Newsjournal of the Society for Industrial and Applied Mathematics

siam news

Volume 39/ Number 1
January/February 2006

A Moveable Feast: Researchers Seek Stability in Lability

By Barry A. Cipra

Nothing endures but change.
—Heraclitus

Eigenvalues. Fixed points. Stable equilibria. Mathematicians like things that stay put. And if they can't stay put, the objects of study should at least repeat themselves on a regular basis, like orbiting planets or populations of predators and prey. Even in the case of chaotic systems, mathematicians have traditionally gravitated toward invariant features, such as strange attractors, stable manifolds, and periodic points.

What makes this tradition possible is that dynamical systems—at least the ones mathematicians favor—are governed by equations that depend on time either cyclically or not at all. But nature doesn't always oblige. Many phenomena require equations whose coefficients are non-periodic functions of time. Indeed, many—arguably most—phenomena can be described not by equations at all, but only as an amalgam of time-varying data.

Are such dynamical systems beyond the reach of analysis? Hardly. Applied mathematicians are developing new tools for the study of time-dependent, data-driven dynamical systems. In the process, they are stretching and bending some of the traditional concepts of dynamical systems—but with an eye on retaining their invariant core.

In a wide-ranging John von Neumann lecture on geometric mechanics and computational dynamical systems at the 2005 SIAM Annual Meeting, Jerry Marsden of Caltech highlighted one of the tools currently taking shape: Lagrangian coherent structures. Introduced in 2000 by George Haller, now at MIT, in a paper with Guo-Cheng Yuan of Brown University, and elaborated by Haller in a series of subsequent papers, Lagrangian coherent structures enable researchers to spot non-obvious boundaries in complicated flows. These quasi-invariant objects have been studied in a variety of computational dynamical systems and applications, including a model of the spread of pollution by ocean currents off the coast of Florida, and even a model of the biochemical process of apoptosis, i.e., cell death.

Timing Is Everything

Roughly speaking, a Lagrangian coherent structure (LCS) is a mobile separatrix with, at worst, a slow leak. More precisely, an LCS is a material line (or surface) that remains hyperbolic for a locally maximal amount of time. An LCS can be identified, for example, as a "ridge" in a scalar field of "direct," finite-time Lyapunov exponents associated with the dynamical system. Indeed, following earlier work by Haller and co-workers, Marsden, François Lekien of Princeton, and Shawn Shadden of Caltech have proposed ridges as the definition of a Lagrangian coherent structure.

Making sense of these descriptions calls for some unpacking of terminology.

One of the basic steps in analyzing a dynamical system is to identify regions of qualitatively different dynamics and find the boundaries between them. That's what a separatrix does for "autonomous" systems—i.e., systems whose equations don't depend on time. The classic example is the mathematical pendulum, defined by the (normalized) equations $dx/dt = y$ and $dy/dt = -\sin x$. The separatrix is the curve separating normal, back-and-forth oscillation and

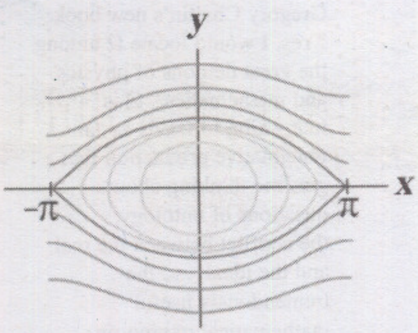

Figure 1. *Eye on the prize. The phase diagram for the simple pendulum features a separatrix. Figure adapted from Shawn Shadden, www.cds.caltech.edu/~shawn/LCS-tutorial.*

high-speed clockwise or counterclockwise spinning (see Figure 1). Because there's no time dependence in the equations of the pendulum, this boundary doesn't move.

For non-autonomous systems, all bets are seemingly off. But here too, regions with qualitative differences exist, at least for a while. It's just that the boundaries between them tend to wander and, in some cases, disappear. Their waywardness suggests a Lagrangian, as opposed to Eulerian, approach to the analysis. (In fluid dynamics, a Lagrangian approach follows particle trajectories, whereas the Eulerian viewpoint sticks to a single, fixed frame of reference.) Their tendency to disappear suggests abandoning, or at least modifying, any analytic tool based on asymptotic limits in time.

