the current frontier in attempts to systematically classify local bifurcations of increasing codimension.

Similar principles to the ones discussed above apply to global bifurcations, but the defining equations are expressed in terms of the flow maps instead of directly in terms of the vector field. Chapter XXX of this volume contains more specific information about global codimension one and two bifurcations, including discussion of bifurcations of homoclinic and heteroclinic orbits and the numerical methods implemented in the HomCont package [27] that is part of the 1997 version of AUTO [50]. If flow maps and their derivatives can be computed accurately with numerical integration, then similar numerical methods can be used to compute bifurcations of periodic orbits. There are aspects of global bifurcations that have no counterparts in the theory of local bifurcation. One example is the breakdown of invariant tori. In generic two parameters of vector fields, invariant two dimensional tori with fixed irrational winding number may be present along curves in the parameter space. These parameter space curves corresponding to invariant tori with an irrational winding number may have endpoints beyond which the invariant torus “breaks down” into a Cantor set or a chaotic invariant set. The singularity theory based methods described above are inadequate to analyze the break down process. Renormalization methods that search for self-similar patterns in these phenomena have been used [60].

4.2 Continuation Methods

Continuation methods solve underdetermined systems of equations

$$F = c; \quad F : R^m \rightarrow R^n$$

with $m > n$. They usually assume F is smooth and regular; i.e., DF has rank n on the level set of c . In these circumstances, the Implicit Function Theorem implies that the level set is a smooth manifold of dimension $m - n$ whose tangent spaces are given by the null spaces of DF . Sard’s Theorem [92] implies that for almost all c (with respect to Lebesgue measure in R^n), the level set is a manifold. Continuation algorithms are best developed when $m = n + 1$ and the solution manifolds are curves. Multiparameter continuation with $m > n + 1$ is not yet in widespread use, but remains an active research area [23, 5]. The topological complexity of higher dimensional level sets has not been fully incorporated into robust algorithms.

Single parameter continuation can be formulated as a combination of numerical integration and root finding. The goal is to compute the level curve $F = c$. On the set of regular points of F , one can define the line field that assigns to x the null space of $DF(x)$. This line field can be represented by vector fields in a variety of ways: for example, as a unit vector field determined by an orientation of the null space or via a parametrization of the level curve in the form $y(x)$ where (x, y) are coordinates on R^{n+1} with $D_y F$ a regular $n \times n$ minor of DF and $D_x y = -(D_y F)^{-1} D_x F$. The integral curves of this vector field are level curves of F . Continuation methods exploit this fact to choose predicted steps along a level curve, but they then utilize root finding methods to refine these steps so that they once again satisfy $F = c$. Without this cycle of prediction and correction, numerically integrated curves will likely drift away from

the level curve on which they start. The use of the initial prediction step (typically an Euler step that gives a tangent approximation to the level curve) helps pick seeds for iterative root finders that are close to the desired solutions. This is important when using a root finder like Newton’s method that is not globally convergent.

As with numerical integration, the choice of step length in a continuation method is important. Large step lengths tend to make the root finding less reliable or slower. Small step lengths take more steps to traverse a level set. Choosing unit vectors to parametrize the level curves leads to *pseudo-arclength* continuation. Fixed step sizes yield points along the curve that are approximately equidistant. If the level curve has tight folds with areas of large curvature, then pseudo-arclength continuation is likely to require very small steps. Therefore, adaptive strategies typically monitor the curvature of the level curve and adjust the step length to control the estimated error from each prediction step. The final choice that needs to be made in implementing a continuation method is the choice of subspace in which to perform the root finding. To obtain a “square” system, the original system of equations is constrained to a hypersurface on which F is regular. Common choices are to fix one coordinate; i.e., use a subspace parallel to a coordinate subspace, or to use the hypersurface that is orthogonal to the continuation step.

Continuation methods have been extremely useful in the study of dynamical systems. Here, we examine here use in computing information about local bifurcations. Consider the system of differential equations

$$\dot{x} = f(x, \lambda)$$

with $x \in R^n$ and $\lambda \in R^k$. Local bifurcations locate parameter values λ at which equilibria of this dynamical system have qualitative changes. The Transversality Theorem [92] implies that the equilibrium set of f is a smooth manifold for generic f . In this case, continuation methods can be used to compute the equilibrium manifold. As we compute the set of equilibria $f = 0$ with continuation, we expect to occasionally see bifurcations along the branch. These occur when $D_x f$ is singular or has eigenvalues along the imaginary axis. Thus the problem of computing local bifurcations consists of a continuation problem together with solving additional equations that explicitly depend upon the derivatives of f .

Bifurcations of high codimension serve as “organizing centers” where multiple types of lower codimension bifurcations meet. The branching patterns of many bifurcations of codimension 2, 3 and 4 have been analyzed by first computing their normal forms and then studying the dynamics exhibited by the normal form families. Single parameter continuation has been used to locate and identify high codimension bifurcations with the following strategy, implemented in CONTENT [107]. In computing branches of equilibria for a generic dynamical system with one active parameter, one expects to meet saddle-node and Hopf bifurcations. These are detected by evaluating a function that changes sign at the bifurcation. When a bifurcation point is located, a new continuation can be started to follow the bifurcation curves. The defining equation is added to the equilibrium equations and a second parameter is made active,

producing $n + 1$ equations in $n + 2$ variables. At selected points along these bifurcation curves, codimension two bifurcations may be encountered and detected by evaluation of suitable functions. When this happens, a new continuation is established with a pair of defining equations for the codimension two bifurcation and three active parameters. This process can bootstrap from codimension j to codimension $j + 1$ bifurcations as long as explicit defining equations for the bifurcations have been formulated and the root finding converges. The package CONTENT implements computations of all local codimension two bifurcations of vector fields and discrete maps (and much more as well).

