from *SIAM News*, Volume 32, Number 8

Beyond Simulation—Computing Dynamical Systems

For John Guckenheimer, his term as SIAM president (1997–98) was a unique opportunity "to observe and participate in the process by which we establish scientific policy in this country." Inevitably, he says, he brought his personal perspective to these activities. In this article, an adaptation of his Past President's Address, given in Atlanta at the 1999 SIAM Annual Meeting, he describes some of his experiences, expresses some opinions, and speculates on research trends.

SIAM PAST PRESIDENT'S ADDRESS

By John Guckenheimer

To achieve fidelity in simulations, says SIAM past president John Guckenheimer, "we need conceptual frameworks that facilitate deep understanding of how the components of complex systems interact. Part of the richness of dynamical systems theory is that it does precisely that."

Perhaps my biggest struggle as SIAM president was to articulate the "*reasonable* effectiveness of mathematics in computational science." I believe that many of our colleagues in other disciplines take mathematics for granted as a foundation for computation, even when their attitude is unwarranted. I have been surprised repeatedly during my career by situations in which the interplay between mathematics and applications is subtle. New mathematical discoveries and computational methods have been needed to resolve discrepancies, leading to deeper insight into both the mathematics and the science.

In this article, I emphasize this interplay, using problems I have encountered during my research on dynamical systems as examples. I take a broad view, illustrating how we can exploit the power of mathematics to distill the essence of problems and find solutions based on core principles and critical ideas.

Nonlinear dynamics is a thriving research area, one that I believe will continue to grow in importance. In part, this stems from the use of dynamical systems analysis to interpret simulations in three areas:

- Computational science. Computational science, much of it devoted to simulation, has joined theory and experiment as a mode of doing science.
- Industrial design. Using computer simulation as a design tool, companies seek to reduce time to market and improve reliability for new products.
- Public policy and systems. We simulate the influence of human activity in producing climate change as an input to government policy. U.S. national policy is to replace nuclear weapons testing with simulation.

Simulation has become an important part of the fabric of science, engineering, and the technological infrastructure we depend on every day. In all these efforts, fidelity of the simulations is important. The usefulness of our computer models and simulations depends on the degree to which they reproduce the phenomena being simulated. Nonlinear dynamics contributes to these efforts by discovering common patterns that appear in real, artificial, and simulated systems.

I am fascinated by the questions raised by simulation. Mathematics and science for me have always begun with the search for deep truths about the world. The simulation process itself is an intriguing subject of investigation, posing a rich set of mathematical questions that have stimulated my research. The exact reproducibility of numerical calculations sets computation apart from other forms of experiment. I am also intrigued by the phenomena that we simulate and their relation to the simulations themselves. I find myself dissatisfied by pronouncements of a "good fit" between simulation and data in cases in which little effort has been made to objectively evaluate the fidelity of the simulations. Dynamical systems theory has been successful in revealing universal patterns of "nonlinear" behavior that serve as signposts in efforts to match data with simulations. As we attempt to improve computational methods for bringing theory to bear on simulations of complex systems, I want to recall some of the subtle issues that have arisen in the interactions between mathematics, modeling, and experimental science through a series of examples.

Example 1: Genericity in Experiment and Computation

I started to do research on dynamical systems theory as a graduate student. A nice feature of the subject was that it was supposed to say something about the "real" world, but that was incidental to my involvement with mathematics. My professional interest in applications began when I read René Thom's book *Structural Stability and Morphogenesis*. I was fascinated with his ideas about the role of mathematics in biology, in particular with the thesis that genericity leads to ubiquitous patterns.

Genericity is a subtle concept, hard to explain in nontechnical terms and difficult to apply quantitatively. Indeed, there was

substantial controversy about catastrophe theory during the 1970s. The starting point for the use of genericity as a modeling strategy is a universe of possible models. We imagine this universe as being very large, although concrete representations of models have a finite number of parameters. Thom and others developed singularity theory, elegant mathematics that describes how static patterns of generic models change with varying parameters. In the setting of dynamical systems, bifurcation theory classifies generic changes in the properties of vector fields and their flows.

