Periodic Orbits of Vector Fields: Computational Challenges *

John Guckenheimer
Mathematics Department, Ithaca, NY 14853

Abstract

Simulation of vector fields is ubiquitous: examples occur in every discipline of science and engineering. Periodic orbits are frequently encountered as trajectories. We use solutions of initial value problems, computed via numerical integration, as a means of finding stable periodic orbits of vector fields. We expect numerical integration algorithms to be reliable and their output to be consistent with other means of analyzing the properties of vector fields. These expectations are not always met. This lecture describes several examples of this type, followed by a brief description of a new set of boundary value solvers that appear to give significantly improved methods for computing these difficult periodic orbits.

Introduction

Vector fields in Euclidean space are defined by systems of differential equations

$$\dot{x} = f(x, \lambda),\ x \in R^n,\ \lambda \in R^k$$

with λ a vector of parameters. Periodic orbits are nonequilibrium trajectories $x(t)$ that satisfy $x(T) = x(0)$ for some $T > 0$. The smallest such T is the period of the orbit. The local dynamics near a periodic orbit are typically determined by return maps. A cross-section Σ to a periodic orbit is an $n-1$ dimensional surface that intersects the orbit transversally. The return map is a discrete map of Σ to itself that associates to each point x the first point of intersection of $x(t)$ with Σ. For small enough cross-sections, the return map has a fixed point at its intersection with a periodic orbit. If the Jacobian of the return map at this fixed point has eigenvalues inside the unit circle, the orbit is asymptotically stable. Initial conditions in the basin of attraction of the periodic orbit have trajectories whose limit set is the periodic orbit.

We explore the dynamics of vector fields by computing solutions of the initial value problem with numerical integration algorithms. Theory and implementation of these

*Research partially supported by the Air Force Office of Scientific Research, the Department of Energy and the National Science Foundation.

algorithms has been highly refined. They have a long history of successful use in the study of myriad problems. We expect the computations with numerical integration algorithms to produce reliable results over moderate time intervals, especially when local errors are estimated and step lengths are adjusted to control the estimated errors. Our expectations are usually met, but failure is more prevalent than we generally acknowledge. Dynamical systems theory challenges the capability of numerical integration algorithms to provide comprehensive answers to questions about vector fields and flows. Its emphasis upon qualitative properties and bifurcation leads us to places where the properties of flows are difficult to resolve numerically. We use three examples as illustrations of this phenomenon.

Example 1: A planar cubic vector field

I studied the following four parameter family of planar vector fields with Gerhard Danglemayr over fifteen years ago [4].

$$\begin{aligned} \dot{x} &= y \\ \dot{y} &= -(x^3 + rx^2 + nx + m) + (b - x^2)y \end{aligned}$$

This system contains the unfolding of a codimension two bifurcation of an equilibrium point with a double eigenvalue zero in the presence of a rotational symmetry of the plane [4]. We were interested in exploring the effects of imperfections of the symmetry upon the unfolding. The analysis turned out to be significantly more complicated than we expected it would be when we began our investigations. Most of our conclusions were based upon perturbation analysis of the system as it was rescaled to be nearly integrable. There were aspects of our asymptotic analysis that we had difficulty verifying numerically. One of the most intriguing was the conclusion that there were parameter values for which this system has a single equilibrium point and four nested limit cycles. Polynomial vector fields with many limit cycles are of interest in connection with the still unsolved Hilbert's Sixteenth Problem [14].

My attempts to locate the parameter region with four nested limit cycles were unsuccessful at the time we did our work. (I was constrained in this effort by a computer that performed only a few thousand floating point operations per second. By contrast the workstation, I use now is over 10,000 times faster.) Ten years after we did this work, Salvador Malo reexamined this vector field as a project for a course that I taught. He discovered that the sought for parameter region and estimated its size. As the parameter b is varied with $(r, m, n) = (0.87, -1, -1.127921667)$, the width of the region is approximately $3 * 10^{-9}$. Moreover, the three inner limit cycles are very close to one another. Figure 1 shows a plot of the four cycles with an inset that shows an enlarged view of a region where the three inner cycles cross the x-axis. With orbits so close to one another, we are concerned about the accuracy of the numerical integration algorithms used to compute the orbits.