There are important cases in which we want to study systems whose equilibrium sets are not manifolds. For example, systems that are equivariant with respect to a symmetry group are common in varied applications. Equivariance can force the zero level set of a vector field or family of vector fields to have singularities. This complicates the computation of bifurcations substantially. The analysis of these systems is framed in terms of group theoretic concepts. Dynamical analysis of the normal forms of even moderately complex normal forms of symmetric systems is incomplete [61]. Computation of the equilibria and local bifurcations in these systems can require substantial amounts of algebra [153]. Gatermann [65] and Sanders [133] have made initial steps towards the construction of general software for the computation and analysis of normal forms of symmetric systems.

Continuation methods have been used to track curves of periodic orbits as well as equilibria. AUTO [49] implements continuation methods superimposed on collocation algorithms for periodic orbits. The basic advantages of using a continuation method to compute periodic orbits is that initial conditions close to the desired orbit are used for each point along the continuation path after the first. Thus, convergence of the method to the desired orbit is much more likely than with random or fixed data to start each periodic orbit calculation, and fewer calculations are required at each step along the continuation path. There are circumstances in which the use of a global boundary value solver like that employed in AUTO offers additional advantages when coupled with continuation. First, unstable periodic orbits can be computed. As an iterative method, the algorithm has stable fixed points corresponding to all approximate periodic orbits, not only those that are attracting in the flow. With suitable procedures for choosing dependent and independent parameters in the root finding, curves can be followed around folds in which the periodic orbits do not vary smoothly with the parameters. Second, in problems with multiple time scales, one finds families of stable periodic orbits that cannot be computed readily with numerical integration. The canard example in Section 3.1 shows periodic orbits that can be computed with boundary value solvers but not with numerical integration.

4.3 Numerical Methods for Computing Bifurcations

Bifurcation theory provides a mathematical foundation for algorithms that locate bifurcations in specific families. Implementation of methods based upon singularity theory

encounters three types of numerical issues:

1. Formulation of regular systems of defining equations
2. Accurate evaluation of defining equations that depend upon derivatives of a vector field
3. Numerical condition number of the defining equations, especially for large systems and systems with multiple time scales

These issues can be viewed from both theoretical and practical perspectives. Practically, the most desirable numerical methods are those that give accurate answers for large classes of interesting systems. There are several important choices that enter the construction of software for computing bifurcations, so the potential number of distinct methods is large. Yet, different methods are seldom compared carefully with one another. There have been few attempts to gather suites of test problems in this domain as there has been for numerical integration [95]. Picking parameters in bifurcation algorithms that make methods work well remains an art. Thus, opinions of different methods tend to be very subjective, based upon the experience of users and the skill they develop in adjusting algorithmic parameters when a method initially fails.

Regular systems of defining equations for saddle-node bifurcations are presented above in terms of the determinant of the Jacobian of a vector field. This choice of defining function is natural from a theoretical perspective, but may lead to numerical problems. If the Jacobian has eigenvalues of large magnitude, then these eigenvalues contribute to the condition number of the determinant and may make it difficult to satisfy the defining equations to a desired tolerance. The larger the system, the worse this problem becomes. Thus there are circumstances in which it is desirable to seek alternate defining equations for saddle-node bifurcation that avoid calculation of the determinant. The singular matrices do not form a smooth submanifold of the space $R^{n \times n}$ of $n \times n$ matrices, so there is no regular function whose values measure the distance of a matrix from being singular. However, the corank one matrices are a smooth hypersurface in $R^{n \times n}$. Saddle-node bifurcations occur at matrices of corank one, and this fact can be used in the formulation of defining equations. The following result about bordered matrices is the basis for one method.

Theorem 3 [70] *Let A be an $n \times n$ matrix that has a single eigenvalue zero. For most choices of n vectors B and C and scalar D the $(n+1) \times (n+1)$ block matrix*

$$M = \begin{bmatrix} A & B \\ C^t & D \end{bmatrix}$$

is nonsingular. There are constants $c_1 > 0$ and c_2 and a neighborhood U of A so that if $\bar{A} \in U$ with smallest singular value σ , then the unique solution (u^t, v) of the system of equations

$$\begin{bmatrix} \bar{A} & B \\ C^t & D \end{bmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

satisfies $c_1\sigma < |v| < c_2\sigma$.

Applying the theorem with A the Jacobian of the vector field gives the quantity v as a measure of the distance of the Jacobian from the set of singular matrices. The system of linear equations can be solved using Gaussian elimination with partial pivoting, an algorithm that is efficient and reliable for most systems. Moreover, when v is zero, u is the right zero eigenvector of A , an object needed to compute the normal form of the bifurcation. High dimensional vector fields often have sparse Jacobians. For these, iterative methods can be used to compute the solution of the system of linear equations, avoiding the need to calculate a full factorization of the matrix M . Thus, this method is feasible for discretized systems of partial differential equations for which computation of the determinant of the Jacobian can hardly be done.

A second approach to computing saddle-node bifurcations is to rely upon numerical methods for computing low dimensional invariant subspaces of a matrix. Subspace iteration and Arnoldi methods [142] are effective techniques for identifying invariant subspaces that are associated with the eigenvalues of largest magnitude for a matrix. Inverse iterations can be used in this framework to identify invariant subspaces associated with eigenvalues close to the origin. Cayley transforms [64, 122] extend these methods to compute invariant subspaces for any cluster of eigenvalues on the Riemann sphere. This can be especially useful in finding Hopf bifurcations, but subspaces associated to eigenvalues of large magnitude on the imaginary axis cannot be readily separated from subspaces associated with negative eigenvalues of large magnitude. If appropriate invariant subspaces are computed, then the bifurcation calculations can be reduced to these subspaces. On the remaining small problems, the choice of function that vanishes on singular matrices matters less than it does for large problems.