My emerging interest in biology led me to a collaborative study of population models with George Oster, beginning in 1973. The models we were studying twenty-five years ago were discrete mappings of two variables, noninvertible mappings that generalize iterations of the quadratic function $f(x) = rx(1 - x)$, known as the logistic map by population biologists. I thought I had proved in my thesis that for a dense set of parameters in the interval $0 < r < 4$, this map has a stable periodic orbit that attracts almost all initial conditions. (There were serious gaps in my proof, and it is only in the last five years that techniques from complex analysis and the theory of quasicon-formal mappings have been used to prove the result.)

Numerical studies of the logistic map led to several remarkable observations, including Feigenbaum's discovery of the "metric universality" of period-doubling sequences of bifurcations. I was most puzzled that many values of r produced trajectories that did not settle down to periodic cycles. This seemed to be at odds with the results on periodic orbits from my thesis, but it is not. In 1981, Jakobson demonstrated that the set of parameter values for which there is no stable periodic orbit has positive Lebesgue measure. This reveals a paradox that affects the way in which we use the mathematical theory to interpret the results of numerical computations.

In dynamical systems theory as then developed by the "Smale school" and in Russia, "genericity" was defined in terms of Baire category, not in terms of measure. This leads to a difficult question about computation: Given a set C of real numbers that has positive measure but is nowhere dense, will computations tend to reflect properties of C or not? When we attempt to pick numbers at random, e.g., with a pseudo-random number generator, *experience* has shown that numerical computations with dynamical systems tend to reflect the mathematics of measure more than Baire category. This is not a matter of philosophy or mathematics—it is an experimental fact about pseudo-random number generators and systems of computer arithmetic obtained in the face of our ignorance about how round-off errors really affect long iterative calculations.

Example 2: Qualitative Phenomena Can Be Small

In this example, theory pointed the way to dynamical features that were hard to find. The Ruelle–Takens–Newhouse theory of the transition to turbulence predicted that weakly nonlinear systems of three coupled oscillators would exhibit chaotic behavior. In a numerical study based on simulation of a system with randomly chosen parameters and initial conditions, Celso Grebogi at al. found chaotic behavior only when the coupling between oscillators was very strong. In another study of mode locking in models of weakly coupled oscillators, Swan Kim, Claude Baesens, Robert MacKay, and I sought to understand the bifurcations that occur in these systems.

Bifurcation theory predicted where chaotic behavior would be found. Searching systematically, we found the predicted regions. They were tiny: a strip of width 10^{-10} in a parameter space that was the unit square. The transversal intersections of stable and unstable manifolds characteristic of chaotic behavior were visible only in regions of about 10^{-6} in diameter within the phase space. The theory, although indeed correct, was quantitatively insignificant for many purposes.

■ ■ ■

These first two examples involve chaotic dynamics. The central issues in many applications of dynamical systems theory, however, do not focus on chaos. Sensitivity to initial conditions limits the ability of simulations of chaotic systems to make long-term predictions, but the challenge of modeling complex systems for even short times tends to preclude direct comparisons between long time series of data and simulations of chaotic systems. We are more concerned with predicting the type of attractors a system will have than we are with describing details of chaotic motions over long times. We are especially interested in identifying bifurcations where one type of attractor is replaced by another. In the simplest case, we want to determine do-mains in a parameter space where an equilibrium point is stable. Our ability to compute bifurcations varies greatly from example to example. Even in seemingly simple systems, these problems can be surprisingly hard. That has been an obsta-cle to even more widespread use of dynamical systems theory in applied settings.

Example 3: Data Fitting Qualitative Dynamics

For the past several years, I have been engaged with Ronald Harris-Warrick and his laboratory in a stimulating investigation of models of a small neural network, the stomatogastric ganglion. Neuronal models with elements describing individual electrical currents were first constructed by Hodgkin and Huxley about fifty years ago. In our work on the stomatogastric ganglion of lobsters, the models of single neurons have at least twenty-five parameters. The dynamical properties of the system change as the parameters are varied. Stable oscillations, stable equilibria, and chaotic attractors all occur in these models in overlapping parameter regimes. The behaviors of individual neurons and the network are *modulated*, reflecting different rhythmic motions of the lobster foregut.

