

Computing Periodic Orbits *

John Guckenheimer

Mathematics Department, Ithaca, NY 14853

Abstract

The dynamics of fluids are often studied through the reduction of partial differential equations to systems with only a few degrees of freedom. Analysis of these low dimensional vector fields remains difficult in many examples. The simplest dynamical behaviors are equilibria representing steady flow and periodic orbits. 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 a new set of boundary value solvers that appear to give significantly improved methods for computing “difficult” periodic orbits.

Introduction

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

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

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 algorithms has been highly refined. They have a long history of successful use in the study of myriad problems.

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

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 Dangelmayr 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 arose in the context of bifurcations of “doubly diffusive” convection problems. These are fluid systems in which convection is driven by buoyancy due to temperature variations in the fluid and a second force, for example Coriolis forces arising from rotation, magnetic forces arising from moving electrical charges within the fluid, or additional buoyancy forces arising from concentration variations. At particular balances of the different forces, the system can undergo “codimension two” bifurcation in which the linearized system has a two dimensional invariant subspace with eigenvalue zero. With some boundary conditions having modest symmetry, the reduction of the system yields the two dimensional vector field above with $r = m = 0$ [12]. The full system above incorporates the effects of imperfect symmetry on the dynamics.

The analysis of this vector field turned out to be significantly more complicated than Dangelmayr and I 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. Some 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 [15].

My attempts to locate the parameter region with four nested limit cycles were unsuccessful at the time we did our work. 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.




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

Example2: Canards in the van der Pol Equation

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 $\epsilon \rightarrow 0$. 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.




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.

Example 3: Josephson Junctions

Equations of fluid dynamics often have spatial symmetries that affect their dynamics. The viewpoint of dynamical systems theory is to study “generic” properties, but the presence of symmetry constrains a system to behave in ways that are not generic. Understanding the consequences of symmetry has been a focal point for the subject for twenty years. A set of tools have been developed to explore this issue, but these tools do not give a comprehensive analysis of the dynamics expected in the presence of a particular symmetry. Indeed, there are examples [13] that demonstrate that a full range of dynamical behavior, including chaos, can result from bifurcations of equilibria in symmetric systems. Numerical methods are needed to examine the dynamics of symmetric systems, including the normal forms associated with bifurcations. Here we illustrate the problems of examining periodic orbits in symmetric systems with an example of coupled oscillators.

Swift and Watanabe [16] studied the dynamics of arrays of N Josephson junction oscillators shunted by a series LRC load, a system symmetric with respect to all interchanges of the

oscillators. These systems have “splay-phase” solutions in which the N oscillators all undergo the same motion but out of phase with one another. Denoting the motion of oscillator i in the splay phase state by $\phi_i(t)$, the oscillators can be numbered so that $\phi_{i+1}(t) = \phi_i(t + 2\pi/N)$. In the limit that the Stewart-McCumber parameter of these systems is zero, the splay phase state is neutrally stable. Numerical integration of the system with other values of this parameter suggested that the splay phase states might be neutrally stable for positive values of the Stewart-McCumber parameter as well, with a return map possibly having eigenvalue 1 with multiplicity $N - 3$. Swift and Watanabe studied the case of four coupled junctions carefully and demonstrated that the splay phase states are represented by hyperbolic periodic orbits for generic values of the Stewart-McCumber parameter. They did so by computing the splay phase states using the program AUTO86 [6] and then computing the monodromy of the periodic orbit outside of AUTO. (The computation of monodromy matrices has been greatly improved in AUTO97 compared to AUTO86, but AUTO97 was not available when Swift and Watanabe did their work.)

The vector field f that represents an array of four junctions is defined by the ten dimensional system:

$$\begin{aligned}
\dot{x}_1 &= x_5 \\
\dot{x}_2 &= x_6 \\
\dot{x}_3 &= x_7 \\
\dot{x}_4 &= x_8 \\
\dot{x}_5 &= (I - x_5 - \sin(x_1 - x_{10}))/b \\
\dot{x}_6 &= (I - x_6 - \sin(x_2 - x_{10}))/b \\
\dot{x}_7 &= (I - x_7 - \sin(x_3 - x_{10}))/b \\
\dot{x}_8 &= (I - x_8 - \sin(x_4 - x_{10}))/b \\
\dot{x}_9 &= x_{10} \\
\dot{x}_{10} &= ((x_1 + x_2 + x_3 + x_4)/4 - rx_{10} - x_9/c)/l
\end{aligned}$$

Figure 3 is a plot of the splay phase solutions to these equations with parameter values $I = 2.5, l = 0.75, r = 0, c = 20$ and $b = 0.2$.