Deriving explicit defining equations for bifurcations other than saddle-nodes requires additional effort. For example, Hopf bifurcation occurs when the Jacobian at an equilibrium has a pair of pure imaginary eigenvalues. There is no familiar function that vanishes when a matrix has pure imaginary eigenvalues analogous to the determinant for zero eigenvalues. Guckenheimer, Myers and Sturmfels [83] described algebraic procedures that produce single augmenting equations analogous to the determinant and the bordered matrix equation for saddle-node bifurcation in Section 4.1. The algebraic equation can be derived from the characteristic polynomial of the Jacobian. A determinant, the Sylvester resultant of two polynomials constructed from the characteristic polynomial, vanishes if and only if the Jacobian matrix has a pair of eigenvalues whose sum is zero. There are two ways in which a real matrix can have a pair of eigenvalues whose sum is zero: they can be real or they can be pure imaginary. There is an explicit algebraic inequality in the coefficients of the characteristic polynomial that distinguishes these two cases. While the formulas that arise from this analysis are suitable for computations with low dimensional systems, they rapidly become unwieldy as the dimension of a vector field grows. They suffer from all of the problems associated with the use of the determinant as a defining equation for saddle-node bifurcations as well as the additional difficulty that computations of the characteristic polynomial tend to suffer from numerical instability [150].

Tensor products yield a procedure for computing Hopf bifurcations without forming the characteristic polynomial of a matrix. Given $n \times n$ matrices A and B , their tensor product is an $n^2 \times n^2$ matrix $A \otimes B$ whose eigenvalues are the products of the eigenvalues of A and B . Therefore, the eigenvalues of the matrix $C = A \otimes I + I \otimes A$ are sums of pairs of the eigenvalues of A . Moreover, C can be decomposed into a symmetric part that commutes with the involution $u \otimes v \rightarrow v \otimes u$ and a skewsymmetric part that anticommutes with this involution. The skewsymmetric part is an $n(n-1)/2 \times n(n-1)/2$ matrix (called the *biproduct* of A) whose eigenvalues are the sums of distinct eigenvalues of A . Therefore, A has a single pair of eigenvalues whose sum is zero if and only if its biproduct has corank one. Applying the bordered matrix construction described above to the biproduct gives a defining function for A to have a single pair of eigenvalues whose sum is zero. Govaerts, Guckenheimer and Khibnik [69] studied the Jordan decomposition of the biproduct of matrices with multiple pairs of eigenvalues whose sum was zero and used a bordering construction to implement a system of defining equations for double Hopf bifurcation.

The methods described above for computing saddle-node and Hopf bifurcations construct minimal augmentations of the defining equations. There are alternative methods that introduce additional independent variables and utilize larger systems of defining equations. For example, in the case of Hopf bifurcation, many methods solve for the pure imaginary Hopf eigenvalues and eigenvectors associated with these. In addition to the equilibrium equations, one method solves the equations $Df v = \omega w$ and $Df w = -\omega v$ for vectors v and w as well as the eigenvalue $i\omega$ [128]. To make this system of equations regular, additional equations that normalize v and w are required. The complexity of the expressions appearing in these defining equations is reduced compared to that of minimal augmentation methods. This advantage is offset by the expense of having larger systems to solve with root finding and the necessity of finding initial seeds for the auxiliary variables. Guckenheimer and Myers [82] give a list of methods for computing Hopf bifurcations and a comparison between their method and the one of Roose and Hvalacek [128].

The defining equations of local bifurcations include derivatives of f . In the cases of some bifurcations of codimension two and larger, the expressions for these defining equations are very complex and involve higher derivatives of f . Consequently, accuracy and efficient evaluation of the defining equations is important. Automatic differentiation [15] provides methods for the accurate evaluation of the derivatives themselves that avoids truncation errors inherent in finite difference formulas. General expressions for defining equations of some types of bifurcations have been derived only recently, so only a small amount of testing has been done with computation of these bifurcations [82].

The description of high codimension singularities of maps has proceeded farther than the description of high codimension bifurcations of dynamical systems. The thesis of Xiang [154] contains results that surmount a technical difficulty in implementing the computation of Thom-Boardman singularities [18]. The problem is that the singularities are defined by equations on submanifolds of a domain: $\Sigma^{i_1, i_2, \dots, i_k}$ is the set on which

the map restricted to $\Sigma^{i_1, i_2, \dots, i_{k-1}}$ has corank i_k . The corank conditions can be expressed in terms of minors of the derivative of the restricted map, but numerical computations only yield approximations to $\Sigma^{i_1, i_2, \dots, i_{k-1}}$. These approximations do not automatically produce good approximations of tangent spaces and regular systems of defining equations. Xiang [154] altered the construction of defining equations to produce a regular systems of equations for $\Sigma^{i_1, i_2, \dots, i_k}$ defined in neighborhoods of $\Sigma^{i_1, i_2, \dots, i_{k-1}}$. These methods were tested with seven dimensional stable maps containing $\Sigma^{2,1}$ singularities, the smallest example of stable maps in which this difficulty arises.