Example2: Canards in the van der Pol Equation

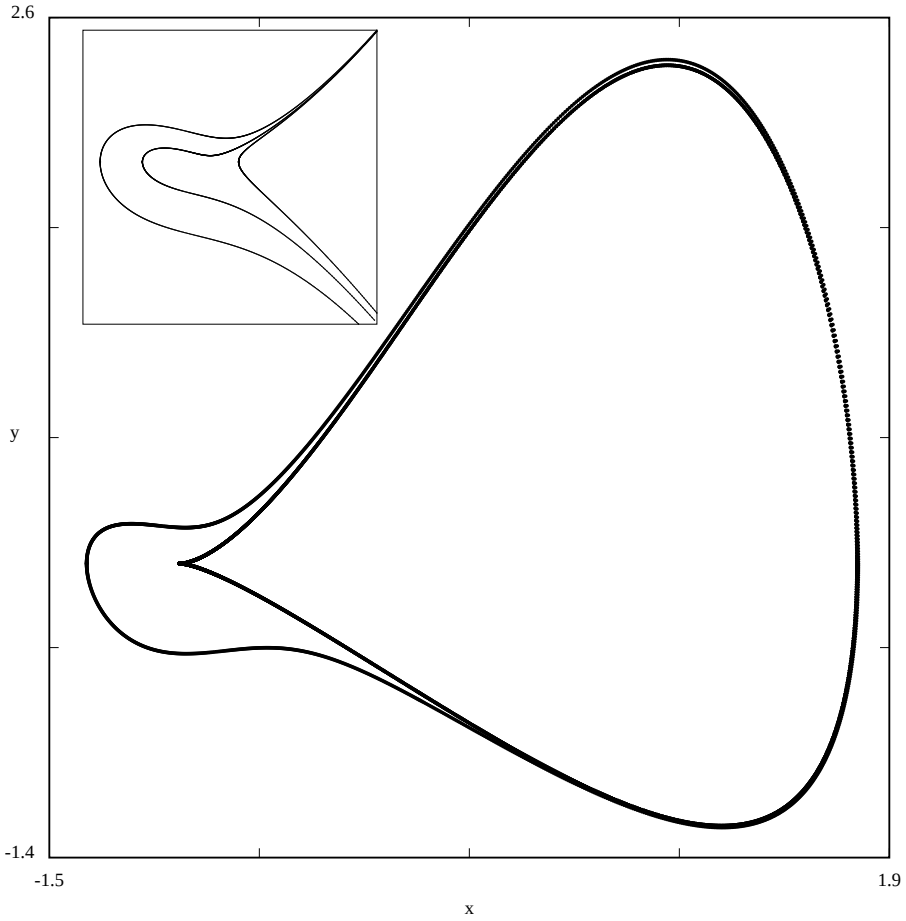

Figure 1: Four nested limit cycles in a cubic planar vector field. The inset shows the left ends of the three cycles on a finer scale

The van der Pol equation is a frequently studied system with relaxation oscillations, a limit cycle during which the speed of the orbit varies dramatically. The system can be transformed to the two dimensional vector field

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

Introducing a new parameter in place of the constant $-1/3$ gives the system

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

As a varies, the "slow manifold" $y = x^2 + x^3$ remains unchanged. On this manifold, the vector field is vertical and has a speed that remains bounded as $\epsilon \to 0$. Away from the

slow manifold, the speed of the vector field is $O(1/\epsilon)$. There is an equilibrium point on the slow manifold at $x = a$. When a passes through 0 or $-2/3$, Hopf bifurcation takes place. The equilibrium is a source for $-2/3 < a < 0$ and a sink for $a < -2/3$ or $0 < a$.

The periodic orbits born in the Hopf bifurcations grow in size to become the relaxation oscillations found at $a = -1/3$. The way in which they do so is a subtle and fascinating story that has been studied extensively from several points of view [5, 9, 8]. The shape of some of these orbits prompted a group of mathematicians from Strasbourg to call them canards [5]. "Canard" has become a technical term, referring to trajectory segments in a multiple time scale system that follow an unstable portion of a slow manifold. Periodic orbits containing canards are often called canards as well.

Numerical computation of canards with numerical integration is problematic. Along an unstable portion of a slow manifold, trajectories diverge from the slow manifold at an exponential rate comparable to $1/\epsilon$. Perturbations due to round-off error can be amplified in times that are $O(\epsilon)$ to large deviations from the canard. Consequently, even an "exact" numerical integration algorithm is unable to compute canards as solutions to the initial value problem when round-off errors prevent one from representing initial conditions that lie exactly on the slow manifold. Note that the slow manifold depends upon ϵ and is not known exactly in typical systems. In particular this is true of the extended van der Pol system. All solutions of the initial value problem can be expected to diverge from the unstable portion of the slow manifold in a characteristic time that may be smaller than the time of canards found in the family of periodic orbits.

The following numerical experiment illustrates the discussion of the last paragraph. Set $\epsilon = 0.001$ in the extended Hopf family. Beginning at the Hopf bifurcation point $a = 0$, compute periodic orbits for decreasing values of a. All of the periodic orbits are stable and are global attractors: their basins of attraction only exclude the equilibrium point. At some value of a, the attractor appears to jump discontinuously in size. For many numerical integration algorithms, these jumps appear in narrow ranges of a for which the numerically computed trajectories are chaotic, qualitative behavior that is impossible for the actual trajectories of a planar vector field. Figure 2 shows an example of such a trajectory, computed with a fourth order Runge-Kutta algorithm. Starting near the local minimum of the slow manifold at the origin, this trajectory begins to climb the unstable portion of the slow manifold. Unable to do so, sometimes it departs to the right and sometimes to the left. When it goes right, it makes a small loop to return to a neighborhood of the origin. When it goes left, the trajectory turns upward along the stable left branch of the slow manifold until it reaches its local maximum. From here, the trajectory turns right until it reaches the stable right branch of the slow manifold where it proceeds downwards until it returns to a neighborhood of the origin.

This example illustrates the limitations of initial value solvers to compute correctly the qualitative behavior of dynamical systems. With varying parameters, even planar vector fields which cannot have chaotic solutions present substantial numerical difficulties. Note that the ratio of time scales in this example is mild compared to typical models of chemical reactors.