We endeavor to relate the models directly to biological issues through data. In particular, we want to obtain maps of the parameter spaces that accurately match model behaviors with those observed in the network and parts of the network when manipulated in

the laboratory. Some of the system parameters can be measured, but many cannot. Thus, we are left with the problem of matching time series data with simulations of an imperfect model. We have had bits of success in this enterprise, obtained through luck and extensive interactive investigation of the models. Still, there are enough unknown parameters (and unknown functions describing components of the model) that we have little understanding of the optimal fit between these models and the data. Making the models more detailed to produce better fidelity is problematic. It adds additional parameters and further confusion about how to fit models to data with large numbers of unknown parameters. It also makes model building and simulation slower, diverting attention from the weakness of our tools for investigating simpler models thoroughly.

Without theory to guide us, we are left with a misleading picture of the dynamics inherent in our models.

My experiences with these neural models are typical of nonlinear systems simulation. For some systems—such as laminar fluid flows in simple geometries of low aspect ratio and electrical circuits—we have very precise models. This is hardly typical, however. We seldom have the ability to describe all the relevant functional relationships between components of a complex system, whether natural or artificial, or the resources to measure them even when it is theoretically possible to do so. We have to contend with incomplete information in describing the system.

How we deal with partial information in modeling complex systems is the flip side of our fascination with "emergent properties," our quest to understand how a system can be more than the sum of its parts. In many cases, system behavior is much more sensitive to some parameters than to others, enticing us with the prospect of building reduced-order models that retain essential variables and discard inessential ones. Nonetheless, we fail more often than we succeed in fitting the dynamics of models to data. Success is a matter of hard work, sound intuition, and lots of luck.

Example 4: Computing Hopf Bifurcations

Mathematics can contribute to the enterprise of high-fidelity simulation by creating better computer methods for the analysis of dynamical systems. This example illustrates the direct computation of bifurcations. Hopf bifurcation occurs in a family of vector fields when the stability of an equilibrium is changed by a pair of eigenvalues crossing the imaginary axis. The set of matrices with a pair of pure imaginary eigenvalues is a "semi-algebraic" set defined by a single polynomial equation and a single inequality in the coefficients of the matrix. This equation becomes complicated as the size of the matrix grows, and it is ill-conditioned for large sets of matrices. Nonetheless, computing Hopf bifurcations means that this equation must be solved in some form. The problem presents numerical challenges as the size of a system grows.

Varied approaches have been taken to the computation of Hopf bifurcations. Mark Myers, Bernd Sturmfels, and I formulated the problem in terms of determining the singularity of a matrix, using classical algebraic constructions that had been largely forgotten. The most elegant version of the method is expressed in terms of tensor products, producing a matrix, called the biproduct, whose eigenvalues are the sums of distinct eigenvalues of the Jacobian. If the Jacobian has pure imaginary eigenvalues, then its biproduct is singular. If the Jacobian is block diagonal, then so is the biproduct.

The smallest singular value of a matrix is a measure of the distance of the matrix from being singular. Computing the smallest singular value of a matrix is a problem that has been intensively studied, and we used the methods created to solve it as part of our strategy for computing Hopf bifurcations. We implemented algorithms for computing Hopf bifurcations based on the calculation of biproducts and, in an alternative approach, resultants of the characteristic polynomial of Df. The implementation with biproducts incorporated subspace iteration to identify low-dimensional invariant subspaces containing the eigen-spaces of Hopf bifurcations. We tested our algorithms on neural models for which which it was also possible to symbolically compute the location of Hopf bifurcations. We also made comparisons with several other methods. In one of our tests, we examined a parameter region close to a point of double Hopf bifurcation, where two pairs of pure imaginary eigenvalues occur simultaneously. There are no clear-cut measures of the effectiveness of the different algorithms, but our methods appeared to be more robust in converging to the solutions closest to the starting values.

Example 5: Computing Periodic Orbits

The final problem I discuss comes from my current work on computing periodic orbits of dynamical systems. This story has many interesting aspects and is far from complete. Periodic orbits are fundamental structures, trajectories of vector fields distinguished by their topology.

In many biological examples, stable periodic behaviors are the desired states of the system: Animal locomotion (walking, flying, and swimming), blood circulation, and circadian rhythms are all fundamental biological processes that are periodic. We design machines to avoid periodic motions because limit-cycle oscillations can cause damaging vibrations, leading to fatigue and fracture of materials. We need to determine the amplitude, stability, and other characteristics of periodic orbits in dynamical systems if we are to operate in periodic regimes. The flexibility and dexterity of biological systems are due in part to their superior ability to exploit periodic motion. For example, legged locomotion has evident advantages over wheeled locomotion in rugged terrain, but we are not yet able to build legged vehicles that have the stability or agility of animals.