References

- [1] H. Abarbanel, R. Brown, J. Sidorowich and L. Tsimring, The analysis of observed chaotic data in physical systems. *Rev. Modern Phys.* 65: 1331-1392, 1993.
- [2] R. Abraham and S. Smale, Nongenericity of ω -stability. In *Global Analysis (Proc. Sympos. Pure Math., Vol. XIV: Berkeley, Calif., 1968)*, pp. 5-8, Amer. Math. Soc., Providence, R.I., 1970.
- [3] Roy L. Adler, Symbolic dynamics and Markov partitions. *Bull. Amer. Math. Soc. (N.S.)*, 35:1-56, 1998.
- [4] J. Alexander, J. Yorke, Z. You and I. Kan, Riddled basins. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.* 2: 795-813, 1992.
- [5] E. Allgower and K. Georg, *Numerical Continuation Methods: An Introduction*, Springer Verlag 13, 1990.
- [6] F. Argoul, A. Arneodo and G. Grasseau, Fractal dimensions and $f(\alpha)$ spectrum for strange attractors. *Z. Angew. Math. Mech.* 68: 519-522, 1988.
- [7] V. Arnold, *Singularity Theory, Selected Papers*. London Mathematical Society Lecture Notes 53, 1981.
- [8] D. G. Aronson, M. A. Chory, G. R. Hall, and R. P. McGehee, A discrete dynamical system with subtly wild behavior. In *New approaches to nonlinear problems in dynamics (Proc. Conf., Pacific Grove, Calif., 1979)*, pp. 339-359, SIAM, Philadelphia, Pa., 1980.
- [9] Uri M. Ascher, Robert M. M. Mattheij, and Robert D. Russell, *Numerical solution of boundary value problems for ordinary differential equations*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1995. Corrected reprint of the 1988 original.
- [10] U. M. Ascher and L. R. Petzold, *Computer methods for ordinary differential equations and differential-algebraic equations*, SIAM, Philadelphia, PA, 1998.

- [11] Michael Benedicks and Lennart Carleson, The dynamics of the Hénon map. *Ann. of Math.* , 133: 73-169, 1991.
- [12] Michael Benedicks and Lai Sang Young, Sinai-Bowen-Ruelle measures for certain Hénon maps. *Invent. Math.*, 112: 541-576, 1993.
- [13] G. Benettin, M. Casartelli, L. Galgani, A. Giorgilli, and J.-M. Strelcyn, On the reliability of numerical studies of stochasticity. II. Identification of time averages. *Nuovo Cimento B* , 50: 211-232, 1979.
- [14] Giancarlo Benettin and Luigi Galgani, Ljapunov characteristic exponents and stochasticity. In *Intrinsic stochasticity in plasmas* (Internat. Workshop, Inst. Études Sci. Cargèse, Cargèse, 1979), pp. 93-114. École Polytech., Palaiseau, 1979.
- [15] M. Berz, C. Bischof, G. Corliss and A. Griewank (eds.), *Computational Differentiation*. SIAM, 1996.
- [16] W. Beyn and J. Kleinkauf, The numerical computation of homoclinic orbits for maps. *SIAM J Numer. Anal.* 34: 1207-1236, 1997.
- [17] W. J. Beyn and W. Kless, Numerical expansions of invariant manifolds in large dynamical systems. *Numer. Math.* 80: 1-38.
- [18] J. Boardman, Singularities of differentiable maps. *Inst. Hautes tudes Sci. Publ. Math.* No. 33: 21-57, 1967.
- [19] R. Bogdanov, Versal deformations of a singular point on the plane in the case of zero eigenvalues. *Functional Analysis and its Applications* 9:144-145, 1975.
- [20] Rufus Bowen, *Equilibrium states and the ergodic theory of Anosov diffeomorphisms*. Springer-Verlag, Berlin, 1975. *Lecture Notes in Mathematics*, Vol. 470.
- [21] Rufus Bowen, Markov partitions are not smooth. *Proc. Amer. Math. Soc.*, 71: 130-132, 1978.
- [22] A. Brandstätter, J. Swift, H. Swinney, A. Wolf and J. Farmer, Low dimensional chaos in a hydrodynamic system, *Phys. Rev. Lett.* 51: 1442-1445, 1983.
- [23] M. Brodzik and W. Rheinboldt, The computation of simplicial approximations of implicitly defined two-dimensional manifolds. *Comput. Math. Appl.* 28: 9-21, 1994.
- [24] H. W. Broer, H. M. Osinga, and G. Vegter, Algorithms for computing normally hyperbolic invariant manifolds. *Z. Angew. Math. Phys.*, 48: 480-524, 1997.
- [25] M. Casdagli and S. Eubank (eds.), *Nonlinear Modeling and Forecasting*, Addison-Wesley, 1992.