Example 3: Hodgkin-Huxley Equations

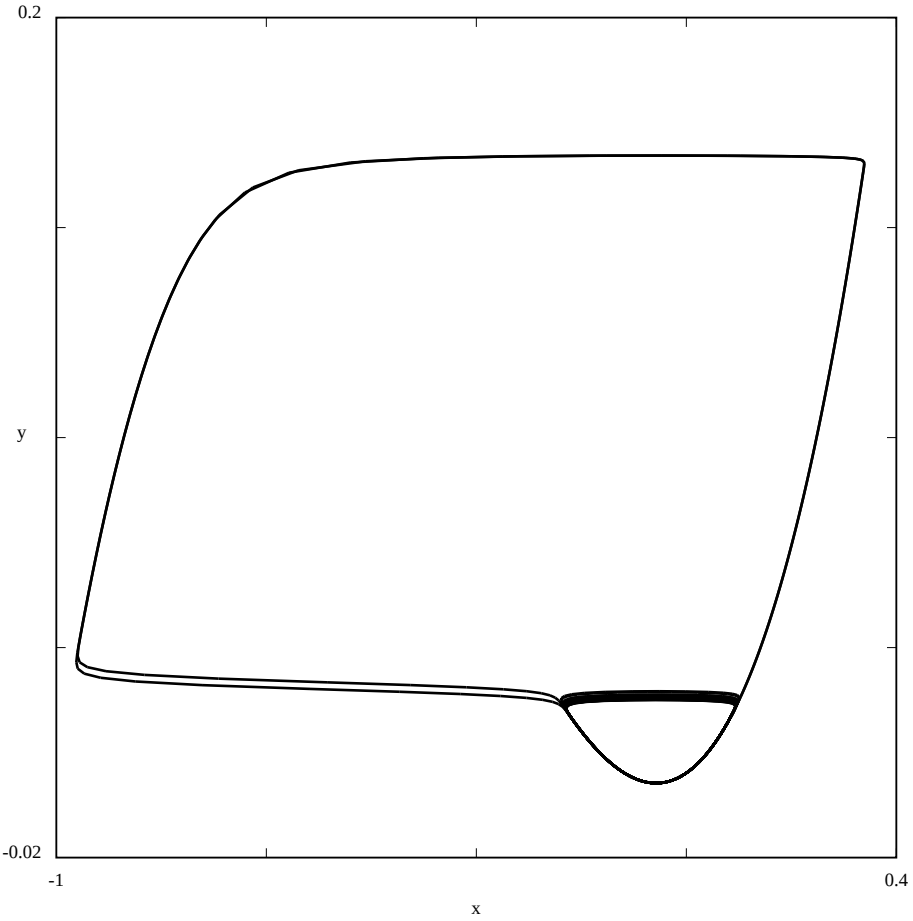

Figure 2: Numerical solutions to the extended van der Pol equation. The numerical solutions are chaotic despite the fact that planar vector fields have no chaotic dynamics.

Hodgkin and Huxley [15] developed a model for the action potential of the squid giant axon that became a paradigm for electrophysiological models of membranes of living organisms. The Hodgkin-Huxley equations relate the difference of electric potential across the cell membrane (V) and gating variables for sodium (m and h) and potassium (n) ion channels through the following equations:,

$$\left\{\begin{array}{rcl} \dot{V} &=& -G(V,m,n,h)+I \\ \dot{m} &=& \Phi(T)\left[(1-m)\alpha_m(V)-m\beta_m(V)\right] \\ \dot{n} &=& \Phi(T)\left[(1-n)\alpha_n(V)-n\beta_n(V)\right] \\ \dot{h} &=& \Phi(T)\left[(1-h)\alpha_h(V)-h\beta_h(V)\right] \end{array}\right. \qquad \text{(HH)}$$

where $\dot{x}$ stands for dx/dt and Φ is given by $\Phi(T) = 3^{(T-6.3)/10}$. The other functions involved are:

$$G(V,m,n,h) = \bar{g}_{\mathrm{Na}} m^3 h (V-\bar{V}_{\mathrm{Na}}) + \bar{g}_{\mathrm{K}} n^4 (V-\bar{V}_{\mathrm{K}}) + \bar{g}_{\mathrm{L}} (V-\bar{V}_{\mathrm{L}})$$

and the equations modeling the variation of membrane permeability:

$$\begin{array}{rclrcl} \alpha_m(V) &=& \Psi(\frac{V+25}{10}) & \beta_m(V) &=& 4e^{V/18} \\ \alpha_n(V) &=& 0.1\Psi(\frac{V+10}{10}) & \beta_n(V) &=& 0.125e^{V/80} \\ \alpha_h(V) &=& 0.07e^{V/20} & \beta_h(V) &=& \left(1+e^{(V+30)/10}\right)^{-1} \end{array}$$

$$\text{with} \qquad \Psi(x) = \begin{cases} x/(e^x-1) & \text{if } x \neq 0 \\ 1 & \text{if } x = 0 \end{cases} .$$

Notice that $\alpha_y(V) + \beta_y(V) \neq 0$ for all V and for $y = m$, n or h. The parameters $\bar{g}_{\text{ion}}$, $\bar{V}_{\text{ion}}$ represent maximum conductance and reversal potential for the ion with the values given below:

$$\begin{array}{lll} \bar{g}_{\text{Na}} = 120 \text{ mS/cm}^2 & \bar{g}_{\text{K}} = 36 \text{ mS/cm}^2 & \bar{g}_{\text{L}} = 0.3 \text{ mS/cm}^2 \\ \bar{V}_{\text{Na}} = -115 \text{ mV} & \bar{V}_{\text{K}} = 12 \text{ mV} & \bar{V}_{\text{L}} = 10.599 \text{ mV} \end{array}$$

The normalization of the potential used by Hodgkin and Huxley differs from modern conventions in sign and origin. The parameter I represents an external current injected into the axon with an electrode and T is the temperature of the preparation, taken to be 6.3 in our computations.

While this model established the principles for electrophysiological descriptions of neurons and cardiac tissue, its dynamics have not been thoroughly investigated [12]. Attempts to compute the dependence of solutions upon parameters have encountered numerical difficulties that have yet to be surmounted. Comparison with the canards of Example (2) is appropriate. As the injected current is increased, the system develops a very stable limit cycle representing tonic action potentials. These oscillations can be described in terms of the influx of sodium ions and the outflux of potassium ions through the membrane channels. The activation of the sodium channels represented through the gating variable m is much faster than the other membrane processes included in the model. As I is varied, a Hopf bifurcation occurs, giving rise to a family of periodic orbits that connects to to the stable limit cycles described above. The Hopf bifurcation is subcritical [16], though the Lyapunov number that determines this fact is small and difficult to calculate.

Computing the branch of periodic orbits that connect the Hopf bifurcation to the stable limit cycles representing tonic action potentials is even more difficult. The system does not include an explicit small parameter, but some of the orbits vary rapidly in amplitude with changes in I and contain highly unstable segments like canards. Since the orbits themselves are of saddle type with Floquet multipliers that are both larger and smaller than one, the orbits cannot be computed as the limit sets of trajectories obtained by solving an initial value problem. We report below new observations about this family of periodic orbits, demonstrating that the family does not take the simplest possible path that could connect the Hopf bifurcation to the stable limit cycles.