The analytic determination of periodic orbits of planar vector fields is one of the classic "grand challenges" of mathematics, included in Hilbert's famous list of problems a century ago. We don't even have sound estimates for the upper bound on the number of limit cycles in polynomial vector fields of degree larger than two. The quadratic vector fields have five independent parameters, after allowances for linear coordinate changes and time rescaling. The cubic vector fields have 13, already a formidable number

for comprehensive numerical studies.

When I got my first computer (a Sun 1 workstation that computed at a few kiloflops), Gerhard Dangelmayr and I, motivated by questions of the effects of imperfect symmetry on bifurcations, undertook the study of a cubic system with four parameters:

$$\begin{aligned} \dot{x} &= y \\ \dot{y} &= -(x^3 + r\,x^2 + n\,x + m) + y\,(b - x^2). \end{aligned}$$

With perturbation methods and symbolic computation, we showed that there are parameters for which this system has four limit cycles surrounding a single equilibrium point. At that time, we were unsuccessful in displaying these cycles with a numerical computation. Several years later, Salvador Malo was able to do so. He estimated that, with fixed r and m, the region of the (b,n) plane with the desired four limit cycles was a strip of width about 3×10^{-9}. As with the example of chaotic behavior in coupled oscillators, the mathematics leads to phenomena that occur in exceedingly small regions of a parameter space. This example is hardly special. The simplest model of oscillatory behavior in chemical reactors, the so-called CSTR, is a planar vector field whose qualitative dynamics are just as challenging to determine.

Here is yet one more example of periodic orbits that are challenging to compute: The vector field

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

gives rise to periodic orbits that have been termed "canards" because of their vaguely duck-like shapes. We study this equation when $\varepsilon > 0$ is small, say 10^{-3}. As a decreases through 0, a Hopf bifurcation occurs. The periodic orbits created in this Hopf bifurcation are stable, and they grow quickly into relaxation oscillations as a decreases.

The intermediate orbits have segments that lie close to the decreasing portion of the characteristic curve $y = x^2 + x^3$. Here, the flow is violently unstable. The divergence of trajectories is so rapid that changes in initial conditions of unit precision are amplified to unit magnitude long before the trajectory has climbed as far along the characteristic curve as the periodic orbit has. Numerically, we observe a jump from small limit cycles to large ones, perhaps with a narrow zone in which the numerical trajectories are chaotic. Qualitatively, this is all wrong because planar vector fields do not have chaotic trajectories. Notice that a numerical integration that is exact apart from round-off errors will still be unable to track the canard solutions of this vector field.

This is thus an example of a vector field for which, without extreme numerical precision, much greater than standard IEEE floating-point arithmetic, a whole set of attractors is inaccessible to numerical integration. Simulation gives a very different picture from the mathematical analysis, producing results that the mathematics proves to be impossible in the system we are trying to simulate. This small system, with only mild stiffness compared with the equations of many chemical reactions, takes us rapidly to places where simulation is unable to reproduce the behavior of the system, quantitatively or qualitatively.

It is reasonable to ask how important these phenomena are for simulation of the real world. Let me give you my opinion. There is a dichotomy, drawn beautifully by Peter Henrici, between mathematics done for the sake of achieving a deeper understanding of the world and mathematics done to act on the world. Clearly, the issues I have raised in this example need to be resolved if we are to achieve a deep understanding of dynamical systems and their bifurcations. I believe that they are also important as we bring our theoretical insight to bear on specific problems through computation and simulation. Canards, for example, are common in systems with multiple time scales as they undergo bifurcation.

Without theory to guide us, we are left with a misleading picture of the dynamics inherent in our models. Because there are many sources of uncertainty in our models, the challenge of producing high levels of fidelity in our simulations is enormous. To meet that challenge, we need conceptual frameworks that facilitate deep understanding of how the components of complex systems interact. Part of the richness of dynamical systems theory is that it does precisely that. To develop powerful computational tools that help analyze how system behavior depends on parameters and system "architecture," we must confront subtle issues in the mathematics. These issues are already apparent in low-dimensional settings and are most easily studied there. We can simulate complex systems and learn interesting things as we do, but we must expect the same issues—and others—to impede our progress in producing high-fidelity simulations. Let me reiterate that this is especially true when our models are inaccurate and when we are unable to measure relevant parameters for models. In short, this is true of almost every complex system we seek to understand. The severity of this problem is comparable to that of winning a lottery by guessing correctly, say, a 20-digit number. It's possible, but hardly likely, even with millions of attempts. Without a strategy to guide us in searching parameter spaces, that is the type of challenge we face in fitting data to models with unkown parameters. Surely, we need to do better than blind or random search.