- [26] E. Cawley, Smooth Markov partitions and toral automorphisms. *Ergodic Theory Dynam. Systems*, 11: 633–651, 1991.
- [27] A. R. Champneys, Yu. A. Kuznetsov, and B. Sandstede, A numerical toolbox for homoclinic bifurcation analysis. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 6: 867-887, 1996.
- [28] S. Chow, X. Lin and K. Palmer, A shadowing lemma with applications to semi-linear parabolic equations. *SIAM J. Math. Anal.*, 20: 547-557, 1989.
- [29] S. Chow and E. Van Vleck, A shadowing lemma approach to global error analysis for initial value ODEs. *SIAM J. Sci. Comput.* 15: 959-976, 1994.
- [30] B. Coombes, H. Kocak and K. Palmer, Rigorous computational shadowing of orbits of ordinary differential equations. *Numer. Math.* 69: 401-421, FEB 1995.
- [31] B. Coombes, H. Kocak and K. Palmer, Long periodic shadowing. *Numer Algorithms* 14: 55-78, 1997.
- [32] George Corliss and Y. F. Chang, Solving ordinary differential equations using Taylor series. *ACM Trans. Math. Software*, 8: 114-144, 1982.
- [33] George F. Corliss, Survey of interval algorithms for ordinary differential equations. *Appl. Math. Comput.*, 31: 112-120, 1989. Numerical ordinary differential equations (Albuquerque, NM, 1986).
- [34] G. Corliss, Theory and numerics of ODE and PDE in M. Ainsworth, J. Levesley, W.A. Light, M. Marletta (Eds.) *Advances in Numerical Analysis*, Vol. IV: Oxford U. Press, 1-76, 1995.
- [35] G. Dahlquist, A special stability problem for linear multistep methods. *BIT*: 27-43, 1963.
- [36] Welington de Melo and Sebastian van Strien, One-dimensional dynamics. Springer-Verlag, Berlin, 1993.
- [37] Michael Dellnitz and Andreas Hohmann, The computation of unstable manifolds using subdivision and continuation. In *Nonlinear dynamical systems and chaos* (Groningen, 1995), pp. 449-459, Birkhäuser, Basel, 1996.
- [38] A. Denjoy, Sur les courbes définies par les équations différentielles à la surface du tore. *J. Math.*, 17(IV): 333-375, 1932.
- [39] L. Dieci, Positive definiteness in the numerical solution of Riccati differential equations. *Numer. Math.* 67: 303-313, 1994.
- [40] L. Dieci and G. Bader, Solutions of the systems associated with the invariant tori approximation 2. Multigrid methods. *Siam J. Sci. Comput.* 15: 1375-1400, 1994.

- [41] L. Dieci and T. Eirola, On smooth decomposition of matrices. *SIAM J. Matrix Anal. Appl.* 20: 800-819, 1999.
- [42] L. Dieci and J. Lorenz, Computation of invariant tori by the method of characteristics. *SIAM J. Numer. Anal.* 32: 1436-1474, 1995.
- [43] L. Dieci and E. Van Vleck, Computation of invariant tori by the method of characteristics. *SIAM J. Numer. Anal.* 32: 1436-1474, 1995.
- [44] L. Dieci and J. Lorenz, Lyapunov-type numbers and torus breakdown: Numerical aspects and a case study. *Numer. Algorithms* 14: 79-102, 1997.
- [45] L. Dieci and J. Lorenz and R. Russell, Numerical calculation of invariant tori. *SIAM J. Sci. Stat. Comput.* 12: 607-647, 1991.
- [46] L. Dieci, R. Russell, and E. VanVleck, On the computation of Lyapunov exponents for continuous dynamical systems. *SIAM J Numer. Anal.* 34: 402-423, 1997.
- [47] Francine Diener and Marc Diener, Chasse au canard. I. Les canards. *Collect. Math.*, 32: 37-74, 1981.
- [48] Marc Diener, The canard unchained or how fast/slow dynamical systems bifurcate. *Math. Intelligencer*, 6: 38-49, 1984.
- [49] E. Doedel, AUTO: a program for the automatic bifurcation analysis of autonomous systems. In *Proceedings of the Tenth Manitoba Conference on Numerical Mathematics and Computing, Vol. I (Winnipeg, Man., 1980)*, volume 30: pp. 265-284, 1981.
- [50] E. Doedel, AUTO, <ftp://ftp.cs.concordia.ca/pub/doedel/auto>.
- [51] E. Doedel, personal communication.
- [52] Eusebius Doedel, Herbert B. Keller, and Jean-Pierre Kernévez, Numerical analysis and control of bifurcation problems. I. Bifurcation in finite dimensions. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 1: 493-520, 1991.
- [53] Eusebius Doedel, Herbert B. Keller, and Jean-Pierre Kernévez, Numerical analysis and control of bifurcation problems. II. Bifurcation in infinite dimensions. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 1: 745-772, 1991.
- [54] G. F. D. Duff, Limit cycles and rotated vector fields. *Ann. Math.* 57: 15-31, 1953.
- [55] F. Dumortier, R. Roussarie, and J. Sotomayor, Generic 3-parameter families of vector fields on the plane, unfolding a singularity with nilpotent linear part. The cusp case of codimension 3. *Ergodic Theory Dynamical Systems*, 7: 375-413, 1987.