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 boundary 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




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.

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.

Both shooting and global methods yield sets of equations that are solved with standard algorithms, usually Newton's method. Newton's method uses the Jacobian J of the set of equations being solved and includes solution of the linear system $Ju = v$. 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. Here we describe the most direct implementation of multiple shooting. Shooting codes layer root finding on top of numerical integration, seeking periodic orbits through a set of points $x_1, \dots, x_N$ and times $s_1, \dots, s_N$ 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. In the simplest methods, the polynomials are the Taylor polynomials of the trajectories. The order of accuracy for the numerical integration in these methods 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. Step lengths for the numerical integration of Taylor polynomials are determined by the apparent growth rates of the Taylor series coefficients. We employ a varying number of integration steps in each mesh interval to maintain coarse meshes and limit the size of the systems solved with Newton's method. Mesh sizes are determined by evaluating the norm of the Jacobian for Φ . We do not want the condition number of the system of equations $\Phi(x_i, s_i) = x_{i+1}$ to become excessively large. Placing a limit on the norm of $D\Phi$ is a strategy to achieve this. The computation of $D\Phi$ and the Jacobian appearing in Newton's method 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




Figure 4: 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.

than the shooting methods, so we emphasize the shooting methods here. Details about the implementations are reported elsewhere [14].

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 4 plots the value of b versus the left hand intersection of the orbits with the x -axis. Since the cross product $(b_1 - b_2)y^2$




Figure 6: 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.

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.

Since Swift and Watanabe had difficulty computing the splay phase states of the Josephson Junction equations and their stability, we sought to confirm their results with an automatic differentiation based solver. We used the Hermite interpolation global method to compute the splay phase state, starting with linear varying phase for each oscillator. The parameter values were $I = 2.5, l = 0.75, r = 0, c = 20$ and Stewart-McCumber parameter $b = 0.2$. With twenty mesh intervals and degree sixteen Taylor series approximations at the mesh points, the algorithm converged in six Newton steps to an approximate periodic orbit $p(t)$. Using the uniform norm in R^{10} , the norm of $p'(t) - f(p(t))$ appears to be smaller than 2×10^{-13} . The

eigenvalues of the monodromy matrix were computed to be

$$\begin{aligned}
&1.212564610112479e - 06 \pm 5.700237500982539e - 08i \\
&1.390021921820548e - 06 \\
&-1.172334117548194e - 03 \pm 4.455213497385255e - 04i \\
&8.826221531499485e - 01 \\
&1.0000000000000006e + 00 \\
&1.003009060195232e + 00 \\
&1.149723251975266e + 00 \pm 5.356810539765165e - 02i
\end{aligned}$$

The monodromy matrix of a periodic orbit has an eigenvalue 1. The precision with which this eigenvalue is computed is a measure for the accuracy of the computation of the periodic orbit and its monodromy matrix. In our calculation, the error for this eigenvalue is 6×10^{-15} . We also note that the eigenvalue 1.003009060195232 is in excellent agreement with the perturbation analysis of Swift and Watanabe [16], Figure 8. As a final check on the precision of these calculations, we repeated them with a coarser mesh having ten mesh intervals. The algorithm converged in seven Newton steps. The monodromy matrix agreed with the values in the table above to ten decimal places. The error in the eigenvalue 1 was approximately 5.5×10^{11} . Thus, these periodic orbit calculations seem to be converging quickly and accurately.

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. My experience with the examples shown here and others makes me confident that we can compute information about periodic orbits of dynamical systems with almost the same degree of precision that we compute information about equilibria of such systems.

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 P. Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcation of Vector Fields*, Springer-Verlag, 1983.
- [13] J. Guckenheimer and P. Worfolk, Instant Chaos, *Nonlinearity* 5, 1211–1222, 1992.
- [14] J. Guckenheimer and B. Meloon, Computing Periodic Orbits and their Bifurcations with Automatic Differentiation, preprint, 1999.
- [15] D. Hilbert, Mathematical problems, *Bull. Amer. math. Soc.*, 8, 437-479, 1902.
- [16] S. Watanabe and J. Swift, Stability of periodic solutions in series arrays of Josephson junctions with internal capacitance, *J. Nonlinear Sci.* 7, 503-536, 1997.