Boundary Value Methods

The examples described above motivate the use of boundary value methods to find periodic orbits. However, there is a smaller literature and less software devoted to bound-

ary value solvers for periodic orbits [7] than for two point boundary value problems [1]. Periodic boundary value problems can be recast as larger two point boundary value problems with separated boundary conditions, but this is seldom done and there is little evidence that such methods work well. Apart from the code AUTO [6] and varied implementations of single shooting algorithms, there is a dearth of software for computing periodic orbits. As with the solution of initial value problems, I believe that no single method or code for solving these problems will be optimal in all situations. Therefore, I have undertaken development of a new set of boundary value solvers for finding periodic orbits. I describe here the philosophy and structure of these algorithms and then examine their effectiveness in addressing the issues raised in the examples presented above.

Boundary value methods fall into two main classes: shooting methods and global methods. The distinction between the two is the following. Shooting methods compute approximate trajectory segments with an initial value solver, matching the ends of these trajectory segments with each other and the boundary conditions. One interpretation of shooting methods is that the initial value solver produces approximations to the flow map of the vector field, and the boundary conditions are expressed in terms of the approximate flow map. Global methods project the differential equations onto a finite dimensional space of curves that satisfy the boundary conditions. Collocation methods have become the predominant global methods for finding periodic orbits, especially through the use of the computer code AUTO [6]. In collocation methods, the discretization of the differential equations consists in satisfying the differential equations along polynomial curves at specific mesh points.

Discussions of boundary value methods tend to focus upon the size of the problem being solved and the order of asymptotic accuracy. Here my interest is in how well algorithms work when implemented on typical computers and tested on a broad selection of dynamical systems. There are choices for alternate ways of implementing each part of an algorithm and for many algorithmic parameters that affect the performance of the algorithm. This makes a general comparison of algorithms difficult. Therefore, I focus upon the ability of the methods we have developed to solve problems that I have found difficult to solve in the past.

Both shooting and global methods yield sets of equations that are solved with standard algorithms, usually Newton's method. This requires that the Jacobian J of the set of equations be computed and that linear systems of the form $Ju = v$ be solved for u. The robustness of Newton's method depends upon several factors, including the condition number and accuracy of J. If J is obtained from finite difference calculations with numerical integrations that themselves have only moderate precision, the root finding tends to become erratic. Proper formulation of the boundary value problems and highly accurate implementation of the methods is required to obtain algorithms that perform robustly. The size and structure of J depend upon the dimension of the vector field, the number of points used in the discretization of the periodic orbit and the algorithm. Since the dimension of the vector field is given as part of the problem, we concentrate on using coarse meshes for the root finding while maintaining high orders of accuracy.

Finite difference calculations of derivatives have limited accuracy that stems from the balance between truncation and round-off errors. Truncation errors increase with

increment size while round-off errors decrease with increment size. To circumvent these difficulties, we have employed *automatic differentiation* in our algorithms. Automatic differentiation [2] is a technique for computing derivatives of elementary functions that makes use of explicit formulas for the derivatives of basic functions. By allocating storage for intermediate expressions, differentiation rules can be applied to codes without generating symbolic expressions for the derivatives. I have used the code ADOL-C [10] to compute Taylor series expansions of solutions to ordinary differential equations and to compute derivatives of the Taylor coefficients with respect to phase space variables and parameters. These computations can be carried out to high order without truncation error: the results appear to have accuracy close to the floating point precision. This enables Taylor series methods to be used effectively in our algorithms.

Conceptually, our methods are straightforward implementations of simple ideas. Shooting codes layer root finding on top of numerical integration, seeking periodic orbits through a set of points $x_1, \ldots x_N$ and times $s_1, \ldots s_N$ with the property that the flow of the system satisfies $\Phi(x_i, s_i) = x_{i+1}$ where $x_{N+1} = x_1$. In our codes Φ and its derivatives are approximated by piecewise polynomial functions whose coefficients are computed with automatic differentiation. Step lengths for the piecewise polynomials are determined by the apparent growth rates of the Taylor series coefficients. Mesh sizes are determined so that the condition number of the Jacobian for system of equations $\Phi(x_i, s_i) = x_{i+1}$ does not become excessively large. One of the big advantages of these methods is that the order of accuracy for the numerical integration depends upon the degrees of the Taylor polynomials that are computed. This can be chosen to be large, compared to the order of accuracy of prevailing numerical methods. A second advantage of the method is that the computation of derivatives of the flow map requires almost no additional programming effort beyond that required to describe the system itself.

We have also used Taylor series in the formulation of global methods for computing periodic orbits. The global algorithms that worked best with the examples we studied were based upon Hermite polynomial interpolation. Hermite polynomials were computed as the minimal degree polynomials that matched the Taylor series expansions computed at endpoints of each mesh interval. The discretized equations that we used in these algorithms stated that the tangent vector to the Hermite polynomial was given by the differential equation at the midpoint of each mesh interval. These methods work well, but are somewhat more complex than the shooting methods, so we emphasize the shooting methods here. Details about the implementations are reported elsewhere [13].

Test Results

I used our multiple shooting algorithms to track periodic orbits through the family of periodic orbits obtained by varying b in the first example above. A tangent approximation to the family of periodic orbits is computed at each set of parameter values and used to set initial choices for the next orbit in the family. Tracking the family smoothly near folds (saddle-node bifurcations of the periodic orbits) requires that the parameter b be "free" and allowed to vary during the root finding. This is done by restricting updates to the subspace orthogonal to the computed tangent vector to the family of periodic orbits.