- [56] Freddy Dumortier, Robert Roussarie, and Jorge Sotomayor, Bifurcations of cuspidal loops. *Nonlinearity*, 10: 1369-1408, 1997.
- [57] W. Eckhaus, A standard chase on French ducks. *Lect. Notes in Math.* 985, 449-494, 1983.
- [58] J.-P. Eckmann and D. Ruelle, Addendum: "Ergodic theory of chaos and strange attractors". *Rev. Modern Phys.*, 57: 1115, 1985.
- [59] J.-P. Eckmann and D. Ruelle, Ergodic theory of chaos and strange attractors. *Rev. Modern Phys.*, 57: 617-656, 1985.
- [60] Mitchell J. Feigenbaum, Universal behavior in nonlinear systems. *Phys. D*, 7: 16-39, 1983. *Order in chaos* (Los Alamos, N.M., 1982).
- [61] M. Field, Symmetry breaking for compact Lie groups. *Mem. Amer. Math. Soc.* 120 (1996), no. 574: viii+170 pp.
- [62] M. Field, Lectures on bifurcations, dynamics and symmetry. *Pitman Research Notes in Mathematics Series*, 356. Longman, Harlow, 1996. iv+159 pp.
- [63] H. Froehling, J. Crutchfield, J. Farmer, N. Packard and R. Shaw, On determining the dimension of chaotic flows. *Physica 3D*: 605-617, 1981.
- [64] T. J. Garratt, G. Moore, and A. Spence, A generalised Cayley transform for the numerical detection of Hopf bifurcations in large systems. In *Contributions in numerical mathematics*, pp. 177-195, World Sci. Publishing, River Edge, NJ, 1993.
- [65] K. Gatermann and R. Lauterbach, Automatic classification of normal forms. *Nonlinear Anal.* 34: 157-190, 1998.
- [66] C. William Gear, Numerical initial value problems in ordinary differential equations. Prentice-Hall Inc., Englewood Cliffs, N.J., 1971.
- [67] M. Golubitsky and V. Guillemin, Stable mappings and their singularities. Springer-Verlag, New York, 1973. *Graduate Texts in Mathematics*, Vol. 14.
- [68] W. Govaerts, Numerical Methods for Bifurcations of Dynamical Equilibria, SIAM, 1999.
- [69] W. Govaerts, J. Guckenheimer, and A. Khibnik, Defining functions for multiple Hopf bifurcations. *SIAM J. Numer. Anal.*, 34: 1269-1288, 1997.
- [70] W. Govaerts and Pryce, J.D., A singular value inequality for block matrices. *Lin. Alg. Appl.* 125: 141-148, 1989.
- [71] C. Grebogi, E. Ott, and J. Yorke, Attractors on an N -torus: quasiperiodicity versus chaos. *Phys. D*, 15: 354-373, 1985.

- [72] J. M. Greene, R. S. MacKay, F. Vivaldi, and M. J. Feigenbaum, Universal behaviour in families of area-preserving maps. *Phys. D*, 3: 468-486, 1981.
- [73] A. Griewank, Juedes, and J. Utke, ‘ADOL-C: A Package for the Automatic Differentiation of Algorithms Written in C/C++’, Version 1.7, Argonne National Laboratory.
- [74] J. Guckenheimer, Dimension estimates for attractors. In *Fluids and plasmas: geometry and dynamics* (Boulder, Colo., 1983), pp. 357-367. Amer. Math. Soc., Providence, RI, 1984.
- [75] J. Guckenheimer, Phase portraits of planar vector fields: computer proofs. *Experiment. Math.*, 4: 153-165, 1995.
- [76] J. Guckenheimer and G. Buzyna, Dimension measurements for geostrophic turbulence, *Phys. Rev. Lett.* 51: 1438-1441, 1983.
- [77] J. Guckenheimer and W. G. Choe, Computing periodic orbits with high accuracy. *Computer Methods in Applied Mechanics and Engineering*, 170: 331-341, 1999.
- [78] John Guckenheimer and Philip Holmes, *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*. Springer-Verlag, New York, 1983.
- [79] J. Guckenheimer and S. Johnson, Distortion of S-unimodal maps. *Annals of Mathematics*, 132: 71-130, 1990.
- [80] John Guckenheimer and Salvador Malo, Computer-generated proofs of phase portraits for planar systems. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 6: 889-892, 1996.
- [81] J. Guckenheimer and B. Meloon, Computing Periodic Orbits and their Bifurcations with Automatic Differentiation, preprint 1999.
- [82] J. Guckenheimer and M. Myers, Computing Hopf Bifurcations II. *SIAM J. Sci. Comp.*, 17: 1275-1301, 1996.
- [83] John Guckenheimer, Mark Myers, and Bernd Sturmfels, Computing Hopf bifurcations. I. *SIAM J. Numer. Anal.*, 34: 1-21, 1997.
- [84] John Guckenheimer and R. F. Williams, Structural stability of Lorenz attractors. *Inst. Hautes Études Sci. Publ. Math.*, 50: 59-72, 1979.
- [85] John Guckenheimer and Patrick Worfolk, Dynamical systems: some computational problems. In *Bifurcations and periodic orbits of vector fields* (Montreal, PQ, 1992), pp. 241-277. Kluwer Acad. Publ., Dordrecht, 1993.
- [86] E. Hairer, S. P. Nørsett, and G. Wanner, *Solving ordinary differential equations. I. Springer-Verlag, Berlin, 1987. Nonstiff problems.*

- [87] E. Hairer and G. Wanner, Solving ordinary differential equations. II. Springer-Verlag, Berlin, 1991. Stiff and differential-algebraic problems.
- [88] S. M. Hammel, J. A. Yorke and C. Grebogi, Numerical orbits of chaotic processes represent true orbits. *Bull. Amer. Math. Soc. (N.S.)* 19: 465-469, 1988.
- [89] M. Henon, A two dimensional map with a strange attractor. *Comm. Math. Phys.* 50: 69-77, 1976.
- [90] M. Herman, Sur la conjugaison différentiable des difféomorphismes du cercle des rotations. *Inst. Hautes études Sci. Publ. Math. No. 49*: 5-233, 1979.
- [91] Peter Henrici, Discrete variable methods in ordinary differential equations. John Wiley & Sons Inc., New York, 1962.
- [92] Morris W. Hirsch, Differential topology. Springer-Verlag, New York, 1976. Graduate Texts in Mathematics, No. 33.
- [93] M. Hirsch, C. Pugh and M. Shub, Invariant manifolds. *Lecture Notes in Mathematics*, Vol. 583. Springer-Verlag, Berlin-New York, 1977. ii+149 pp.
- [94] John H. Hubbard and Beverly H. West, Differential equations: a dynamical systems approach. Springer-Verlag, New York, 1995. Ordinary differential equations, Corrected reprint of the 1991 edition.
- [95] T. Hull, W. Enright, B. Fellen and A. Sedgwick, Comparing numerical methods for ordinary differential equations. *SIAM J. Numer. Anal.* 9: 603-637, 1972; errata, *ibid.* 11, 681, 1974.
- [96] Y. Ilyashenko and S. Yakovenko, Finitely smooth normal forms of local families of diffeomorphisms and vector fields. (Russian) *Uspekhi Mat. Nauk* 46: 3-39, 1991; translation in *Russian Math. Surveys* 46, 1-43, 1991.
- [97] Mark E. Johnson, Michael S. Jolly, and Ioannis G. Kevrekidis, Two-dimensional invariant manifolds and global bifurcations: some approximation and visualization studies. *Numer. Algorithms*, 14: 125-140, 1997. *Dynamical numerical analysis* (Atlanta, GA, 1995).
- [98] A. Jorba and J. Villanueva, Numerical computation of normal forms around some periodic orbits of the Restricted Three Body Problem. *Physica D* 114: 197-229, 1998.
- [99] Anatole Katok and Boris Hasselblatt, Introduction to the modern theory of dynamical systems. Cambridge University Press, Cambridge, 1995. With a supplementary chapter by Katok and Leonardo Mendoza.
- [100] Richard Kenyon and Anatoly Vershik, Arithmetic construction of sofic partitions of hyperbolic toral automorphisms. *Ergodic Theory Dynam. Systems*, 18: 357-372, 1998.