Let us return to what *can* be done to compute periodic orbits of vector fields. Two types of strategies are commonly used. First, numerical integration of trajectories does converge to periodic orbits in many cases. Second, the computation of periodic orbits can be formulated as a boundary value problem with periodic boundary conditions. There is one dominant code for finding periodic orbits with "global" boundary value methods—AUTO, written and maintained by Sebius Doedel. In principle, it is possible to use two-point boundary value codes to solve problems with periodic boundary conditions, but my attempts to do so have been very disappointing. (Explaining why they fail is a nice research problem.) For years, driven by my desire to deal with examples like those we have just seen, I have wanted to establish methods that would be more robust and more accurate than AUTO. Two of the difficulties with AUTO are that it requires very good approximations to a periodic orbit as initial data, and that it has many algorithmic parameters that affect convergence in ways I find hard to understand. I set out a couple of years ago to do something

better, and I want to describe progress I have made.

Boundary value methods for computing periodic orbits formulate discretized versions of the differential equations and boundary conditions satisfied by the periodic orbits. These discretized equations are then solved with standard methods, typically Newton iteration. There is a tension between the size of the system of discretized equations and the accuracy with which the periodic orbit is represented. The minimal size of the system to be solved is given approximately by the product of the number of mesh points and the dimension of the phase space. With collocation methods, each mesh interval contains collocation points that add yet more equations. Finer meshes yield larger systems of equations that are time-consuming to solve and require good initial data for convergence. My objective has been to create methods that will achieve good accuracy with coarse meshes. Working with Won Gyu Choe and Brian Meloon, I have approached this goal by making use of a computer method called *automatic differentiation*.

We use automatic differentiation to compute the Taylor series of trajectories; specifically, we use the package ADOL-C, developed at Argonne National Laboratory by Andreas Griewank et al. to compute Taylor series of trajectories and their derivatives with respect to phase space variables. The Taylor series are used to construct highly accurate approximations to trajectories and periodic orbits of vector fields. The limit of increasing degree on a fixed mesh corresponds to the use of coarse meshes to compute periodic orbits in boundary value solvers. Approximate periodic orbits of increasing accuracy are parameterized by the same data, namely the coefficients of mesh points. We have been experimenting with implementations of both shooting and global methods based on Taylor series.

I find the results amazing. In test examples, we are able to compute periodic orbits on coarse meshes more accurate than those produced by fourth-order Runge–Kutta integration *at any step sizes*. The precision of the results consistently exceeds that of AUTO computations by orders of magnitude, and always with coarser meshes. Our test examples include the canard example described earlier, showing that the methods work with systems that have multiple time scales.

Final Remarks

Dynamical systems theory gives us a framework for investigating the process of simulation. As we build algorithms and computational tools on these mathematical foundations, we repeatedly encounter new challenges. In some cases, numerical investigations lead us to discover new phenomena and deepen our theoretical understanding. In other situations, the theory guides us to phenomena that we would overlook without its help. In still others, we encounter limits to straightforward simulation. In all these cases, I take delight in the interplay between mathematics, computation, and science.

Let me come back for a few moments to the realm of scientific politics and policy. We ask ourselves how much we can rely on the results we obtain from simulations. The term *uncertainty* has been adopted to describe technical aspects of this important question. Quantifying uncertainty in the simulation of complex systems is an emerging re-search theme that cuts across all disciplines. It behooves SIAM to articulate the common issues that arise when we study different systems and to propose strategies for seeking common solutions. Many areas of mathematics have tools to contribute to this enterprise, but we have hardly begun to develop comprehensive theories or needed tools. We need to remind ourselves and others that mathematics remains an effective intellectual viewpoint for tackling interdisciplinary problems and will be an essential part of this effort. Moreover, the mathematics found in these problems is both fascinating and deep.