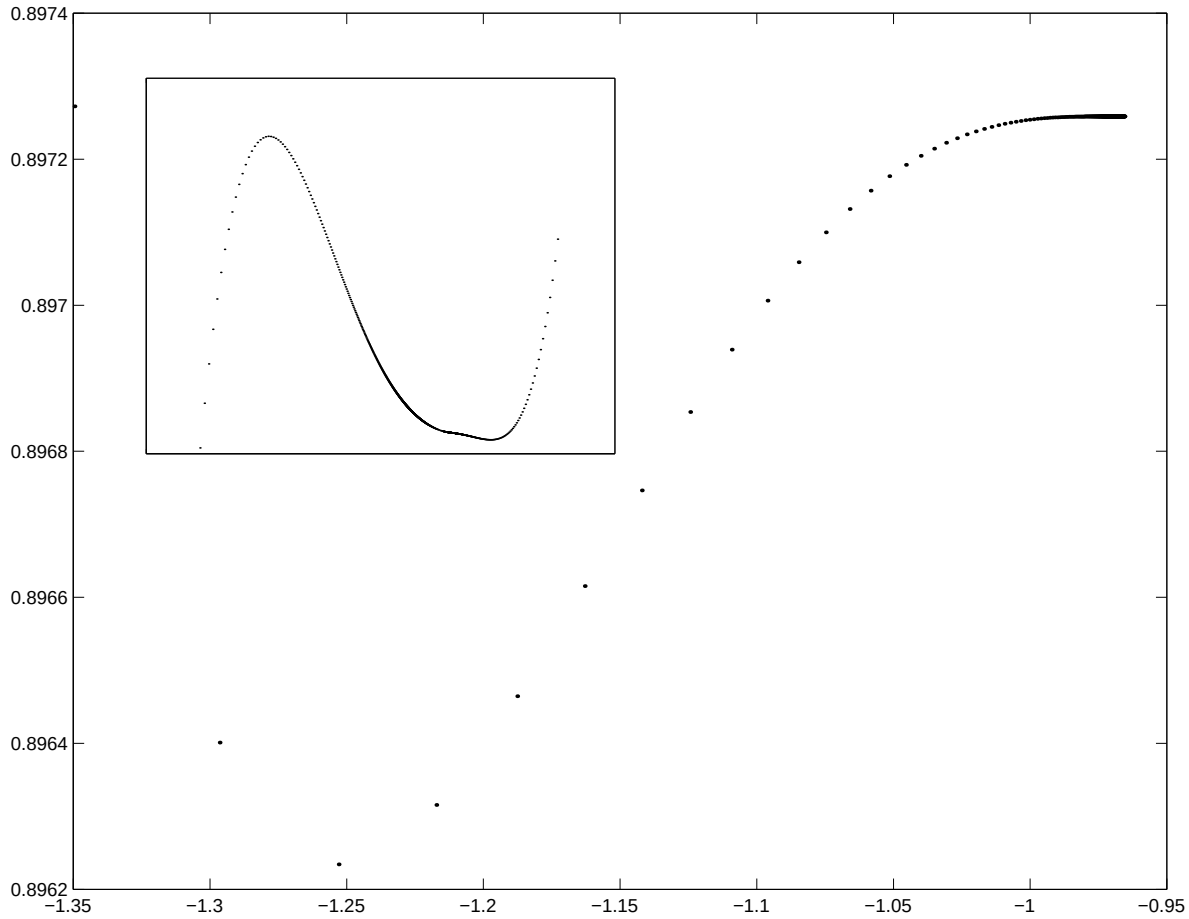

Figure 3: The parameter b versus the intersection of limit cycles with the x-axis for the family of periodic orbits that connects the limit cycles shown in Figure 1. The inset shows an expanded view of the data on the far right of the figure.

Figure 3 plots the value of b versus the left hand intersection of the orbits with the x-axis. Since the cross product $(b_1 - b_2)y^2$ of the vector fields for two values b_1, b_2 of b does not change sign, the periodic orbits in the family are all disjoint from one another. The inset shows the right hand portion of the curve on a finer scale. This computation was quite challenging because the periods of the periodic orbits increase to approximately 1000 in the middle of the family. There are a pair of complex equilibrium points with very small imaginary parts (approximately 0.00007) and the vector field slows in the vicinity of these equilibrium points. Furthermore, the nearby presence of the equilibrium points limits the radius of convergence of the Taylor series for the trajectories, so that a large number of time steps are required to compute the periodic orbits that pass through this region. Nonetheless, the algorithms appear to do an excellent job of tracking the family of periodic orbits. Figure 4 plots the periods of the periodic orbits as a function of the left hand intersection of the orbits with the x-axis. The large periods have no apparent relationship to the saddle-node bifurcations in the family, only to close passage near the