- [101] I. Kevrekidis, R. Aris, L. Schmidt and S. Pelikan, Numerical computation of invariant circles of maps. *Phys. D* 16, 243-251, 1985.
- [102] Alexander I. Khibnik, Yuri A. Kuznetsov, Victor V. Levitin, and Eugene V. Nikolaev, Continuation techniques and interactive software for bifurcation analysis of ODEs and iterated maps. *Phys. D*, 62: 360-371, 1993. Homoclinic chaos (Brussels, 1991).
- [103] Seung-hwan Kim, R. S. MacKay, and J. Guckenheimer, Resonance regions for families of torus maps. *Nonlinearity*, 2: 391-404, 1989.
- [104] Bernd Krauskopf and Hinke Osinga, Growing 1D and quasi-2D unstable manifolds of maps. *J. Comput. Phys.*, 146: 404-419, 1998.
- [105] Ivan Kupka, Contribution à la théorie des champs génériques. *Contributions to Differential Equations*, 2: 457-484, 1963.
- [106] Y. Kuznetsov, Elements of applied bifurcation theory. Applied Mathematical Sciences, 112. Springer-Verlag, New York, 1995.
- [107] V. Kuznetsov and V. Levitin, CONTENT, <ftp://ftp.cwi.nl/pub/CONTENT>.
- [108] O. Lanford, A computer-assisted proof of the Feigenbaum conjectures. *Bull. Amer. Math. Soc. (N.S.)*, 6: 427-434, 1982.
- [109] O. Lanford, Computer-assisted proofs in analysis. In *Proceedings of the International Congress of Mathematicians, Vol. 1, 2* (Berkeley, Calif., 1986), pp. 1385-1394, Providence, RI, 1987. Amer. Math. Soc.
- [110] O. Lanford, Renormalization group methods for circle mappings. *Statistical mechanics and field theory: mathematical aspects*. (Groningen, 1985), 176-189, *Lecture Notes in Phys.*, 257: Springer, Berlin-New York, 1986.
- [111] A. Laub and P. gahinet, Numerical improvements for solving Riccati equations. *IEEE Trans. Automat. Control* 42: 1303–1308, 1997.
- [112] Rudolf J. Löhner, Computation of guaranteed enclosures for the solutions of ordinary initial and boundary value problems. In *Computational ordinary differential equations* (London, 1989), pp. 425-435. Oxford Univ. Press, New York, 1992.
- [113] E. Lorenz, Deterministic nonperiodic flow. *J. Atmos. Sci.*, 20: 130-141, 1963.
- [114] Z. Lou, E. J. Kostelich, and J. A. Yorke, Erratum: "Calculating stable and unstable manifolds". *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 2: 215, 1992.
- [115] Kyoko Makino and Martin Berz, Remainder differential algebras and their applications. In *Computational differentiation* (Santa Fe, NM, 1996), pp. 63-74. SIAM, Philadelphia, PA, 1996.

- [116] J. N. Mather, Stability of C^∞ mappings. VI: The nice dimensions. pp. 207-253. Lecture Notes in Math., Vol. 192, 1971.
- [117] John N. Mather, Stability of C^∞ mappings. I. The division theorem. Ann. of Math. , 87: 89-104, 1968.
- [118] John N. Mather, Stability of C^∞ mappings. III. Finitely determined mapgerms. Inst. Hautes Études Sci. Publ. Math. No., 35: 279-308, 1968.
- [119] John N. Mather, Stability of C^∞ mappings. II. Infinitesimal stability implies stability. Ann. of Math. , 89: 254-291, 1969.
- [120] John N. Mather, Stability of C^∞ mappings. IV. Classification of stable germs by R -algebras. Inst. Hautes Études Sci. Publ. Math. No., 37: 223-248, 1969.
- [121] John N. Mather, Stability of C^∞ mappings. V. Transversality. Advances in Math., 4: 301-336, 1970.
- [122] K. Meerbergen, A. Spence, and D. Roose, Shift-invert and Cayley transforms for detection of rightmost eigenvalues of nonsymmetric matrices. BIT, 34: 409-423, 1994.
- [123] V. Melnikov, On the stability of the center for time periodic perturbations, Trans. Moscow Math. Soc. 12: 1-57, 1963.
- [124] A multiplicative ergodic theorem. Liapunov charactersitic numbers for dynamical systems. Trans. Moscow Math. Soc. 19:197-231, 1968.
- [125] H. Osinga, Computing Invariant manifolds, Thesis, Groningen, 1996.
- [126] H. Osinga, <http://www.cds.caltech.edu/hinke/lorenz/>.
- [127] C. Pugh and M. Shub, Stably ergodic dynamical systems and partial hyperbolicity. J. Complexity 13: 125-179, 1997.
- [128] Roose, D. and V. Hlavacek, A Direct Method for the Computation of Hopf Bifurcation Points. SIAM Journal of Applied Mathematics, Vol. 45: pps. 879-894, 1985.
- [129] David Ruelle, Ergodic theory of differentiable dynamical systems. Inst. Hautes Études Sci. Publ. Math., 50: 27-58, 1979.
- [130] D. Ruelle and F. Takens, On the nature of turbulence. Comm. Math. Phys. 20: 167-192, 23: 343-344, 1971.
- [131] Lorenz attractors through Silnikov-type bifurcation. I. Ergodic Theory Dynamical Systems 10: 793-821, 1990.
- [132] J. Sanders, Melnikov's method and averaging. Celestial Mech. 28: 171-181, 1982.