Enter finite-time Lyapunov exponents. The traditional Lyapunov exponent is an asymptotic object; roughly speaking, it tracks the extent to which infinitely close particles separate in an infinite amount of time. For autonomous systems, this has immediate, short-term significance. For non-autonomous, not to mention data-

See **Stability** *on page 12*

Stability

continued from page 1

defined, systems, it's meaningless. But it's still possible to measure the change in separation over a finite time interval. The formal definition of a finite-time Lyapunov exponent (FTLE) is fairly technical, but the upshot is the assignment of a number to each point (x,y) that measures how strongly the trajectory starting there at time t will separate from nearby trajectories by time $t + T$. The definition of a ridge is also fairly technical, but the term itself offers an intuitive explanation: A ridge is a path in the Lyapunov landscape that, while it may (and usually does) go up and down in its tangent direction, definitely drops off steeply on either side.

Finite-time Lyapunov exponents are also called direct Lyapunov exponents (DLEs), because they can be determined directly from particle trajectories. (Marsden prefers the acronym FTLE, in part, he explains, to distinguish FTLEs from finite-space Lyapunov exponents, or FSLEs.) This attribute makes them especially suitable for the computational analysis of real-world data. Lekien, Haller, Marsden, and colleagues Chad Coulliette of Caltech and Arthur Mariano, Edward Ryan, and Lynn Shay of the University of Miami have used them with radar data measuring ocean surface currents along the coast of Florida near Fort Lauderdale.

The analysis revealed an LCS attached to the coast and extending to the southeast (see Figure 2). This LCS separates the Florida Current from a zone of recirculation. Its existence—and especially the fact that it moves—has obvious implications for the fate of any pollutants released in the area. In particular, Haller points out, it matters not only where you dump your effluent, but also when. The ability to predict the motion of the LCS, the researchers have shown, provides the basis for a real-time pollution-control algorithm.

Go with the Flow

An LCS is not, in general, a perfect barrier; over time, particles on one side of an LCS may make it to the other side—a Lagrangian coherent structure, that is, might more properly be described as quasi-Lagrangian. The flux across an LCS is usually very small (or non-existent) in practice, however. Marsden, Lekien, and Shadden quantified this for their finite-time, Lyapunov-exponent-based definition. They showed that, up to an error term that scales as $1/T$, the flux across an LCS depends on the product of two terms: one that measures the sharpness of the ridge and one that represents the difference between the local rotation rate of the LCS and that of the ambient, Eulerian velocity field. In the Florida coastline example, the computed crossing rate was less than two meters per hour, or about 0.05% of the average flow speed in the vicinity of the LCS.

Some of the most exciting applications of Lagrangian coherent structures are in the life sciences. Marsden, Shadden, and John Dabiri, a professor of aeronautics and bioengineering at Caltech, have computed an LCS in the fluid surrounding a free-swimming jellyfish. Superimposed on a video of the jellyfish, the LCS shows how fluid—and nutrients—are entrained within the critter's jellybelly (more properly called the subumbrellar region). The animation can be found at http://www.cds.caltech.edu/~marsden/research/demos/fluidtransport.php (or by googling "shadden jellyfish").

Marsden's group is also working with Charles Taylor of Stanford University on computational studies of cardiovascular