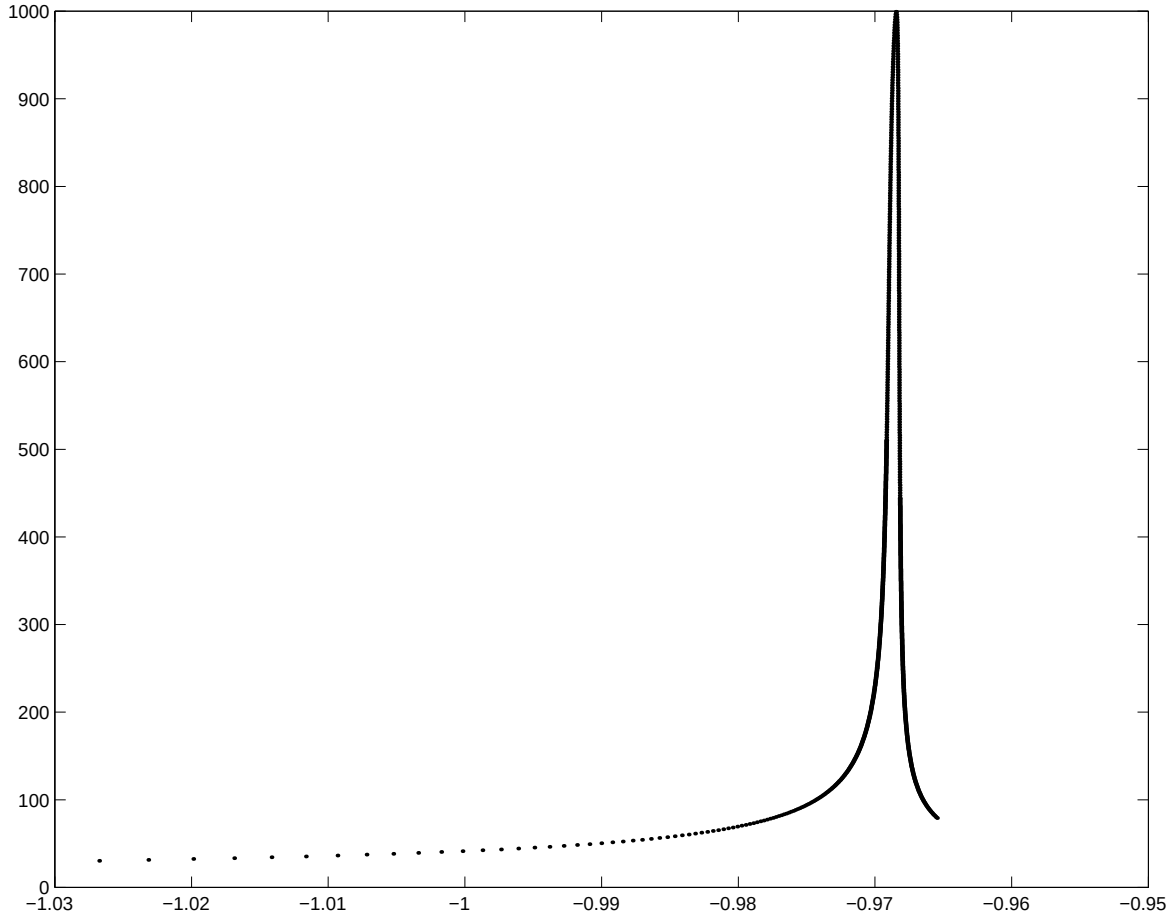

Figure 4: Periods of the limit cycles versus the intersection of limit cycles with the x-axis for part of the family of periodic orbits shown in Figure 3.

complex equilibrium points.

Figure 5 shows the family of canard solutions to the extended van der Pol family computed with our multiple shooting code. The parameter $\epsilon = 0.001$. The asymptotic theory for this singularly perturbed system indicates that the canards occur over a range of parameter values of length smaller than 10^{-20}, too small to be resolved in our computations with IEEE-754 double precision arithmetic. Therefore, we again let the parameter a vary, using the left hand intersection of the periodic orbits with the cubic characteristic $y = x^2 + x^3$ to determine the amplitude of the periodic orbit. Without a quantity to measure this amplitude, it is hard to consistently choose a direction along the tangent to the family of periodic orbits that varies monotonely. At selected places along the family of periodic orbits, the convergence of Newton's method slowed markedly from its normal quadratic convergence. This appeared to be related to the points of the computed discrete periodic orbit being out of order, with some time intervals going in a negative direction. By monitoring for this behavior and deleting the offending points, most difficulties with convergence were avoided.

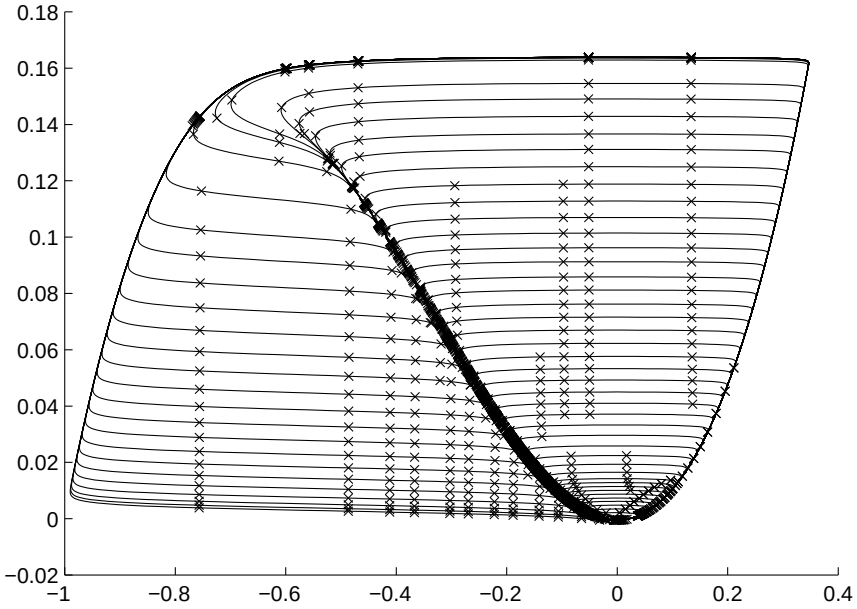

Figure 5: The family of canard solutions that connect the Hopf bifurcation in the extended van-der Pol equation to relaxation oscillations. The symbol "x" denotes mesh points of the discrete closed curve produced by the multiple shooting algorithm.

Figure 6 plots the maximum membrane potential v along a family of periodic orbits versus the injected current parameter I in the Hodgkin-Huxley equations. The Hopf bifurcation occurs near the bottom of the curve with $(I, v) \approx (16.14, -5.35)$. Following the curve upwards, it reaches a minimum value of I at a turning point where there is a saddle-node of periodic orbits. The branch of orbits that slopes gently down towards the right consists of large amplitude stable periodic orbits. The branch of periodic orbits that connects the turning point with the Hopf bifurcation is more complicated. This branch has two additional turning points that are hard to distinguish at this scale. To understand better the qualitative features of this branch of periodic orbits, the eigenvalues of the monodromy were computed. (The *monodromy* of a periodic orbit of period T is the Jacobian of the time T map at a point of the orbit. Choosing different basepoints changes the monodromy by a similarity transformation.)