- [133] J. Sanders, Versal normal form computations and representation theory. Computer algebra and differential equations (1992), 185-210, London Math. Soc. Lecture Note Ser., 193: Cambridge Univ. Press, Cambridge, 1994.
- [134] J. M. Sanz-Serna, Solving numerically Hamiltonian systems. In Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994), pp. 1468-1472, Basel, 1995. Birkhäuser.
- [135] J. Scheurle, J. Marsden and P. Holmes, Exponentially small estimates for separatrix splittings. Asymptotics beyond all orders. (La Jolla, CA, 1991), 187-195, NATO Adv. Sci. Inst. Ser. B Phys., 284, Plenum, New York, 1991.
- [136] M. Shub, Dynamical systems, filtrations and entropy. Bull. Amer. Math. Soc. 80: 27-41, 1974.
- [137] Nikita Sidorov and Anatoly Vershik, Bijective Arithmetic Codings of Hyperbolic Automorphisms of the 2-Torus, and Binary Quadratic Forms, preprint 1998.
- [138] Simó C.: Analytical and numerical computation of invariant manifolds. in D. Benest et C. Froeschlé (editors): Modern Methods in Celestial Mechanics, Editions Frontières (1990), pp. 285-330.
- [139] C. Simó Effective computations in Hamiltonian dynamics. Mécanique céleste, 23 pp., SMF Journ. Annu., 1996, Soc. Math. France, Paris, 1996.
- [140] S. Smale, Differentiable dynamical systems. Bull. Amer. Math. Soc., 73: 747-817, 1967.
- [141] Stephen Smale. Diffeomorphisms with many periodic points. In Differential and Combinatorial Topology (A Symposium in Honor of Marston Morse), pp. 63-80. Princeton Univ. Press, Princeton, N.J., 1965.
- [142] D. C. Sorensen, The k -step Arnoldi process. In Large-scale numerical optimization (Ithaca, NY, 1989), pp. 228-237. SIAM, Philadelphia, PA, 1990.
- [143] A. M. Stuart and A. R. Humphries, Dynamical systems and numerical analysis. Cambridge University Press, Cambridge, 1996.
- [144] Floris Takens, Forced oscillations and bifurcations. Comm. math. Inst. Rijksuniversiteit Utrecht, 3: 1-59, 1974.
- [145] Floris Takens, Detecting strange attractors in turbulence. In Dynamical systems and turbulence, Warwick 1980 (Coventry, 1979/1980), pp. 366-381. Springer, Berlin, 1981.
- [146] W. Tucker, The Lorenz attractor exists. Thesis, Uppsala, 1998.
- [147] M. van Veldhuizen, A new algorithm for the numerical approximation of an invariant curve. SIAM J. Sci. Statist. Comput. 8: 951-962, 1987.

- [148] F. Warner, Foundations of Differentiable Manifolds and Lie Groups. Scott, Foresman and Co., 1971.
- [149] A. Weigend and N. Gershenfeld, Time Series Prediction, Addison-Wesley, 1994.
- [150] J. Wilkinson, The Algebraic Eigenvalue Problem, Oxford University Press, 1965.
- [151] J. Wisdom, Two lectures on chaotic dynamics in the solar system. Nonlinear evolution and chaotic phenomena (Noto, 1987), 185-187, NATO Adv. Sci. Inst. Ser. B Phys., 176, Plenum, New York, 1988.
- [152] A. Wolf et al., Determining Lyapunov exponents from a time series. Phys. D 16: 285-317, 1985.
- [153] P. Worfolk, Zeros of equivariant vector fields: algorithms for an invariant approach. J. Symbolic Comput. 17: 487-511, 1994.
- [154] Y. Xiang, Computing Thom-Boardman singularities. Thesis, Cornell University, 1998.
- [155] Jean-Christophe Yoccoz, Recent developments in dynamics. In Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994), pp. 246-265, Basel, 1995. Birkhäuser.
- [156] Lai Sang Young, Dimension, entropy and Lyapunov exponents. Ergodic Theory Dynamical Systems, 2: 109-124, 1982.
- [157] Lai Sang Young, Entropy, Lyapunov exponents, and Hausdorff dimension in differentiable dynamical systems. IEEE Trans. Circuits and Systems, 30: 599-607, 1983.
- [158] Lai-Sang Young, Ergodic theory of attractors. In Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994), pp. 1230-1237, Basel, 1995. Birkhäuser.