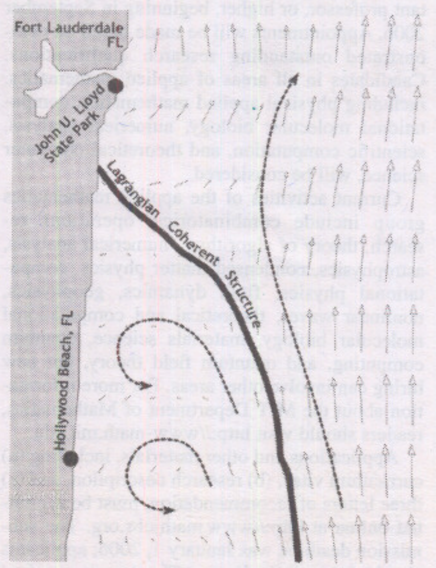

Figure 2. *An LCS off the coast of Florida indicates the boundary between two regions of qualitatively different dynamics. The meandering of the LCS has implications for the fate of pollutants released in the vicinity. Figure courtesy of François Lekien, based on OSCR data collected during the 4D current experiment by Tom Cook, Brian Haus, Arthur Mariano, Jorge Martinez, Lynn Shay, and Ed Ryan at the Rosentiel School of Marine and Atmospheric Science in Miami.*

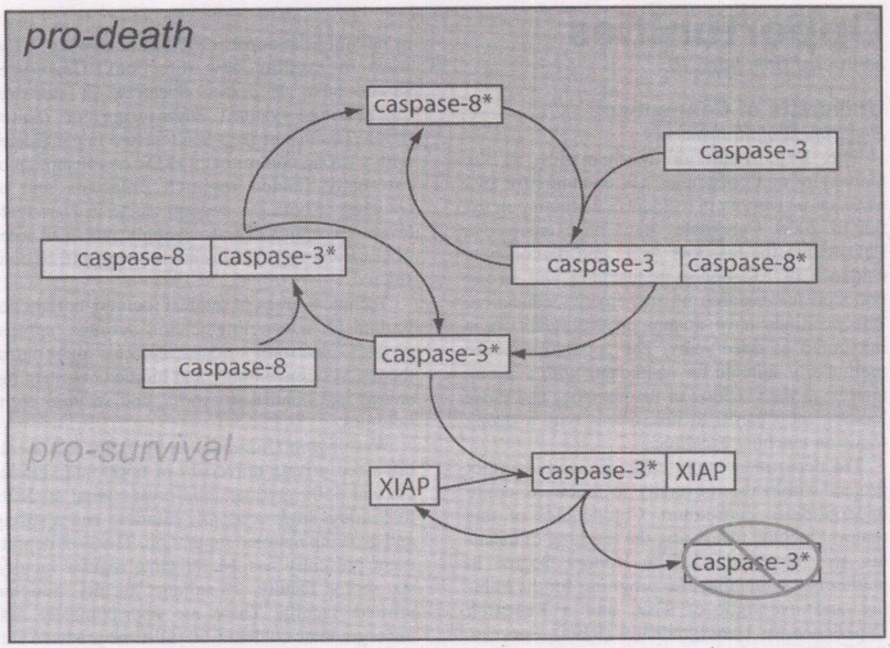

Figure 3. *A schematic of the programmed cell death decision pathway. Figure courtesy of Bree Aldridge, Peter Sorger, and Douglas Lauffenburger.*

flows (see *SIAM News*, October 2005, page 1; http://www.siam.org/news/news.php?id=160). The computation of LCSs from, say, MRI data can show if a zone of recirculation is lingering in one spot—a bad thing, in that recirculation promotes blood clot formation and plaque build-up, also known as hardening of the arteries. (According to Taylor, one of the benefits of exercising is that it breaks up these zones.) In the (non-asymptotic) future, your cardiologist may judge the state of your health by measuring your Lyapunov exponents.

Death, too, is coming into the fold. Haller and colleagues Bree Aldridge, Peter Sorger, and Douglas Lauffenburger, all of the biological engineering division at MIT, recently used Lyapunov exponents to analyze a mathematical model of transient signaling in a protein network involved in apoptosis. The network consists of eight forms or combinations of three proteins, two (caspase-3 and caspase-8) that promote cell death and one (XIAP) that inhibits it (see Figure 3). They were able to find a DLE-defined LCS that separates apoptosis from survival. With functional genomics and proteomics shouldering more and more of the burden of biomedicine, the non-steady-state analysis of cellular networks will likely loom large. As John Maynard Keynes said in the *econ-* branch of *-omics*, "In the long run, we are all dead."

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.