Figure 7 shows some details of the monodromy in the region near the middle fold of the branch. The horizontal axis of the figure gives the injected current I shifted by 14.28137, while the vertical axis represents the logarithm of the magnitude of two eigenvalues of the monodromy. The eigenvalues not shown are those of largest and smallest magnitude. There is always one very small eigenvalue of the monodromy whose magnitude is not resolved by our computations but is smaller than 10^{-9}. The magnitude of the largest eigenvalue of the monodromy is much larger than 1 on this portion of our family of orbits. Thus, the periodic orbits on the middle portion of the branch in Figure 6 with negative slope are not stable. One of the two remaining eigenvalues of the monodromy is always 1, with eigenvector pointing along the vector field. Computing this eigenvalue correctly

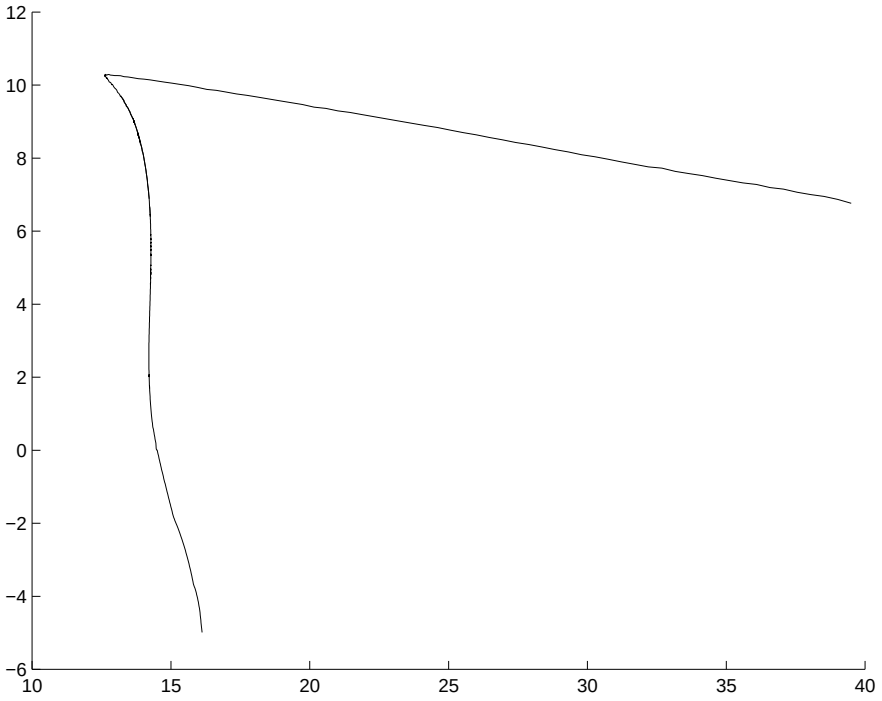

Figure 6: Maximum membrane potential v versus injected current I along a family of periodic orbits in the Hodgkin-Huxley model. The Hopf bifurcation occurs at the lower end of the curve near $I = 16$.

provides a check for our calculations. The final eigenvalue of the monodromy is the one that follows the curve in Figure 7. This eigenvalue starts with a positive value less than 1 on the part of the branch in Figure 6 with positive slope. As it reaches the saddle-node of the family of periodic orbits, it crosses the unit circle and the log of its magnitude becomes positive. This happens at the right hand crossing of the horizontal axis in Figure 7. This eigenvalue then increases in magnitude until it meets the eigenvalue of largest magnitude. The two eigenvalues become complex at this point and have the same magnitude. The two complex eigenvalues cross the imaginary axis, acquiring negative real parts before they become real and split once more. The complex eigenvalues have magnitude larger than 50. They give rise to the nearly horizontal segment at the top of the eigenvalue curve in Figure 7. After the complex eigenvalues become real once again, the magnitude of the smaller one decreases and becomes smaller than 1 in magnitude once again. It is negative in this regime, so a period doubling bifurcation occurs as it crosses -1.

The behavior revealed in Figures 6 and 7 is substantially more complicated than I expected to find in this example. Conventional wisdom is that the Hodgkin-Huxley equations reduce to a two dimensional system, with an attracting two dimensional manifold in

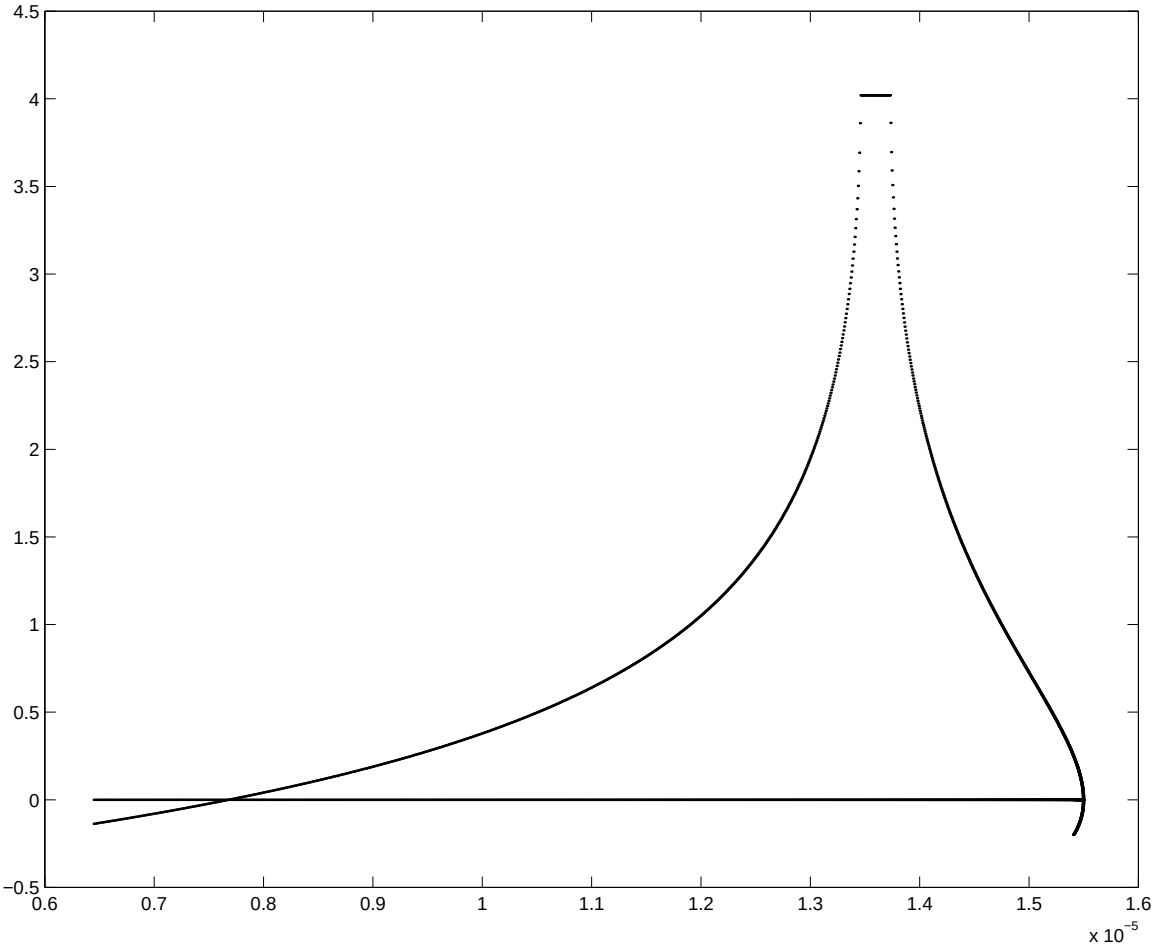

Figure 7: Logarithms of the magnitudes of the two intermediate eigenvalues of the monodromy versus injected current I for a small portion of the family of periodic orbits shown in Figure 6. One of the eigenvalues becomes unstable at a saddle-node bifurcation, then becomes complex at the plateau in the figure and then becomes stable again at a period doubling bifurcation.

its phase space. These computations show that this is not at all the case when we examine unstable as well as stable limit sets for the system. Periodic orbits of two dimensional vector fields do not have complex eigenvalues for their monodromy. The behavior that we have found also suggests that even more complex dynamics may be lurking within this system. I have not yet tried to follow the branch of periodic orbits that arises in the period doubling bifurcation displayed here to see whether it undergoes further bifurcations that yield additional periodic orbits. The point I wish to emphasize here is that thorough analysis of this classical model in neurophysiology requires extremely good computational tools and careful guidance stemming from our theoretical understanding of dynamical systems theory.

The Future

Taylor series methods have not been widely adopted in solving differential equations, despite evidence that they are more accurate than other methods and are as fast to compute [3]. Presumably, their slow acceptance has been due to the additional programming effort required for their implementation. Robust software packages are commonly used for numerical integration, but ones that incorporate automatic differentiation and use Taylor series have not existed. I hope that the examples in this lecture - and many more similar examples - will stimulate wider adoption of Taylor series methods. We have the computing resources needed to make these methods easy to use. The development of software that implements algorithms for computing periodic orbits and their bifurcations is an ongoing process. I hope that by the next Equadiff conference, we will have a set of tools that enable us to examine periodic orbits of vector fields with similar precision and detail to our investigations of equilibrium points.

References

[1] U. Ascher, R. Mattheij and R. Russell, Numerical Solution of Boundary Value Problems, Prentice Hall, 1988.

[2] M. Berz, C. Bischof, G. Corliss and A. Griewank (eds.), Computational Differentiation, SIAM, 1996.

[3] George Corliss and Y. F. Chang. Solving ordinary differential equations using Taylor series. ACM Trans. Math. Software, 8(2):114-144, 1982.

[4] G. Dangelmayr and J. Guckenheimer, On a four parameter family of planar vector fields, Arch. Rat. Mech. Anal., 97, 321-352, 1987.

[5] M. Diener, Canards et bifurcations, in *Mathematical tools and models for control, systems analysis and signal processing, Vol. 3 (Toulouse/Paris, 1981/1982)*, 289–313, CNRS, Paris.

[6] E. Doedel, AUTO, ftp://ftp.cs.concordia.ca/pub/doedel/auto.

[7] Eusebius Doedel, Herbert B. Keller, and Jean-Pierre Kernévez. Numerical analysis and control of bifurcation problems. Bifurcation in finite dimensions. Internat. J. Bifur. Chaos Appl. Sci. Engrg., 1(3):493–520, 1991 and 1(4):745–772, 1991.

[8] F. Dumortier and R. Roussarie, Canard cycles and center manifolds, Mem. Amer. Math. Soc. **121** (1996), no. 577, x+100 pp.

[9] W. Eckhaus, A standard chase on French ducks, Lect. Notes in Math. 985, 449-494, 1983.

[10] Griewank, A., Juedes, and J. Utke, 'ADOL-C: A Package for the Automatic Differentiation of Algorithms Written in C/C++', Version 1.7, Argonne National Laboratory.

[11] J. Guckenheimer and W. G. Choe, Computing periodic orbits with high accuracy, Computer Methods in Applied Mechanics and Engineering, 170, 331-341, 1999.

[12] J. Guckenheimer and I. Labouriau, Bifurcation of the Hodgkin Huxley Equations: A New Twist, Bulletin of Math. Biol.,55, 937–952, 1993.

[13] J. Guckenheimer and B. Meloon, Computing Periodic Orbits and their Bifurcations with Automatic Differentiation, preprint, 1999.

[14] D. Hilbert, Mathematical problems, Bull. Amer. math. Soc., 8, 437-479, 1902.

[15] A.L.Hodgkin and A.F.Huxley, *A quantitative description of membrane current and its applications to conduction and excitation in nerve,* J.Physiology 117, 500-544, 1952.

[16] I.S.Labouriau, *Degenerate Hopf bifurcation and nerve impulse, Part II,* SIAM J. Math. Anal. 20, 1-12, 1989.